



OpenShift Container Platform 4.9

Logging

OpenShift Logging installation, usage, and release notes

OpenShift Container Platform 4.9 Logging

OpenShift Logging installation, usage, and release notes

Legal Notice

Copyright © 2023 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[®] is the registered trademark of Linus Torvalds in the United States and other countries.

Java[®] is a registered trademark of Oracle and/or its affiliates.

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

MySQL[®] is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js[®] 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[®] 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 instructions for installing, configuring, and using OpenShift Logging, which aggregates logs for a range of OpenShift Container Platform services.

Table of Contents

CHAPTER 1. RELEASE NOTES FOR LOGGING	10
1.1. LOGGING 5.5.8	10
1.1.1. Bug fixes	10
1.1.2. CVEs	10
1.2. LOGGING 5.5.7	10
1.2.1. Bug fixes	10
1.2.2. CVEs	11
1.3. LOGGING 5.5.6	11
1.3.1. Known issues	11
1.3.2. Bug fixes	11
1.3.3. CVEs	12
1.4. LOGGING 5.5.5	12
1.4.1. Bug fixes	12
1.4.2. CVEs	13
1.5. LOGGING 5.5.4	15
1.5.1. Bug fixes	16
1.5.2. CVEs	16
1.6. LOGGING 5.5.3	17
1.6.1. Bug fixes	17
1.6.2. CVEs	17
1.7. LOGGING 5.5.2	18
1.7.1. Bug fixes	18
1.7.2. CVEs	19
1.8. LOGGING 5.5.1	19
1.8.1. Enhancements	19
1.8.2. Bug fixes	19
1.8.3. CVEs	20
1.9. LOGGING 5.5	20
1.9.1. Enhancements	20
1.9.2. Bug fixes	21
1.9.3. CVEs	21
1.10. LOGGING 5.4.14	22
1.10.1. Bug fixes	22
1.10.2. CVEs	22
1.11. LOGGING 5.4.13	22
1.11.1. Bug fixes	22
1.11.2. CVEs	22
1.12. LOGGING 5.4.12	22
1.12.1. Bug fixes	22
1.12.2. CVEs	23
1.13. LOGGING 5.4.11	23
1.13.1. Bug fixes	23
1.13.2. CVEs	23
1.14. LOGGING 5.4.10	23
1.14.1. Bug fixes	24
1.14.2. CVEs	24
1.15. LOGGING 5.4.9	24
1.15.1. Bug fixes	24
1.15.2. CVEs	25
1.16. LOGGING 5.4.8	27
1.16.1. Bug fixes	27

1.16.2. CVEs	27
1.17. LOGGING 5.4.6	28
1.17.1. Bug fixes	28
1.17.2. CVEs	28
1.18. LOGGING 5.4.5	29
1.18.1. Bug fixes	29
1.18.2. CVEs	29
1.19. LOGGING 5.4.4	29
1.19.1. Bug fixes	30
1.19.2. CVEs	30
1.20. LOGGING 5.4.3	30
1.20.1. Elasticsearch Operator deprecation notice	30
1.20.2. Bug fixes	30
1.20.3. CVEs	31
1.21. LOGGING 5.4.2	31
1.21.1. Bug fixes	31
1.21.2. CVEs	32
1.22. LOGGING 5.4.1	33
1.22.1. Bug fixes	33
1.22.2. CVEs	34
1.23. LOGGING 5.4	34
1.23.1. Technology Previews	35
1.23.2. About Vector	35
1.23.2.1. Enabling Vector	35
1.23.3. About Loki	36
1.23.3.1. Deploying the Lokistack	36
1.23.4. Bug fixes	37
1.23.5. CVEs	38
1.24. LOGGING 5.3.14	39
1.24.1. Bug fixes	39
1.24.2. CVEs	39
1.25. LOGGING 5.3.13	41
1.25.1. Bug fixes	41
1.25.2. CVEs	42
1.26. LOGGING 5.3.12	42
1.26.1. Bug fixes	42
1.26.2. CVEs	42
1.27. LOGGING 5.3.11	43
1.27.1. Bug fixes	43
1.27.2. CVEs	43
1.28. LOGGING 5.3.10	43
1.28.1. Bug fixes	43
1.28.2. CVEs	43
1.29. LOGGING 5.3.9	44
1.29.1. Bug fixes	44
1.29.2. CVEs	44
1.30. LOGGING 5.3.8	45
1.30.1. Bug fixes	45
1.30.2. CVEs	45
1.31. OPENSIFT LOGGING 5.3.7	47
1.31.1. Bug fixes	47
1.31.2. CVEs	47
1.32. OPENSIFT LOGGING 5.3.6	48

1.32.1. Bug fixes	48
1.33. OPENSIFT LOGGING 5.3.5	48
1.33.1. Bug fixes	48
1.33.2. CVEs	48
1.34. OPENSIFT LOGGING 5.3.4	49
1.34.1. Bug fixes	49
1.34.2. CVEs	49
1.35. OPENSIFT LOGGING 5.3.3	50
1.35.1. Bug fixes	50
1.35.2. CVEs	50
1.36. OPENSIFT LOGGING 5.3.2	50
1.36.1. Bug fixes	50
1.36.2. CVEs	51
1.37. OPENSIFT LOGGING 5.3.1	51
1.37.1. Bug fixes	51
1.37.2. CVEs	51
1.38. OPENSIFT LOGGING 5.3.0	54
1.38.1. New features and enhancements	54
1.38.2. Bug fixes	54
1.38.3. Known issues	54
1.38.4. Deprecated and removed features	55
1.38.4.1. Forwarding logs using the legacy Fluentd and legacy syslog methods have been removed	55
1.38.4.2. Configuration mechanisms for legacy forwarding methods have been removed	55
1.38.5. CVEs	55
1.39. LOGGING 5.2.13	59
1.39.1. Bug fixes	59
1.39.2. CVEs	59
1.40. LOGGING 5.2.12	60
1.40.1. Bug fixes	60
1.40.2. CVEs	60
1.41. LOGGING 5.2.11	60
1.41.1. Bug fixes	60
1.41.2. CVEs	61
1.42. OPENSIFT LOGGING 5.2.10	62
1.42.1. Bug fixes	62
1.42.2. CVEs	63
1.43. OPENSIFT LOGGING 5.2.9	63
1.43.1. Bug fixes	63
1.44. OPENSIFT LOGGING 5.2.8	64
1.44.1. Bug fixes	64
1.44.2. CVEs	64
1.45. OPENSIFT LOGGING 5.2.7	64
1.45.1. Bug fixes	64
1.45.2. CVEs	64
1.46. OPENSIFT LOGGING 5.2.6	65
1.46.1. Bug fixes	65
1.46.2. CVEs	65
1.47. OPENSIFT LOGGING 5.2.5	65
1.47.1. Bug fixes	65
1.47.2. CVEs	65
1.48. OPENSIFT LOGGING 5.2.4	66
1.48.1. Bug fixes	66
1.48.2. CVEs	66

1.49. OPENSIFT LOGGING 5.2.3	68
1.49.1. Bug fixes	68
1.49.2. CVEs	68
1.50. OPENSIFT LOGGING 5.2.2	70
1.50.1. Bug fixes	70
1.50.2. CVEs	70
1.51. OPENSIFT LOGGING 5.2.1	71
1.51.1. Bug fixes	71
1.51.2. CVEs	71
1.52. OPENSIFT LOGGING 5.2.0	71
1.52.1. New features and enhancements	71
1.52.2. Bug fixes	72
1.52.3. Known issues	74
1.52.4. Deprecated and removed features	74
1.52.5. Forwarding logs using the legacy Fluentd and legacy syslog methods have been deprecated	74
1.52.6. CVEs	74
CHAPTER 2. UNDERSTANDING THE LOGGING SUBSYSTEM FOR RED HAT OPENSIFT	76
2.1. GLOSSARY OF COMMON TERMS FOR OPENSIFT CONTAINER PLATFORM LOGGING	76
2.2. ABOUT DEPLOYING THE LOGGING SUBSYSTEM FOR RED HAT OPENSIFT	78
2.2.1. About JSON OpenShift Container Platform Logging	78
2.2.2. About collecting and storing Kubernetes events	78
2.2.3. About updating OpenShift Container Platform Logging	79
2.2.4. About viewing the cluster dashboard	79
2.2.5. About troubleshooting OpenShift Container Platform Logging	79
2.2.6. About uninstalling OpenShift Container Platform Logging	79
2.2.7. About exporting fields	79
2.2.8. About logging subsystem components	79
2.2.9. About the logging collector	80
2.2.10. About the log store	80
2.2.11. About logging visualization	81
2.2.12. About event routing	81
2.2.13. About log forwarding	82
CHAPTER 3. INSTALLING THE LOGGING SUBSYSTEM FOR RED HAT OPENSIFT	83
3.1. INSTALLING THE LOGGING SUBSYSTEM FOR RED HAT OPENSIFT USING THE WEB CONSOLE	83
3.2. POST-INSTALLATION TASKS	88
3.3. INSTALLING THE LOGGING SUBSYSTEM FOR RED HAT OPENSIFT USING THE CLI	88
3.4. POST-INSTALLATION TASKS	96
3.4.1. Defining Kibana index patterns	96
3.4.2. Allowing traffic between projects when network isolation is enabled	97
CHAPTER 4. CONFIGURING YOUR LOGGING DEPLOYMENT	99
4.1. ABOUT THE CLUSTER LOGGING CUSTOM RESOURCE	99
4.1.1. About the ClusterLogging custom resource	99
4.2. CONFIGURING THE LOGGING COLLECTOR	100
4.2.1. About unsupported configurations	100
4.2.2. Viewing logging collector pods	101
4.2.3. Configure log collector CPU and memory limits	101
4.2.4. Advanced configuration for the log forwarder	102
4.2.5. Removing unused components if you do not use the default Elasticsearch log store	106
4.3. CONFIGURING THE LOG STORE	107
4.3.1. Forwarding audit logs to the log store	107
4.3.2. Configuring log retention time	109

4.3.3. Configuring CPU and memory requests for the log store	111
4.3.4. Configuring replication policy for the log store	113
4.3.5. Scaling down Elasticsearch pods	114
4.3.6. Configuring persistent storage for the log store	114
4.3.7. Configuring the log store for emptyDir storage	115
4.3.8. Performing an Elasticsearch rolling cluster restart	116
4.3.9. Exposing the log store service as a route	119
4.4. CONFIGURING THE LOG VISUALIZER	121
4.4.1. Configuring CPU and memory limits	122
4.4.2. Scaling redundancy for the log visualizer nodes	123
4.5. CONFIGURING LOGGING SUBSYSTEM STORAGE	123
4.5.1. Storage considerations for the logging subsystem for Red Hat OpenShift	123
4.5.2. Additional resources	124
4.6. CONFIGURING CPU AND MEMORY LIMITS FOR LOGGING SUBSYSTEM COMPONENTS	124
4.6.1. Configuring CPU and memory limits	124
4.7. USING TOLERATIONS TO CONTROL OPENSIFT LOGGING POD PLACEMENT	126
4.7.1. Using tolerations to control the log store pod placement	127
4.7.2. Using tolerations to control the log visualizer pod placement	128
4.7.3. Using tolerations to control the log collector pod placement	129
4.7.4. Additional resources	130
4.8. MOVING LOGGING SUBSYSTEM RESOURCES WITH NODE SELECTORS	130
4.8.1. Moving OpenShift Logging resources	130
4.9. CONFIGURING SYSTEMD-JOURNALD AND FLUENTD	134
4.9.1. Configuring systemd-journald for OpenShift Logging	134
4.10. MAINTENANCE AND SUPPORT	137
4.10.1. About unsupported configurations	137
4.10.2. Unsupported configurations	137
4.10.3. Support policy for unmanaged Operators	138
CHAPTER 5. VIEWING LOGS FOR A RESOURCE	140
5.1. VIEWING RESOURCE LOGS	140
CHAPTER 6. VIEWING CLUSTER LOGS BY USING KIBANA	142
6.1. DEFINING KIBANA INDEX PATTERNS	142
6.2. VIEWING CLUSTER LOGS IN KIBANA	143
CHAPTER 7. FORWARDING LOGS TO EXTERNAL THIRD-PARTY LOGGING SYSTEMS	146
7.1. ABOUT FORWARDING LOGS TO THIRD-PARTY SYSTEMS	146
Fluentd log handling when the external log aggregator is unavailable	150
Supported Authorization Keys	150
7.1.1. Creating a Secret	151
7.2. SUPPORTED LOG DATA OUTPUT TYPES IN OPENSIFT LOGGING 5.1	151
7.3. SUPPORTED LOG DATA OUTPUT TYPES IN OPENSIFT LOGGING 5.2	152
7.4. SUPPORTED LOG DATA OUTPUT TYPES IN OPENSIFT LOGGING 5.3	152
7.5. SUPPORTED LOG DATA OUTPUT TYPES IN OPENSIFT LOGGING 5.4	153
7.6. SUPPORTED LOG DATA OUTPUT TYPES IN OPENSIFT LOGGING 5.5	154
7.7. SUPPORTED LOG DATA OUTPUT TYPES IN OPENSIFT LOGGING 5.6	154
7.8. FORWARDING LOGS TO AN EXTERNAL ELASTICSEARCH INSTANCE	155
7.9. FORWARDING LOGS USING THE FLUENTD FORWARD PROTOCOL	158
7.9.1. Enabling nanosecond precision for Logstash to ingest data from fluentd	160
7.10. FORWARDING LOGS USING THE SYSLOG PROTOCOL	160
7.10.1. Adding log source information to message output	163
7.10.2. Syslog parameters	163
7.10.3. Additional RFC5424 syslog parameters	165

7.11. FORWARDING LOGS TO AMAZON CLOUDWATCH	165
7.12. FORWARDING LOGS TO LOKI	170
7.12.1. Troubleshooting Loki "entry out of order" errors	172
7.13. FORWARDING APPLICATION LOGS FROM SPECIFIC PROJECTS	175
7.14. FORWARDING APPLICATION LOGS FROM SPECIFIC PODS	177
7.15. TROUBLESHOOTING LOG FORWARDING	179
CHAPTER 8. ENABLING JSON LOGGING	180
8.1. PARSING JSON LOGS	180
8.2. CONFIGURING JSON LOG DATA FOR ELASTICSEARCH	181
8.3. FORWARDING JSON LOGS TO THE ELASTICSEARCH LOG STORE	183
CHAPTER 9. COLLECTING AND STORING KUBERNETES EVENTS	185
9.1. DEPLOYING AND CONFIGURING THE EVENT ROUTER	185
CHAPTER 10. UPDATING OPENSIFT LOGGING	189
10.1. SUPPORTED VERSIONS	189
10.2. UPDATING LOGGING TO THE CURRENT VERSION	189
CHAPTER 11. VIEWING CLUSTER DASHBOARDS	194
11.1. ACCESSING THE ELASTICSEARCH AND OPENSIFT LOGGING DASHBOARDS	194
11.2. ABOUT THE OPENSIFT LOGGING DASHBOARD	194
11.3. CHARTS ON THE LOGGING/ELASTICSEARCH NODES DASHBOARD	196
CHAPTER 12. TROUBLESHOOTING LOGGING	202
12.1. VIEWING OPENSIFT LOGGING STATUS	202
12.1.1. Viewing the status of the Red Hat OpenShift Logging Operator	202
12.1.1.1. Example condition messages	204
12.1.2. Viewing the status of logging subsystem components	206
12.2. VIEWING THE STATUS OF THE ELASTICSEARCH LOG STORE	207
12.2.1. Viewing the status of the log store	207
12.2.1.1. Example condition messages	209
12.2.2. Viewing the status of the log store components	211
12.2.3. Elasticsearch cluster status	215
12.3. UNDERSTANDING LOGGING SUBSYSTEM ALERTS	216
12.3.1. Viewing logging collector alerts	216
12.3.2. About logging collector alerts	216
12.3.3. About Elasticsearch alerting rules	217
12.4. COLLECTING LOGGING DATA FOR RED HAT SUPPORT	218
12.4.1. About the must-gather tool	218
12.4.2. Prerequisites	219
12.4.3. Collecting OpenShift Logging data	219
12.5. TROUBLESHOOTING FOR CRITICAL ALERTS	219
12.5.1. Elasticsearch Cluster Health is Red	219
12.5.2. Elasticsearch Cluster Health is Yellow	222
12.5.3. Elasticsearch Node Disk Low Watermark Reached	222
12.5.4. Elasticsearch Node Disk High Watermark Reached	223
12.5.5. Elasticsearch Node Disk Flood Watermark Reached	224
12.5.6. Elasticsearch JVM Heap Use is High	226
12.5.7. Aggregated Logging System CPU is High	226
12.5.8. Elasticsearch Process CPU is High	226
12.5.9. Elasticsearch Disk Space is Running Low	226
12.5.10. Elasticsearch FileDescriptor Usage is high	227
CHAPTER 13. UNINSTALLING OPENSIFT LOGGING	228

13.1. UNINSTALLING THE LOGGING SUBSYSTEM FOR RED HAT OPENSIFT	228
CHAPTER 14. LOG RECORD FIELDS	231
CHAPTER 15. MESSAGE	232
CHAPTER 16. STRUCTURED	233
CHAPTER 17. @TIMESTAMP	234
CHAPTER 18. HOSTNAME	235
CHAPTER 19. IPADDR4	236
CHAPTER 20. IPADDR6	237
CHAPTER 21. LEVEL	238
CHAPTER 22. PID	239
CHAPTER 23. SERVICE	240
CHAPTER 24. TAGS	241
CHAPTER 25. FILE	242
CHAPTER 26. OFFSET	243
CHAPTER 27. KUBERNETES	244
27.1. KUBERNETES.POD_NAME	244
27.2. KUBERNETES.POD_ID	244
27.3. KUBERNETES.NAMESPACE_NAME	244
27.4. KUBERNETES.NAMESPACE_ID	244
27.5. KUBERNETES.HOST	244
27.6. KUBERNETES.CONTAINER_NAME	244
27.7. KUBERNETES.ANNOTATIONS	245
27.8. KUBERNETES.LABELS	245
27.9. KUBERNETES.EVENT	245
27.9.1. kubernetes.event.verb	245
27.9.2. kubernetes.event.metadata	245
27.9.2.1. kubernetes.event.metadata.name	245
27.9.2.2. kubernetes.event.metadata.namespace	245
27.9.2.3. kubernetes.event.metadata.selfLink	246
27.9.2.4. kubernetes.event.metadata.uid	246
27.9.2.5. kubernetes.event.metadata.resourceVersion	246
27.9.3. kubernetes.event.involvedObject	246
27.9.3.1. kubernetes.event.involvedObject.kind	246
27.9.3.2. kubernetes.event.involvedObject.namespace	247
27.9.3.3. kubernetes.event.involvedObject.name	247
27.9.3.4. kubernetes.event.involvedObject.uid	247
27.9.3.5. kubernetes.event.involvedObject.apiVersion	247
27.9.3.6. kubernetes.event.involvedObject.resourceVersion	247
27.9.4. kubernetes.event.reason	248
27.9.5. kubernetes.event.source_component	248
27.9.6. kubernetes.event.firstTimestamp	248
27.9.7. kubernetes.event.count	248
27.9.8. kubernetes.event.type	248

CHAPTER 28. OPENSIFT	250
28.1. OPENSIFT.LABELS	250

CHAPTER 1. RELEASE NOTES FOR LOGGING

Logging Compatibility

The logging subsystem for Red Hat OpenShift is provided as an installable component, with a distinct release cycle from the core OpenShift Container Platform. The [Red Hat OpenShift Container Platform Life Cycle Policy](#) outlines release compatibility.

1.1. LOGGING 5.5.8

This release includes [OpenShift Logging Bug Fix Release 5.5.8](#).

1.1.1. Bug fixes

- Before this update, the **priority** field was missing from **systemd** logs due to an error in how the collector set **level** fields. With this update, these fields are set correctly, resolving the issue. ([LOG-3630](#))

1.1.2. CVEs

- [CVE-2020-10735](#)
- [CVE-2021-28861](#)
- [CVE-2022-2873](#)
- [CVE-2022-4415](#)
- [CVE-2022-24999](#)
- [CVE-2022-40897](#)
- [CVE-2022-41222](#)
- [CVE-2022-41717](#)
- [CVE-2022-43945](#)
- [CVE-2022-45061](#)
- [CVE-2022-48303](#)

1.2. LOGGING 5.5.7

This release includes [OpenShift Logging Bug Fix Release 5.5.7](#).

1.2.1. Bug fixes

- Before this update, the LokiStack Gateway Labels Enforcer generated parsing errors for valid LogQL queries when using combined label filters with boolean expressions. With this update, the LokiStack LogQL implementation supports label filters with boolean expression and resolves the issue. ([LOG-3534](#))

- Before this update, the **ClusterLogForwarder** custom resource (CR) did not pass TLS credentials for syslog output to Fluentd, resulting in errors during forwarding. With this update, credentials pass correctly to Fluentd, resolving the issue. ([LOG-3533](#))

1.2.2. CVEs

[CVE-2021-46848](#)[CVE-2022-3821](#)[CVE-2022-35737](#)[CVE-2022-42010](#)[CVE-2022-42011](#)[CVE-2022-42012](#)[CVE-2022-42898](#)[CVE-2022-43680](#)

1.3. LOGGING 5.5.6

This release includes [OpenShift Logging Bug Fix Release 5.5.6](#).

1.3.1. Known issues

1.3.2. Bug fixes

- Before this update, the Pod Security admission controller added the label **podSecurityLabelSync = true** to the **openshift-logging** namespace. This resulted in our specified security labels being overwritten, and as a result Collector pods would not start. With this update, the label **podSecurityLabelSync = false** preserves security labels. Collector pods deploy as expected. ([LOG-3340](#))
- Before this update, the Operator installed the console view plugin, even when it was not enabled on the cluster. This caused the Operator to crash. With this update, if an account for a cluster does not have the console view enabled, the Operator functions normally and does not install the console view. ([LOG-3407](#))
- Before this update, a prior fix to support a regression where the status of the Elasticsearch deployment was not being updated caused the Operator to crash unless the **Red Hat Elasticsearch Operator** was deployed. With this update, that fix has been reverted so the Operator is now stable but re-introduces the previous issue related to the reported status. ([LOG-3428](#))
- Before this update, the Loki Operator only deployed one replica of the LokiStack gateway regardless of the chosen stack size. With this update, the number of replicas is correctly configured according to the selected size. ([LOG-3478](#))
- Before this update, records written to Elasticsearch would fail if multiple label keys had the same prefix and some keys included dots. With this update, underscores replace dots in label keys, resolving the issue. ([LOG-3341](#))
- Before this update, the logging view plugin contained an incompatible feature for certain versions of OpenShift Container Platform. With this update, the correct release stream of the plugin resolves the issue. ([LOG-3467](#))
- Before this update, the reconciliation of the **ClusterLogForwarder** custom resource would incorrectly report a degraded status of one or more pipelines causing the collector pods to restart every 8-10 seconds. With this update, reconciliation of the **ClusterLogForwarder** custom resource processes correctly, resolving the issue. ([LOG-3469](#))
- Before this change the spec for the **outputDefaults** field of the ClusterLogForwarder custom resource would apply the settings to every declared Elasticsearch output type. This change corrects the behavior to match the enhancement specification where the setting specifically applies to the default managed Elasticsearch store. ([LOG-3342](#))

- Before this update, the the OpenShift CLI (oc) **must-gather** script did not complete because the OpenShift CLI (oc) needs a folder with write permission to build its cache. With this update, the OpenShift CLI (oc) has write permissions to a folder, and the **must-gather** script completes successfully. ([LOG-3472](#))
- Before this update, the Loki Operator webhook server caused TLS errors. With this update, the Loki Operator webhook PKI is managed by the Operator Lifecycle Manager's dynamic webhook management resolving the issue. ([LOG-3511](#))

1.3.3. CVEs

- [CVE-2021-46848](#)
- [CVE-2022-2056](#)
- [CVE-2022-2057](#)
- [CVE-2022-2058](#)
- [CVE-2022-2519](#)
- [CVE-2022-2520](#)
- [CVE-2022-2521](#)
- [CVE-2022-2867](#)
- [CVE-2022-2868](#)
- [CVE-2022-2869](#)
- [CVE-2022-2953](#)
- [CVE-2022-2964](#)
- [CVE-2022-4139](#)
- [CVE-2022-35737](#)
- [CVE-2022-42010](#)
- [CVE-2022-42011](#)
- [CVE-2022-42012](#)
- [CVE-2022-42898](#)
- [CVE-2022-43680](#)

1.4. LOGGING 5.5.5

This release includes [OpenShift Logging Bug Fix Release 5.5.5](#) .

1.4.1. Bug fixes

- Before this update, Kibana had a fixed **24h** OAuth cookie expiration time, which resulted in 401

errors in Kibana whenever the **accessTokenInactivityTimeout** field was set to a value lower than **24h**. With this update, Kibana's OAuth cookie expiration time synchronizes to the **accessTokenInactivityTimeout**, with a default value of **24h**. ([LOG-3305](#))

- Before this update, Vector parsed the message field when JSON parsing was enabled without also defining **structuredTypeKey** or **structuredTypeName** values. With this update, a value is required for either **structuredTypeKey** or **structuredTypeName** when writing structured logs to Elasticsearch. ([LOG-3284](#))
- Before this update, the **FluentdQueueLengthIncreasing** alert could fail to fire when there was a cardinality issue with the set of labels returned from this alert expression. This update reduces labels to only include those required for the alert. ([LOG-3226](#))
- Before this update, Loki did not have support to reach an external storage in a disconnected cluster. With this update, proxy environment variables and proxy trusted CA bundles are included in the container image to support these connections. ([LOG-2860](#))
- Before this update, OpenShift Container Platform web console users could not choose the **ConfigMap** object that includes the CA certificate for Loki, causing pods to operate without the CA. With this update, web console users can select the config map, resolving the issue. ([LOG-3310](#))
- Before this update, the CA key was used as volume name for mounting the CA into Loki, causing error states when the CA Key included non-conforming characters (such as dots). With this update, the volume name is standardized to an internal string which resolves the issue. ([LOG-3332](#))

1.4.2. CVEs

- [CVE-2016-3709](#)
- [CVE-2020-35525](#)
- [CVE-2020-35527](#)
- [CVE-2020-36516](#)
- [CVE-2020-36558](#)
- [CVE-2021-3640](#)
- [CVE-2021-30002](#)
- [CVE-2022-0168](#)
- [CVE-2022-0561](#)
- [CVE-2022-0562](#)
- [CVE-2022-0617](#)
- [CVE-2022-0854](#)
- [CVE-2022-0865](#)
- [CVE-2022-0891](#)

- [CVE-2022-0908](#)
- [CVE-2022-0909](#)
- [CVE-2022-0924](#)
- [CVE-2022-1016](#)
- [CVE-2022-1048](#)
- [CVE-2022-1055](#)
- [CVE-2022-1184](#)
- [CVE-2022-1292](#)
- [CVE-2022-1304](#)
- [CVE-2022-1355](#)
- [CVE-2022-1586](#)
- [CVE-2022-1785](#)
- [CVE-2022-1852](#)
- [CVE-2022-1897](#)
- [CVE-2022-1927](#)
- [CVE-2022-2068](#)
- [CVE-2022-2078](#)
- [CVE-2022-2097](#)
- [CVE-2022-2509](#)
- [CVE-2022-2586](#)
- [CVE-2022-2639](#)
- [CVE-2022-2938](#)
- [CVE-2022-3515](#)
- [CVE-2022-20368](#)
- [CVE-2022-21499](#)
- [CVE-2022-21618](#)
- [CVE-2022-21619](#)
- [CVE-2022-21624](#)
- [CVE-2022-21626](#)

- [CVE-2022-21628](#)
- [CVE-2022-22624](#)
- [CVE-2022-22628](#)
- [CVE-2022-22629](#)
- [CVE-2022-22662](#)
- [CVE-2022-22844](#)
- [CVE-2022-23960](#)
- [CVE-2022-24448](#)
- [CVE-2022-25255](#)
- [CVE-2022-26373](#)
- [CVE-2022-26700](#)
- [CVE-2022-26709](#)
- [CVE-2022-26710](#)
- [CVE-2022-26716](#)
- [CVE-2022-26717](#)
- [CVE-2022-26719](#)
- [CVE-2022-27404](#)
- [CVE-2022-27405](#)
- [CVE-2022-27406](#)
- [CVE-2022-27950](#)
- [CVE-2022-28390](#)
- [CVE-2022-28893](#)
- [CVE-2022-29581](#)
- [CVE-2022-30293](#)
- [CVE-2022-34903](#)
- [CVE-2022-36946](#)
- [CVE-2022-37434](#)
- [CVE-2022-39399](#)

1.5. LOGGING 5.5.4

This release includes [RHSA-2022:7434-OpenShift Logging Bug Fix Release 5.5.4](#).

1.5.1. Bug fixes

- Before this update, an error in the query parser of the logging view plugin caused parts of the logs query to disappear if the query contained curly brackets `{}`. This made the queries invalid, leading to errors being returned for valid queries. With this update, the parser correctly handles these queries. ([LOG-3042](#))
- Before this update, the Operator could enter a loop of removing and recreating the collector daemonset while the Elasticsearch or Kibana deployments changed their status. With this update, a fix in the status handling of the Operator resolves the issue. ([LOG-3049](#))
- Before this update, no alerts were implemented to support the collector implementation of Vector. This change adds Vector alerts and deploys separate alerts, depending upon the chosen collector implementation. ([LOG-3127](#))
- Before this update, the secret creation component of the Elasticsearch Operator modified internal secrets constantly. With this update, the existing secret is properly handled. ([LOG-3138](#))
- Before this update, a prior refactoring of the logging **must-gather** scripts removed the expected location for the artifacts. This update reverts that change to write artifacts to the **/must-gather** folder. ([LOG-3213](#))
- Before this update, on certain clusters, the Prometheus exporter would bind on IPv4 instead of IPv6. After this update, Fluentd detects the IP version and binds to **0.0.0.0** for IPv4 or **[::]** for IPv6. ([LOG-3162](#))

1.5.2. CVEs

- [CVE-2020-35525](#)
- [CVE-2020-35527](#)
- [CVE-2022-0494](#)
- [CVE-2022-1353](#)
- [CVE-2022-2509](#)
- [CVE-2022-2588](#)
- [CVE-2022-3515](#)
- [CVE-2022-21618](#)
- [CVE-2022-21619](#)
- [CVE-2022-21624](#)
- [CVE-2022-21626](#)
- [CVE-2022-21628](#)
- [CVE-2022-23816](#)
- [CVE-2022-23825](#)

- [CVE-2022-29900](#)
- [CVE-2022-29901](#)
- [CVE-2022-32149](#)
- [CVE-2022-37434](#)
- [CVE-2022-40674](#)

1.6. LOGGING 5.5.3

This release includes [OpenShift Logging Bug Fix Release 5.5.3](#) .

1.6.1. Bug fixes

- Before this update, log entries that had structured messages included the original message field, which made the entry larger. This update removes the message field for structured logs to reduce the increased size. ([LOG-2759](#))
- Before this update, the collector configuration excluded logs from **collector**, **default-log-store**, and **visualization** pods, but was unable to exclude logs archived in a **.gz** file. With this update, archived logs stored as **.gz** files of **collector**, **default-log-store**, and **visualization** pods are also excluded. ([LOG-2844](#))
- Before this update, when requests to an unavailable pod were sent through the gateway, no alert would warn of the disruption. With this update, individual alerts will generate if the gateway has issues completing a write or read request. ([LOG-2884](#))
- Before this update, pod metadata could be altered by fluent plugins because the values passed through the pipeline by reference. This update ensures each log message receives a copy of the pod metadata so each message processes independently. ([LOG-3046](#))
- Before this update, selecting **unknown** severity in the OpenShift Console Logs view excluded logs with a **level=unknown** value. With this update, logs without level and with **level=unknown** values are visible when filtering by **unknown** severity. ([LOG-3062](#))
- Before this update, log records sent to Elasticsearch had an extra field named **write-index** that contained the name of the index to which the logs needed to be sent. This field is not a part of the data model. After this update, this field is no longer sent. ([LOG-3075](#))
- With the introduction of the new built-in [Pod Security Admission Controller](#) , Pods not configured in accordance with the enforced security standards defined globally or on the namespace level cannot run. With this update, the Operator and collectors allow privileged execution and run without security audit warnings or errors. ([LOG-3077](#))
- Before this update, the Operator removed any custom outputs defined in the **ClusterLogForwarder** custom resource when using LokiStack as the default log storage. With this update, the Operator merges custom outputs with the default outputs when processing the **ClusterLogForwarder** custom resource. ([LOG-3095](#))

1.6.2. CVEs

- [CVE-2015-20107](#)
- [CVE-2022-0391](#)

- [CVE-2022-2526](#)
- [CVE-2022-21123](#)
- [CVE-2022-21125](#)
- [CVE-2022-21166](#)
- [CVE-2022-29154](#)
- [CVE-2022-32206](#)
- [CVE-2022-32208](#)
- [CVE-2022-34903](#)

1.7. LOGGING 5.5.2

This release includes [OpenShift Logging Bug Fix Release 5.5.2](#).

1.7.1. Bug fixes

- Before this update, alerting rules for the Fluentd collector did not adhere to the OpenShift Container Platform monitoring style guidelines. This update modifies those alerts to include the namespace label, resolving the issue. ([LOG-1823](#))
- Before this update, the index management rollover script failed to generate a new index name whenever there was more than one hyphen character in the name of the index. With this update, index names generate correctly. ([LOG-2644](#))
- Before this update, the Kibana route was setting a **caCertificate** value without a certificate present. With this update, no **caCertificate** value is set. ([LOG-2661](#))
- Before this update, a change in the collector dependencies caused it to issue a warning message for unused parameters. With this update, removing unused configuration parameters resolves the issue. ([LOG-2859](#))
- Before this update, pods created for deployments that Loki Operator created were mistakenly scheduled on nodes with non-Linux operating systems, if such nodes were available in the cluster the Operator was running in. With this update, the Operator attaches an additional node-selector to the pod definitions which only allows scheduling the pods on Linux-based nodes. ([LOG-2895](#))
- Before this update, the OpenShift Console Logs view did not filter logs by severity due to a LogQL parser issue in the LokiStack gateway. With this update, a parser fix resolves the issue and the OpenShift Console Logs view can filter by severity. ([LOG-2908](#))
- Before this update, a refactoring of the Fluentd collector plugins removed the timestamp field for events. This update restores the timestamp field, sourced from the event's received time. ([LOG-2923](#))
- Before this update, absence of a **level** field in audit logs caused an error in vector logs. With this update, the addition of a **level** field in the audit log record resolves the issue. ([LOG-2961](#))

- Before this update, if you deleted the Kibana Custom Resource, the OpenShift Container Platform web console continued displaying a link to Kibana. With this update, removing the Kibana Custom Resource also removes that link. ([LOG-3053](#))
- Before this update, each rollover job created empty indices when the **ClusterLogForwarder** custom resource had JSON parsing defined. With this update, new indices are not empty. ([LOG-3063](#))
- Before this update, when the user deleted the LokiStack after an update to Loki Operator 5.5 resources originally created by Loki Operator 5.4 remained. With this update, the resources' owner-references point to the 5.5 LokiStack. ([LOG-2945](#))
- Before this update, a user was not able to view the application logs of namespaces they have access to. With this update, the Loki Operator automatically creates a cluster role and cluster role binding allowing users to read application logs. ([LOG-2918](#))
- Before this update, users with cluster-admin privileges were not able to properly view infrastructure and audit logs using the logging console. With this update, the authorization check has been extended to also recognize users in cluster-admin and dedicated-admin groups as admins. ([LOG-2970](#))

1.7.2. CVEs

- [CVE-2015-20107](#)
- [CVE-2022-0391](#)
- [CVE-2022-21123](#)
- [CVE-2022-21125](#)
- [CVE-2022-21166](#)
- [CVE-2022-29154](#)
- [CVE-2022-32206](#)
- [CVE-2022-32208](#)
- [CVE-2022-34903](#)

1.8. LOGGING 5.5.1

This release includes [OpenShift Logging Bug Fix Release 5.5.1](#).

1.8.1. Enhancements

- This enhancement adds an **Aggregated Logs** tab to the **Pod Details** page of the OpenShift Container Platform web console when the Logging Console Plugin is in use. This enhancement is only available on OpenShift Container Platform 4.10 and later. ([LOG-2647](#))
- This enhancement adds Google Cloud Logging as an output option for log forwarding. ([LOG-1482](#))

1.8.2. Bug fixes

- Before this update, the Operator did not ensure that the pod was ready, which caused the cluster to reach an inoperable state during a cluster restart. With this update, the Operator marks new pods as ready before continuing to a new pod during a restart, which resolves the issue. ([LOG-2745](#))
- Before this update, Fluentd would sometimes not recognize that the Kubernetes platform rotated the log file and would no longer read log messages. This update corrects that by setting the configuration parameter suggested by the upstream development team. ([LOG-2995](#))
- Before this update, the addition of multi-line error detection caused internal routing to change and forward records to the wrong destination. With this update, the internal routing is correct. ([LOG-2801](#))
- Before this update, changing the OpenShift Container Platform web console's refresh interval created an error when the **Query** field was empty. With this update, changing the interval is not an available option when the **Query** field is empty. ([LOG-2917](#))

1.8.3. CVEs


- [CVE-2022-1705](#)
- [CVE-2022-2526](#)
- [CVE-2022-29154](#)
- [CVE-2022-30631](#)
- [CVE-2022-32148](#)
- [CVE-2022-32206](#)
- [CVE-2022-32208](#)

1.9. LOGGING 5.5

The following advisories are available for Logging 5.5:[Release 5.5](#)

1.9.1. Enhancements

- With this update, you can forward structured logs from different containers within the same pod to different indices. To use this feature, you must configure the pipeline with multi-container support and annotate the pods. ([LOG-1296](#))



IMPORTANT

JSON formatting of logs varies by application. Because creating too many indices impacts performance, limit your use of this feature to creating indices for logs that have incompatible JSON formats. Use queries to separate logs from different namespaces, or applications with compatible JSON formats.

- With this update, you can filter logs with Elasticsearch outputs by using the Kubernetes common labels, **app.kubernetes.io/component**, **app.kubernetes.io/managed-by**, **app.kubernetes.io/part-of**, and **app.kubernetes.io/version**. Non-Elasticsearch output types can use all labels included in **kubernetes.labels**. ([LOG-2388](#))

- With this update, clusters with AWS Security Token Service (STS) enabled may use STS authentication to forward logs to Amazon CloudWatch. ([LOG-1976](#))
- With this update, the 'LokiOperator' Operator and Vector collector move from Technical Preview to General Availability. Full feature parity with prior releases are pending, and some APIs remain Technical Previews. See the **Logging with the LokiStack** section for details.

1.9.2. Bug fixes

- Before this update, clusters configured to forward logs to Amazon CloudWatch wrote rejected log files to temporary storage, causing cluster instability over time. With this update, chunk backup for all storage options has been disabled, resolving the issue. ([LOG-2746](#))
- Before this update, the Operator was using versions of some APIs that are deprecated and planned for removal in future versions of OpenShift Container Platform. This update moves dependencies to the supported API versions. ([LOG-2656](#))

Before this update, the Operator was using versions of some APIs that are deprecated and planned for removal in future versions of OpenShift Container Platform. This update moves dependencies to the supported API versions. ([LOG-2656](#))

- Before this update, multiple **ClusterLogForwarder** pipelines configured for multiline error detection caused the collector to go into a **crashloopbackoff** error state. This update fixes the issue where multiple configuration sections had the same unique ID. ([LOG-2241](#))
- Before this update, the collector could not save non UTF-8 symbols to the Elasticsearch storage logs. With this update the collector encodes non UTF-8 symbols, resolving the issue. ([LOG-2203](#))
- Before this update, non-latin characters displayed incorrectly in Kibana. With this update, Kibana displays all valid UTF-8 symbols correctly. ([LOG-2784](#))

1.9.3. CVEs

- [CVE-2021-38561](#)
- [CVE-2022-1012](#)
- [CVE-2022-1292](#)
- [CVE-2022-1586](#)
- [CVE-2022-1785](#)
- [CVE-2022-1897](#)
- [CVE-2022-1927](#)
- [CVE-2022-2068](#)
- [CVE-2022-2097](#)
- [CVE-2022-21698](#)
- [CVE-2022-30631](#)
- [CVE-2022-32250](#)

1.10. LOGGING 5.4.14

This release includes [OpenShift Logging Bug Fix Release 5.4.14](#) .

1.10.1. Bug fixes

None.

1.10.2. CVEs

- [CVE-2022-4304](#)
- [CVE-2022-4450](#)
- [CVE-2023-0215](#)
- [CVE-2023-0286](#)
- [CVE-2023-0361](#)
- [CVE-2023-23916](#)

1.11. LOGGING 5.4.13

This release includes [OpenShift Logging Bug Fix Release 5.4.13](#) .

1.11.1. Bug fixes

- Before this update, a problem with the Fluentd collector caused it to not capture OAuth login events stored in **`/var/log/auth-server/audit.log`**. This led to incomplete collection of login events from the OAuth service. With this update, the Fluentd collector now resolves this issue by capturing all login events from the OAuth service, including those stored in **`/var/log/auth-server/audit.log`**, as expected. ([LOG-3731](#))

1.11.2. CVEs

- [CVE-2022-4304](#)
- [CVE-2022-4450](#)
- [CVE-2023-0215](#)
- [CVE-2023-0286](#)
- [CVE-2023-0767](#)
- [CVE-2023-23916](#)

1.12. LOGGING 5.4.12

This release includes [OpenShift Logging Bug Fix Release 5.4.12](#) .

1.12.1. Bug fixes

None.

1.12.2. CVEs

- [CVE-2020-10735](#)
- [CVE-2021-28861](#)
- [CVE-2022-2873](#)
- [CVE-2022-4415](#)
- [CVE-2022-40897](#)
- [CVE-2022-41222](#)
- [CVE-2022-41717](#)
- [CVE-2022-43945](#)
- [CVE-2022-45061](#)
- [CVE-2022-48303](#)

1.13. LOGGING 5.4.11

This release includes [OpenShift Logging Bug Fix Release 5.4.11](#) .

1.13.1. Bug fixes

- [BZ 2099524](#)
- [BZ 2161274](#)

1.13.2. CVEs

- [CVE-2021-46848](#)
- [CVE-2022-3821](#)
- [CVE-2022-35737](#)
- [CVE-2022-42010](#)
- [CVE-2022-42011](#)
- [CVE-2022-42012](#)
- [CVE-2022-42898](#)
- [CVE-2022-43680](#)

1.14. LOGGING 5.4.10

This release includes [OpenShift Logging Bug Fix Release 5.4.10](#) .

1.14.1. Bug fixes

None.

1.14.2. CVEs

- [CVE-2021-46848](#)
- [CVE-2022-2056](#)
- [CVE-2022-2057](#)
- [CVE-2022-2058](#)
- [CVE-2022-2519](#)
- [CVE-2022-2520](#)
- [CVE-2022-2521](#)
- [CVE-2022-2867](#)
- [CVE-2022-2868](#)
- [CVE-2022-2869](#)
- [CVE-2022-2953](#)
- [CVE-2022-2964](#)
- [CVE-2022-4139](#)
- [CVE-2022-35737](#)
- [CVE-2022-42010](#)
- [CVE-2022-42011](#)
- [CVE-2022-42012](#)
- [CVE-2022-42898](#)
- [CVE-2022-43680](#)

1.15. LOGGING 5.4.9

This release includes [OpenShift Logging Bug Fix Release 5.4.9](#) .

1.15.1. Bug fixes

- Before this update, the Fluentd collector would warn of unused configuration parameters. This update removes those configuration parameters and their warning messages. ([LOG-3074](#))
- Before this update, Kibana had a fixed **24h** OAuth cookie expiration time, which resulted in 401 errors in Kibana whenever the **accessTokenInactivityTimeout** field was set to a value lower than **24h**. With this update, Kibana's OAuth cookie expiration time synchronizes to the

accessTokenInactivityTimeout, with a default value of **24h**. ([LOG-3306](#))

1.15.2. CVEs

- [CVE-2016-3709](#)
- [CVE-2020-35525](#)
- [CVE-2020-35527](#)
- [CVE-2020-36516](#)
- [CVE-2020-36558](#)
- [CVE-2021-3640](#)
- [CVE-2021-30002](#)
- [CVE-2022-0168](#)
- [CVE-2022-0561](#)
- [CVE-2022-0562](#)
- [CVE-2022-0617](#)
- [CVE-2022-0854](#)
- [CVE-2022-0865](#)
- [CVE-2022-0891](#)
- [CVE-2022-0908](#)
- [CVE-2022-0909](#)
- [CVE-2022-0924](#)
- [CVE-2022-1016](#)
- [CVE-2022-1048](#)
- [CVE-2022-1055](#)
- [CVE-2022-1184](#)
- [CVE-2022-1292](#)
- [CVE-2022-1304](#)
- [CVE-2022-1355](#)
- [CVE-2022-1586](#)
- [CVE-2022-1785](#)
- [CVE-2022-1852](#)

- [CVE-2022-1897](#)
- [CVE-2022-1927](#)
- [CVE-2022-2068](#)
- [CVE-2022-2078](#)
- [CVE-2022-2097](#)
- [CVE-2022-2509](#)
- [CVE-2022-2586](#)
- [CVE-2022-2639](#)
- [CVE-2022-2938](#)
- [CVE-2022-3515](#)
- [CVE-2022-20368](#)
- [CVE-2022-21499](#)
- [CVE-2022-21618](#)
- [CVE-2022-21619](#)
- [CVE-2022-21624](#)
- [CVE-2022-21626](#)
- [CVE-2022-21628](#)
- [CVE-2022-22624](#)
- [CVE-2022-22628](#)
- [CVE-2022-22629](#)
- [CVE-2022-22662](#)
- [CVE-2022-22844](#)
- [CVE-2022-23960](#)
- [CVE-2022-24448](#)
- [CVE-2022-25255](#)
- [CVE-2022-26373](#)
- [CVE-2022-26700](#)
- [CVE-2022-26709](#)
- [CVE-2022-26710](#)

- [CVE-2022-26716](#)
- [CVE-2022-26717](#)
- [CVE-2022-26719](#)
- [CVE-2022-27404](#)
- [CVE-2022-27405](#)
- [CVE-2022-27406](#)
- [CVE-2022-27950](#)
- [CVE-2022-28390](#)
- [CVE-2022-28893](#)
- [CVE-2022-29581](#)
- [CVE-2022-30293](#)
- [CVE-2022-34903](#)
- [CVE-2022-36946](#)
- [CVE-2022-37434](#)
- [CVE-2022-39399](#)

1.16. LOGGING 5.4.8

This release includes [RHSA-2022:7435-OpenShift Logging Bug Fix Release 5.4.8](#) .

1.16.1. Bug fixes

None.

1.16.2. CVEs

- [CVE-2016-3709](#)
- [CVE-2020-35525](#)
- [CVE-2020-35527](#)
- [CVE-2020-36518](#)
- [CVE-2022-1304](#)
- [CVE-2022-2509](#)
- [CVE-2022-3515](#)
- [CVE-2022-22624](#)

- [CVE-2022-22628](#)
- [CVE-2022-22629](#)
- [CVE-2022-22662](#)
- [CVE-2022-26700](#)
- [CVE-2022-26709](#)
- [CVE-2022-26710](#)
- [CVE-2022-26716](#)
- [CVE-2022-26717](#)
- [CVE-2022-26719](#)
- [CVE-2022-30293](#)
- [CVE-2022-32149](#)
- [CVE-2022-37434](#)
- [CVE-2022-40674](#)
- [CVE-2022-42003](#)
- [CVE-2022-42004](#)

1.17. LOGGING 5.4.6

This release includes [OpenShift Logging Bug Fix Release 5.4.6](#) .

1.17.1. Bug fixes

- Before this update, Fluentd would sometimes not recognize that the Kubernetes platform rotated the log file and would no longer read log messages. This update corrects that by setting the configuration parameter suggested by the upstream development team. ([LOG-2792](#))
- Before this update, each rollover job created empty indices when the **ClusterLogForwarder** custom resource had JSON parsing defined. With this update, new indices are not empty. ([LOG-2823](#))
- Before this update, if you deleted the Kibana Custom Resource, the OpenShift Container Platform web console continued displaying a link to Kibana. With this update, removing the Kibana Custom Resource also removes that link. ([LOG-3054](#))

1.17.2. CVEs

- [CVE-2015-20107](#)
- [CVE-2022-0391](#)
- [CVE-2022-21123](#)

- [CVE-2022-21125](#)
- [CVE-2022-21166](#)
- [CVE-2022-29154](#)
- [CVE-2022-32206](#)
- [CVE-2022-32208](#)
- [CVE-2022-34903](#)

1.18. LOGGING 5.4.5

This release includes [RHSA-2022:6183-OpenShift Logging Bug Fix Release 5.4.5](#).

1.18.1. Bug fixes

- Before this update, the Operator did not ensure that the pod was ready, which caused the cluster to reach an inoperable state during a cluster restart. With this update, the Operator marks new pods as ready before continuing to a new pod during a restart, which resolves the issue. ([LOG-2881](#))
- Before this update, the addition of multi-line error detection caused internal routing to change and forward records to the wrong destination. With this update, the internal routing is correct. ([LOG-2946](#))
- Before this update, the Operator could not decode index setting JSON responses with a quoted Boolean value and would result in an error. With this update, the Operator can properly decode this JSON response. ([LOG-3009](#))
- Before this update, Elasticsearch index templates defined the fields for labels with the wrong types. This change updates those templates to match the expected types forwarded by the log collector. ([LOG-2972](#))

1.18.2. CVEs

- [CVE-2022-1292](#)
- [CVE-2022-1586](#)
- [CVE-2022-1785](#)
- [CVE-2022-1897](#)
- [CVE-2022-1927](#)
- [CVE-2022-2068](#)
- [CVE-2022-2097](#)
- [CVE-2022-30631](#)

1.19. LOGGING 5.4.4

This release includes [RHBA-2022:5907-OpenShift Logging Bug Fix Release 5.4.4](#) .

1.19.1. Bug fixes

- Before this update, non-latin characters displayed incorrectly in Elasticsearch. With this update, Elasticsearch displays all valid UTF-8 symbols correctly. ([LOG-2794](#))
- Before this update, non-latin characters displayed incorrectly in Fluentd. With this update, Fluentd displays all valid UTF-8 symbols correctly. ([LOG-2657](#))
- Before this update, the metrics server for the collector attempted to bind to the address using a value exposed by an environment value. This change modifies the configuration to bind to any available interface. ([LOG-2821](#))
- Before this update, the **cluster-logging** Operator relied on the cluster to create a secret. This cluster behavior changed in OpenShift Container Platform 4.11, which caused logging deployments to fail. With this update, the **cluster-logging** Operator resolves the issue by creating the secret if needed. ([LOG-2840](#))

1.19.2. CVEs

- [CVE-2022-21540](#)
- [CVE-2022-21541](#)
- [CVE-2022-34169](#)

1.20. LOGGING 5.4.3

This release includes [RHSA-2022:5556-OpenShift Logging Bug Fix Release 5.4.3](#) .

1.20.1. Elasticsearch Operator deprecation notice

In logging subsystem 5.4.3 the Elasticsearch Operator is deprecated and is planned to be removed in a future release. Red Hat will provide bug fixes and support for this feature during the current release lifecycle, but this feature will no longer receive enhancements and will be removed. As an alternative to using the Elasticsearch Operator to manage the default log storage, you can use the Loki Operator.

1.20.2. Bug fixes

- Before this update, the OpenShift Logging Dashboard showed the number of active primary shards instead of all active shards. With this update, the dashboard displays all active shards. ([LOG-2781](#))
- Before this update, a bug in a library used by **elasticsearch-operator** contained a denial of service attack vulnerability. With this update, the library has been updated to a version that does not contain this vulnerability. ([LOG-2816](#))
- Before this update, when configuring Vector to forward logs to Loki, it was not possible to set a custom bearer token or use the default token if Loki had TLS enabled. With this update, Vector can forward logs to Loki using tokens with TLS enabled. ([LOG-2786](#))
- Before this update, the ElasticSearch Operator omitted the **referencePolicy** property of the **ImageStream** custom resource when selecting an **oauth-proxy** image. This omission caused the Kibana deployment to fail in specific environments. With this update, using **referencePolicy**

resolves the issue, and the Operator can deploy Kibana successfully. ([LOG-2791](#))

- Before this update, alerting rules for the **ClusterLogForwarder** custom resource did not take multiple forward outputs into account. This update resolves the issue. ([LOG-2640](#))
- Before this update, clusters configured to forward logs to Amazon CloudWatch wrote rejected log files to temporary storage, causing cluster instability over time. With this update, chunk backup for CloudWatch has been disabled, resolving the issue. ([LOG-2768](#))

1.20.3. CVEs

Example 1.1. Click to expand CVEs

- [CVE-2020-28915](#)
- [CVE-2021-40528](#)
- [CVE-2022-1271](#)
- [CVE-2022-1621](#)
- [CVE-2022-1629](#)
- [CVE-2022-22576](#)
- [CVE-2022-25313](#)
- [CVE-2022-25314](#)
- [CVE-2022-26691](#)
- [CVE-2022-27666](#)
- [CVE-2022-27774](#)
- [CVE-2022-27776](#)
- [CVE-2022-27782](#)
- [CVE-2022-29824](#)

1.21. LOGGING 5.4.2

This release includes [RHBA-2022:4874-OpenShift Logging Bug Fix Release 5.4.2](#)

1.21.1. Bug fixes

- Before this update, editing the Collector configuration using **oc edit** was difficult because it had inconsistent use of white-space. This change introduces logic to normalize and format the configuration prior to any updates by the Operator so that it is easy to edit using **oc edit**. ([LOG-2319](#))
- Before this update, the **FluentdNodeDown** alert could not provide instance labels in the message section appropriately. This update resolves the issue by fixing the alert rule to provide instance labels in cases of partial instance failures. ([LOG-2607](#))

- Before this update, several log levels, such as `critical`, that were documented as supported by the product were not. This update fixes the discrepancy so the documented log levels are now supported by the product. ([LOG-2033](#))

1.21.2. CVEs

Example 1.2. Click to expand CVEs

- [CVE-2018-25032](#)
- [CVE-2020-0404](#)
- [CVE-2020-4788](#)
- [CVE-2020-13974](#)
- [CVE-2020-19131](#)
- [CVE-2020-27820](#)
- [CVE-2021-0941](#)
- [CVE-2021-3612](#)
- [CVE-2021-3634](#)
- [CVE-2021-3669](#)
- [CVE-2021-3737](#)
- [CVE-2021-3743](#)
- [CVE-2021-3744](#)
- [CVE-2021-3752](#)
- [CVE-2021-3759](#)
- [CVE-2021-3764](#)
- [CVE-2021-3772](#)
- [CVE-2021-3773](#)
- [CVE-2021-4002](#)
- [CVE-2021-4037](#)
- [CVE-2021-4083](#)
- [CVE-2021-4157](#)
- [CVE-2021-4189](#)
- [CVE-2021-4197](#)
- [CVE-2021-4203](#)

- [CVE-2021-20322](#)
- [CVE-2021-21781](#)
- [CVE-2021-23222](#)
- [CVE-2021-26401](#)
- [CVE-2021-29154](#)
- [CVE-2021-37159](#)
- [CVE-2021-41617](#)
- [CVE-2021-41864](#)
- [CVE-2021-42739](#)
- [CVE-2021-43056](#)
- [CVE-2021-43389](#)
- [CVE-2021-43976](#)
- [CVE-2021-44733](#)
- [CVE-2021-45485](#)
- [CVE-2021-45486](#)
- [CVE-2022-0001](#)
- [CVE-2022-0002](#)
- [CVE-2022-0286](#)
- [CVE-2022-0322](#)
- [CVE-2022-1011](#)
- [CVE-2022-1271](#)

1.22. LOGGING 5.4.1

This release includes [RHSA-2022:2216-OpenShift Logging Bug Fix Release 5.4.1](#).

1.22.1. Bug fixes

- Before this update, the log file metric exporter only reported logs created while the exporter was running, which resulted in inaccurate log growth data. This update resolves this issue by monitoring `/var/log/pods`. ([LOG-2442](#))
- Before this update, the collector would be blocked because it continually tried to use a stale connection when forwarding logs to fluentd forward receivers. With this release, the `keepalive_timeout` value has been set to 30 seconds (**30s**) so that the collector recycles the

connection and re-attempts to send failed messages within a reasonable amount of time. ([LOG-2534](#))

- Before this update, an error in the gateway component enforcing tenancy for reading logs limited access to logs with a Kubernetes namespace causing "audit" and some "infrastructure" logs to be unreadable. With this update, the proxy correctly detects users with admin access and allows access to logs without a namespace. ([LOG-2448](#))
- Before this update, the **system:serviceaccount:openshift-monitoring:prometheus-k8s** service account had cluster level privileges as a **clusterrole** and **clusterrolebinding**. This update restricts the service account to the **openshift-logging** namespace with a role and rolebinding. ([LOG-2437](#))
- Before this update, Linux audit log time parsing relied on an ordinal position of a key/value pair. This update changes the parsing to use a regular expression to find the time entry. ([LOG-2321](#))

1.22.2. CVEs


Example 1.3. Click to expand CVEs

- [CVE-2018-25032](#)
- [CVE-2021-4028](#)
- [CVE-2021-37136](#)
- [CVE-2021-37137](#)
- [CVE-2021-43797](#)
- [CVE-2022-0778](#)
- [CVE-2022-1154](#)
- [CVE-2022-1271](#)
- [CVE-2022-21426](#)
- [CVE-2022-21434](#)
- [CVE-2022-21443](#)
- [CVE-2022-21476](#)
- [CVE-2022-21496](#)
- [CVE-2022-21698](#)
- [CVE-2022-25636](#)

1.23. LOGGING 5.4

The following advisories are available for logging 5.4: [Logging subsystem for Red Hat OpenShift Release 5.4](#)

1.23.1. Technology Previews



IMPORTANT

Vector is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

1.23.2. About Vector


Vector is a log collector offered as a tech-preview alternative to the current default collector for the logging subsystem.

The following outputs are supported:

- **elasticsearch**. An external Elasticsearch instance. The **elasticsearch** output can use a TLS connection.
- **kafka**. A Kafka broker. The **kafka** output can use an unsecured or TLS connection.
- **loki**. Loki, a horizontally scalable, highly available, multi-tenant log aggregation system.

1.23.2.1. Enabling Vector

Vector is not enabled by default. Use the following steps to enable Vector on your OpenShift Container Platform cluster.



IMPORTANT

Vector does not support FIPS Enabled Clusters.

Prerequisites

- OpenShift Container Platform: 4.10
- Logging subsystem for Red Hat OpenShift: 5.4
- FIPS disabled

Procedure

1. Edit the **ClusterLogging** custom resource (CR) in the **openshift-logging** project:

```
$ oc -n openshift-logging edit ClusterLogging instance
```

2. Add a **logging.openshift.io/preview-vector-collector: enabled** annotation to the **ClusterLogging** custom resource (CR).
3. Add **vector** as a collection type to the **ClusterLogging** custom resource (CR).


```

apiVersion: "logging.openshift.io/v1"
kind: "ClusterLogging"
metadata:
  name: "instance"
  namespace: "openshift-logging"
  annotations:
    logging.openshift.io/preview-vector-collector: enabled
spec:
  collection:
  logs:
    type: "vector"
    vector: {}

```

Additional resources

- [Vector Documentation](#)



IMPORTANT

Loki Operator is a Technology Preview feature only. Technology Preview features are not supported with Red Hat production service level agreements (SLAs) and might not be functionally complete. Red Hat does not recommend using them in production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process.

For more information about the support scope of Red Hat Technology Preview features, see [Technology Preview Features Support Scope](#).

1.23.3. About Loki

Loki is a horizontally scalable, highly available, multi-tenant log aggregation system currently offered as an alternative to Elasticsearch as a log store for the logging subsystem.

Additional resources

- [Loki Documentation](#)

1.23.3.1. Deploying the Lokistack

You can use the OpenShift Container Platform web console to install the LokiOperator.

Prerequisites

- OpenShift Container Platform: 4.10
- Logging subsystem for Red Hat OpenShift: 5.4

To install the LokiOperator using the OpenShift Container Platform web console:

1. Install the LokiOperator:
 - a. In the OpenShift Container Platform web console, click **Operators** → **OperatorHub**.
 - b. Choose **LokiOperator** from the list of available Operators, and click **Install**.

- c. Under **Installation Mode**, select **All namespaces on the cluster**.
- d. Under **Installed Namespace**, select **openshift-operators-redhat**.
You must specify the **openshift-operators-redhat** namespace. The **openshift-operators** namespace might contain Community Operators, which are untrusted and could publish a metric with the same name as an OpenShift Container Platform metric, which would cause conflicts.
- e. Select **Enable operator recommended cluster monitoring on this namespace**
This option sets the **openshift.io/cluster-monitoring: "true"** label in the Namespace object. You must select this option to ensure that cluster monitoring scrapes the **openshift-operators-redhat** namespace.
- f. Select an **Approval Strategy**.
 - The **Automatic** strategy allows Operator Lifecycle Manager (OLM) to automatically update the Operator when a new version is available.
 - The **Manual** strategy requires a user with appropriate credentials to approve the Operator update.
- g. Click **Install**.
- h. Verify that you installed the LokiOperator. Visit the **Operators → Installed Operators** page and look for "LokiOperator."
- i. Ensure that **LokiOperator** is listed in all the projects whose **Status** is **Succeeded**.

1.23.4. Bug fixes

- Before this update, the **cluster-logging-operator** used cluster scoped roles and bindings to establish permissions for the Prometheus service account to scrape metrics. These permissions were created when deploying the Operator using the console interface but were missing when deploying from the command line. This update fixes the issue by making the roles and bindings namespace-scoped. ([LOG-2286](#))
- Before this update, a prior change to fix dashboard reconciliation introduced a **ownerReferences** field to the resource across namespaces. As a result, both the config map and dashboard were not created in the namespace. With this update, the removal of the **ownerReferences** field resolves the issue, and the OpenShift Logging dashboard is available in the console. ([LOG-2163](#))
- Before this update, changes to the metrics dashboards did not deploy because the **cluster-logging-operator** did not correctly compare existing and modified config maps that contain the dashboard. With this update, the addition of a unique hash value to object labels resolves the issue. ([LOG-2071](#))
- Before this update, the OpenShift Logging dashboard did not correctly display the pods and namespaces in the table, which displays the top producing containers collected over the last 24 hours. With this update, the pods and namespaces are displayed correctly. ([LOG-2069](#))
- Before this update, when the **ClusterLogForwarder** was set up with **Elasticsearch OutputDefault** and Elasticsearch outputs did not have structured keys, the generated configuration contained the incorrect values for authentication. This update corrects the secret and certificates used. ([LOG-2056](#))

- Before this update, the OpenShift Logging dashboard displayed an empty CPU graph because of a reference to an invalid metric. With this update, the correct data point has been selected, resolving the issue. ([LOG-2026](#))
- Before this update, the Fluentd container image included builder tools that were unnecessary at run time. This update removes those tools from the image. ([LOG-1927](#))
- Before this update, a name change of the deployed collector in the 5.3 release caused the logging collector to generate the **FluentdNodeDown** alert. This update resolves the issue by fixing the job name for the Prometheus alert. ([LOG-1918](#))
- Before this update, the log collector was collecting its own logs due to a refactoring of the component name change. This led to a potential feedback loop of the collector processing its own log that might result in memory and log message size issues. This update resolves the issue by excluding the collector logs from the collection. ([LOG-1774](#))
- Before this update, Elasticsearch generated the error **Unable to create PersistentVolumeClaim due to forbidden: exceeded quota: infra-storage-quota**. if the PVC already existed. With this update, Elasticsearch checks for existing PVCs, resolving the issue. ([LOG-2131](#))
- Before this update, Elasticsearch was unable to return to the ready state when the **elasticsearch-signing** secret was removed. With this update, Elasticsearch is able to go back to the ready state after that secret is removed. ([LOG-2171](#))
- Before this update, the change of the path from which the collector reads container logs caused the collector to forward some records to the wrong indices. With this update, the collector now uses the correct configuration to resolve the issue. ([LOG-2160](#))
- Before this update, clusters with a large number of namespaces caused Elasticsearch to stop serving requests because the list of namespaces reached the maximum header size limit. With this update, headers only include a list of namespace names, resolving the issue. ([LOG-1899](#))
- Before this update, the **OpenShift Container Platform Logging** dashboard showed the number of shards 'x' times larger than the actual value when Elasticsearch had 'x' nodes. This issue occurred because it was printing all primary shards for each Elasticsearch pod and calculating a sum on it, although the output was always for the whole Elasticsearch cluster. With this update, the number of shards is now correctly calculated. ([LOG-2156](#))
- Before this update, the secrets **kibana** and **kibana-proxy** were not recreated if they were deleted manually. With this update, the **elasticsearch-operator** will watch the resources and automatically recreate them if deleted. ([LOG-2250](#))
- Before this update, tuning the buffer chunk size could cause the collector to generate a warning about the chunk size exceeding the byte limit for the event stream. With this update, you can also tune the read line limit, resolving the issue. ([LOG-2379](#))
- Before this update, the logging console link in OpenShift web console was not removed with the ClusterLogging CR. With this update, deleting the CR or uninstalling the Cluster Logging Operator removes the link. ([LOG-2373](#))
- Before this update, a change to the container logs path caused the collection metric to always be zero with older releases configured with the original path. With this update, the plugin which exposes metrics about collected logs supports reading from either path to resolve the issue. ([LOG-2462](#))

1.23.5. CVEs

- [CVE-2022-0759](#)
 - [BZ-2058404](#)
- [CVE-2022-21698](#)
 - [BZ-2045880](#)

1.24. LOGGING 5.3.14

This release includes [OpenShift Logging Bug Fix Release 5.3.14](#).

1.24.1. Bug fixes

- Before this update, the log file size map generated by the **log-file-metrics-exporter** component did not remove entries for deleted files, resulting in increased file size, and process memory. With this update, the log file size map does not contain entries for deleted files. ([LOG-3293](#))

1.24.2. CVEs

- [CVE-2016-3709](#)
- [CVE-2020-35525](#)
- [CVE-2020-35527](#)
- [CVE-2020-36516](#)
- [CVE-2020-36558](#)
- [CVE-2021-3640](#)
- [CVE-2021-30002](#)
- [CVE-2022-0168](#)
- [CVE-2022-0561](#)
- [CVE-2022-0562](#)
- [CVE-2022-0617](#)
- [CVE-2022-0854](#)
- [CVE-2022-0865](#)
- [CVE-2022-0891](#)
- [CVE-2022-0908](#)
- [CVE-2022-0909](#)
- [CVE-2022-0924](#)
- [CVE-2022-1016](#)

- [CVE-2022-1048](#)
- [CVE-2022-1055](#)
- [CVE-2022-1184](#)
- [CVE-2022-1292](#)
- [CVE-2022-1304](#)
- [CVE-2022-1355](#)
- [CVE-2022-1586](#)
- [CVE-2022-1785](#)
- [CVE-2022-1852](#)
- [CVE-2022-1897](#)
- [CVE-2022-1927](#)
- [CVE-2022-2068](#)
- [CVE-2022-2078](#)
- [CVE-2022-2097](#)
- [CVE-2022-2509](#)
- [CVE-2022-2586](#)
- [CVE-2022-2639](#)
- [CVE-2022-2938](#)
- [CVE-2022-3515](#)
- [CVE-2022-20368](#)
- [CVE-2022-21499](#)
- [CVE-2022-21618](#)
- [CVE-2022-21619](#)
- [CVE-2022-21624](#)
- [CVE-2022-21626](#)
- [CVE-2022-21628](#)
- [CVE-2022-22624](#)
- [CVE-2022-22628](#)
- [CVE-2022-22629](#)

- [CVE-2022-22662](#)
- [CVE-2022-22844](#)
- [CVE-2022-23960](#)
- [CVE-2022-24448](#)
- [CVE-2022-25255](#)
- [CVE-2022-26373](#)
- [CVE-2022-26700](#)
- [CVE-2022-26709](#)
- [CVE-2022-26710](#)
- [CVE-2022-26716](#)
- [CVE-2022-26717](#)
- [CVE-2022-26719](#)
- [CVE-2022-27404](#)
- [CVE-2022-27405](#)
- [CVE-2022-27406](#)
- [CVE-2022-27950](#)
- [CVE-2022-28390](#)
- [CVE-2022-28893](#)
- [CVE-2022-29581](#)
- [CVE-2022-30293](#)
- [CVE-2022-34903](#)
- [CVE-2022-36946](#)
- [CVE-2022-37434](#)
- [CVE-2022-39399](#)
- [CVE-2022-42898](#)

1.25. LOGGING 5.3.13

This release includes [RHSA-2022:68828-OpenShift Logging Bug Fix Release 5.3.13](#).

1.25.1. Bug fixes

None.

1.25.2. CVEs

Example 1.4. Click to expand CVEs

- [CVE-2020-35525](#)
- [CVE-2020-35527](#)
- [CVE-2022-0494](#)
- [CVE-2022-1353](#)
- [CVE-2022-2509](#)
- [CVE-2022-2588](#)
- [CVE-2022-3515](#)
- [CVE-2022-21618](#)
- [CVE-2022-21619](#)
- [CVE-2022-21624](#)
- [CVE-2022-21626](#)
- [CVE-2022-21628](#)
- [CVE-2022-23816](#)
- [CVE-2022-23825](#)
- [CVE-2022-29900](#)
- [CVE-2022-29901](#)
- [CVE-2022-32149](#)
- [CVE-2022-37434](#)
- [CVE-2022-39399](#)
- [CVE-2022-40674](#)

1.26. LOGGING 5.3.12

This release includes [OpenShift Logging Bug Fix Release 5.3.12](#) .

1.26.1. Bug fixes

None.

1.26.2. CVEs

- [CVE-2015-20107](#)

- [CVE-2022-0391](#)
- [CVE-2022-21123](#)
- [CVE-2022-21125](#)
- [CVE-2022-21166](#)
- [CVE-2022-29154](#)
- [CVE-2022-32206](#)
- [CVE-2022-32208](#)
- [CVE-2022-34903](#)

1.27. LOGGING 5.3.11

This release includes [OpenShift Logging Bug Fix Release 5.3.11](#).

1.27.1. Bug fixes

- Before this update, the Operator did not ensure that the pod was ready, which caused the cluster to reach an inoperable state during a cluster restart. With this update, the Operator marks new pods as ready before continuing to a new pod during a restart, which resolves the issue. ([LOG-2871](#))

1.27.2. CVEs

- [CVE-2022-1292](#)
- [CVE-2022-1586](#)
- [CVE-2022-1785](#)
- [CVE-2022-1897](#)
- [CVE-2022-1927](#)
- [CVE-2022-2068](#)
- [CVE-2022-2097](#)
- [CVE-2022-30631](#)

1.28. LOGGING 5.3.10

This release includes [RHSA-2022:5908-OpenShift Logging Bug Fix Release 5.3.10](#).

1.28.1. Bug fixes

- [BZ-2100495](#)

1.28.2. CVEs

-

Example 1.5. Click to expand CVEs

- [CVE-2021-38561](#)
- [CVE-2021-40528](#)
- [CVE-2022-1271](#)
- [CVE-2022-1621](#)
- [CVE-2022-1629](#)
- [CVE-2022-21540](#)
- [CVE-2022-21541](#)
- [CVE-2022-22576](#)
- [CVE-2022-25313](#)
- [CVE-2022-25314](#)
- [CVE-2022-27774](#)
- [CVE-2022-27776](#)
- [CVE-2022-27782](#)
- [CVE-2022-29824](#)
- [CVE-2022-34169](#)

1.29. LOGGING 5.3.9

This release includes [RHBA-2022:5557-OpenShift Logging Bug Fix Release 5.3.9](#).

1.29.1. Bug fixes

- Before this update, the logging collector included a path as a label for the metrics it produced. This path changed frequently and contributed to significant storage changes for the Prometheus server. With this update, the label has been dropped to resolve the issue and reduce storage consumption. ([LOG-2682](#))

1.29.2. CVEs

Example 1.6. Click to expand CVEs

- [CVE-2020-28915](#)
- [CVE-2021-40528](#)
- [CVE-2022-1271](#)
- [CVE-2022-1621](#)

- [CVE-2022-1629](#)
- [CVE-2022-22576](#)
- [CVE-2022-25313](#)
- [CVE-2022-25314](#)
- [CVE-2022-26691](#)
- [CVE-2022-27666](#)
- [CVE-2022-27774](#)
- [CVE-2022-27776](#)
- [CVE-2022-27782](#)
- [CVE-2022-29824](#)

1.30. LOGGING 5.3.8

This release includes [RHBA-2022:5010-OpenShift Logging Bug Fix Release 5.3.8](#)

1.30.1. Bug fixes

(None.)

1.30.2. CVEs

Example 1.7. Click to expand CVEs

- [CVE-2018-25032](#)
- [CVE-2020-0404](#)
- [CVE-2020-4788](#)
- [CVE-2020-13974](#)
- [CVE-2020-19131](#)
- [CVE-2020-27820](#)
- [CVE-2021-0941](#)
- [CVE-2021-3612](#)
- [CVE-2021-3634](#)
- [CVE-2021-3669](#)
- [CVE-2021-3737](#)
- [CVE-2021-3743](#)

- [CVE-2021-3744](#)
- [CVE-2021-3752](#)
- [CVE-2021-3759](#)
- [CVE-2021-3764](#)
- [CVE-2021-3772](#)
- [CVE-2021-3773](#)
- [CVE-2021-4002](#)
- [CVE-2021-4037](#)
- [CVE-2021-4083](#)
- [CVE-2021-4157](#)
- [CVE-2021-4189](#)
- [CVE-2021-4197](#)
- [CVE-2021-4203](#)
- [CVE-2021-20322](#)
- [CVE-2021-21781](#)
- [CVE-2021-23222](#)
- [CVE-2021-26401](#)
- [CVE-2021-29154](#)
- [CVE-2021-37159](#)
- [CVE-2021-41617](#)
- [CVE-2021-41864](#)
- [CVE-2021-42739](#)
- [CVE-2021-43056](#)
- [CVE-2021-43389](#)
- [CVE-2021-43976](#)
- [CVE-2021-44733](#)
- [CVE-2021-45485](#)
- [CVE-2021-45486](#)
- [CVE-2022-0001](#)

- [CVE-2022-0002](#)
- [CVE-2022-0286](#)
- [CVE-2022-0322](#)
- [CVE-2022-1011](#)
- [CVE-2022-1271](#)

1.31. OPENSIFT LOGGING 5.3.7

This release includes [RHSA-2022:2217 OpenShift Logging Bug Fix Release 5.3.7](#)

1.31.1. Bug fixes

- Before this update, Linux audit log time parsing relied on an ordinal position of key/value pair. This update changes the parsing to utilize a regex to find the time entry. ([LOG-2322](#))
- Before this update, some log forwarder outputs could re-order logs with the same time-stamp. With this update, a sequence number has been added to the log record to order entries that have matching timestamps. ([LOG-2334](#))
- Before this update, clusters with a large number of namespaces caused Elasticsearch to stop serving requests because the list of namespaces reached the maximum header size limit. With this update, headers only include a list of namespace names, resolving the issue. ([LOG-2450](#))
- Before this update, **system:serviceaccount:openshift-monitoring:prometheus-k8s** had cluster level privileges as a **clusterrole** and **clusterrolebinding**. This update restricts the **serviceaccount** to the **openshift-logging** namespace with a role and rolebinding. ([LOG-2481](#))

1.31.2. CVEs

Example 1.8. Click to expand CVEs

- [CVE-2018-25032](#)
- [CVE-2021-4028](#)
- [CVE-2021-37136](#)
- [CVE-2021-37137](#)
- [CVE-2021-43797](#)
- [CVE-2022-0759](#)
- [CVE-2022-0778](#)
- [CVE-2022-1154](#)
- [CVE-2022-1271](#)
- [CVE-2022-21426](#)

- [CVE-2022-21434](#)
- [CVE-2022-21443](#)
- [CVE-2022-21476](#)
- [CVE-2022-21496](#)
- [CVE-2022-21698](#)
- [CVE-2022-25636](#)

1.32. OPENSIFT LOGGING 5.3.6

This release includes [RHBA-2022:1377 OpenShift Logging Bug Fix Release 5.3.6](#)

1.32.1. Bug fixes

- Before this update, defining a toleration with no key and the existing Operator caused the Operator to be unable to complete an upgrade. With this update, this toleration no longer blocks the upgrade from completing. ([LOG-2126](#))
- Before this change, it was possible for the collector to generate a warning where the chunk byte limit was exceeding an emitted event. With this change, you can tune the readline limit to resolve the issue as advised by the upstream documentation. ([LOG-2380](#))

1.33. OPENSIFT LOGGING 5.3.5

This release includes [RHSA-2022:0721 OpenShift Logging Bug Fix Release 5.3.5](#)

1.33.1. Bug fixes

- Before this update, if you removed OpenShift Logging from OpenShift Container Platform, the web console continued displaying a link to the **Logging** page. With this update, removing or uninstalling OpenShift Logging also removes that link. ([LOG-2182](#))

1.33.2. CVEs

Example 1.9. Click to expand CVEs

- [CVE-2020-28491](#)
- [CVE-2021-3521](#)
- [CVE-2021-3872](#)
- [CVE-2021-3984](#)
- [CVE-2021-4019](#)
- [CVE-2021-4122](#)
- [CVE-2021-4192](#)

- [CVE-2021-4193](#)
- [CVE-2022-0552](#)

1.34. OPENSIFT LOGGING 5.3.4

This release includes [RHBA-2022:0411 OpenShift Logging Bug Fix Release 5.3.4](#)

1.34.1. Bug fixes

- Before this update, changes to the metrics dashboards had not yet been deployed because the **cluster-logging-operator** did not correctly compare existing and desired config maps that contained the dashboard. This update fixes the logic by adding a unique hash value to the object labels. ([LOG-2066](#))
- Before this update, Elasticsearch pods failed to start after updating with FIPS enabled. With this update, Elasticsearch pods start successfully. ([LOG-1974](#))
- Before this update, elasticsearch generated the error "Unable to create PersistentVolumeClaim due to forbidden: exceeded quota: infra-storage-quota." if the PVC already existed. With this update, elasticsearch checks for existing PVCs, resolving the issue. ([LOG-2127](#))

1.34.2. CVEs

Example 1.10. Click to expand CVEs

- [CVE-2021-3521](#)
- [CVE-2021-3872](#)
- [CVE-2021-3984](#)
- [CVE-2021-4019](#)
- [CVE-2021-4122](#)
- [CVE-2021-4155](#)
- [CVE-2021-4192](#)
- [CVE-2021-4193](#)
- [CVE-2022-0185](#)
- [CVE-2022-21248](#)
- [CVE-2022-21277](#)
- [CVE-2022-21282](#)
- [CVE-2022-21283](#)
- [CVE-2022-21291](#)
- [CVE-2022-21293](#)

- [CVE-2022-21294](#)
- [CVE-2022-21296](#)
- [CVE-2022-21299](#)
- [CVE-2022-21305](#)
- [CVE-2022-21340](#)
- [CVE-2022-21341](#)
- [CVE-2022-21360](#)
- [CVE-2022-21365](#)
- [CVE-2022-21366](#)

1.35. OPENSIFT LOGGING 5.3.3

This release includes [RHSA-2022:0227 OpenShift Logging Bug Fix Release 5.3.3](#)

1.35.1. Bug fixes

- Before this update, changes to the metrics dashboards had not yet been deployed because the cluster-logging-operator did not correctly compare existing and desired configmaps containing the dashboard. This update fixes the logic by adding a dashboard unique hash value to the object labels. ([LOG-2066](#))
- This update changes the log4j dependency to 2.17.1 to resolve [CVE-2021-44832](#). ([LOG-2102](#))

1.35.2. CVEs

Example 1.11. Click to expand CVEs

- [CVE-2021-27292](#)
 - [BZ-1940613](#)
- [CVE-2021-44832](#)
 - [BZ-2035951](#)

1.36. OPENSIFT LOGGING 5.3.2

This release includes [RHSA-2022:0044 OpenShift Logging Bug Fix Release 5.3.2](#)

1.36.1. Bug fixes

- Before this update, Elasticsearch rejected logs from the Event Router due to a parsing error. This update changes the data model to resolve the parsing error. However, as a result, previous indices might cause warnings or errors within Kibana. The

kubernetes.event.metadata.resourceVersion field causes errors until existing indices are removed or reindexed. If this field is not used in Kibana, you can ignore the error messages. If you have a retention policy that deletes old indices, the policy eventually removes the old indices and stops the error messages. Otherwise, manually reindex to stop the error messages. ([LOG-2087](#))

- Before this update, the OpenShift Logging Dashboard displayed the wrong pod namespace in the table that displays top producing and collected containers over the last 24 hours. With this update, the OpenShift Logging Dashboard displays the correct pod namespace. ([LOG-2051](#))
- Before this update, if **outputDefaults.elasticsearch.structuredTypeKey** in the **ClusterLogForwarder** custom resource (CR) instance did not have a structured key, the CR replaced the output secret with the default secret used to communicate to the default log store. With this update, the defined output secret is correctly used. ([LOG-2046](#))

1.36.2. CVEs

Example 1.12. Click to expand CVEs

- [CVE-2020-36327](#)
 - [BZ-1958999](#)
- [CVE-2021-45105](#)
 - [BZ-2034067](#)
- [CVE-2021-3712](#)
- [CVE-2021-20321](#)
- [CVE-2021-42574](#)

1.37. OPENSIFT LOGGING 5.3.1

This release includes [RHSA-2021:5129 OpenShift Logging Bug Fix Release 5.3.1](#)

1.37.1. Bug fixes

- Before this update, the Fluentd container image included builder tools that were unnecessary at run time. This update removes those tools from the image. ([LOG-1998](#))
- Before this update, the Logging dashboard displayed an empty CPU graph because of a reference to an invalid metric. With this update, the Logging dashboard displays CPU graphs correctly. ([LOG-1925](#))
- Before this update, the Elasticsearch Prometheus exporter plugin compiled index-level metrics using a high-cost query that impacted the Elasticsearch node performance. This update implements a lower-cost query that improves performance. ([LOG-1897](#))

1.37.2. CVEs

Example 1.13. Click to expand CVEs

- [CVE-2021-21409](#)
 - [BZ-1944888](#)
- [CVE-2021-37136](#)
 - [BZ-2004133](#)
- [CVE-2021-37137](#)
 - [BZ-2004135](#)
- [CVE-2021-44228](#)
 - [BZ-2030932](#)
- [CVE-2018-25009](#)
- [CVE-2018-25010](#)
- [CVE-2018-25012](#)
- [CVE-2018-25013](#)
- [CVE-2018-25014](#)
- [CVE-2019-5827](#)
- [CVE-2019-13750](#)
- [CVE-2019-13751](#)
- [CVE-2019-17594](#)
- [CVE-2019-17595](#)
- [CVE-2019-18218](#)
- [CVE-2019-19603](#)
- [CVE-2019-20838](#)
- [CVE-2020-12762](#)
- [CVE-2020-13435](#)
- [CVE-2020-14145](#)
- [CVE-2020-14155](#)
- [CVE-2020-16135](#)
- [CVE-2020-17541](#)
- [CVE-2020-24370](#)
- [CVE-2020-35521](#)
- [CVE-2020-35522](#)

- [CVE-2020-35523](#)
- [CVE-2020-35524](#)
- [CVE-2020-36330](#)
- [CVE-2020-36331](#)
- [CVE-2020-36332](#)
- [CVE-2021-3200](#)
- [CVE-2021-3426](#)
- [CVE-2021-3445](#)
- [CVE-2021-3481](#)
- [CVE-2021-3572](#)
- [CVE-2021-3580](#)
- [CVE-2021-3712](#)
- [CVE-2021-3800](#)
- [CVE-2021-20231](#)
- [CVE-2021-20232](#)
- [CVE-2021-20266](#)
- [CVE-2021-20317](#)
- [CVE-2021-22876](#)
- [CVE-2021-22898](#)
- [CVE-2021-22925](#)
- [CVE-2021-27645](#)
- [CVE-2021-28153](#)
- [CVE-2021-31535](#)
- [CVE-2021-33560](#)
- [CVE-2021-33574](#)
- [CVE-2021-35942](#)
- [CVE-2021-36084](#)
- [CVE-2021-36085](#)
- [CVE-2021-36086](#)

- [CVE-2021-36087](#)
- [CVE-2021-42574](#)
- [CVE-2021-43267](#)
- [CVE-2021-43527](#)
- [CVE-2021-45046](#)

1.38. OPENSIFT LOGGING 5.3.0

This release includes [RHSA-2021:4627 OpenShift Logging Bug Fix Release 5.3.0](#)

1.38.1. New features and enhancements

- With this update, authorization options for Log Forwarding have been expanded. Outputs may now be configured with SASL, username/password, or TLS.

1.38.2. Bug fixes

- Before this update, if you forwarded logs using the syslog protocol, serializing a ruby hash encoded key/value pairs to contain a '⇒' character and replaced tabs with "#11". This update fixes the issue so that log messages are correctly serialized as valid JSON. ([LOG-1494](#))
- Before this update, application logs were not correctly configured to forward to the proper Cloudwatch stream with multi-line error detection enabled. ([LOG-1939](#))
- Before this update, a name change of the deployed collector in the 5.3 release caused the alert 'fluentnodedown' to generate. ([LOG-1918](#))
- Before this update, a regression introduced in a prior release configuration caused the collector to flush its buffered messages before shutdown, creating a delay the termination and restart of collector Pods. With this update, fluentd no longer flushes buffers at shutdown, resolving the issue. ([LOG-1735](#))
- Before this update, a regression introduced in a prior release intentionally disabled JSON message parsing. This update re-enables JSON parsing. It also sets the log entry "level" based on the "level" field in parsed JSON message or by using regex to extract a match from a message field. ([LOG-1199](#))
- Before this update, the **ClusterLogging** custom resource (CR) applied the value of the **totalLimitSize** field to the Fluentd **total_limit_size** field, even if the required buffer space was not available. With this update, the CR applies the lesser of the two **totalLimitSize** or 'default' values to the Fluentd **total_limit_size** field, resolving the issue. ([LOG-1776](#))

1.38.3. Known issues

- If you forward logs to an external Elasticsearch server and then change a configured value in the pipeline secret, such as the username and password, the Fluentd forwarder loads the new secret but uses the old value to connect to an external Elasticsearch server. This issue happens because the Red Hat OpenShift Logging Operator does not currently monitor secrets for content changes. ([LOG-1652](#))

As a workaround, if you change the secret, you can force the Fluentd pods to redeploy by entering:

```
$ oc delete pod -l component=collector
```

1.38.4. Deprecated and removed features

Some features available in previous releases have been deprecated or removed.

Deprecated functionality is still included in OpenShift Logging and continues to be supported; however, it will be removed in a future release of this product and is not recommended for new deployments.

1.38.4.1. Forwarding logs using the legacy Fluentd and legacy syslog methods have been removed

In OpenShift Logging 5.3, the legacy methods of forwarding logs to Syslog and Fluentd are removed. Bug fixes and support are provided through the end of the OpenShift Logging 5.2 life cycle. After which, no new feature enhancements are made.

Instead, use the following non-legacy methods:

- [Forwarding logs using the Fluentd forward protocol](#)
- [Forwarding logs using the syslog protocol](#)

1.38.4.2. Configuration mechanisms for legacy forwarding methods have been removed

In OpenShift Logging 5.3, the legacy configuration mechanism for log forwarding is removed: You cannot forward logs using the legacy Fluentd method and legacy Syslog method. Use the standard log forwarding methods instead.

1.38.5. CVEs

Example 1.14. Click to expand CVEs

- [CVE-2018-20673](#)
- [CVE-2018-25009](#)
- [CVE-2018-25010](#)
- [CVE-2018-25012](#)
- [CVE-2018-25013](#)
- [CVE-2018-25014](#)
- [CVE-2019-5827](#)
- [CVE-2019-13750](#)
- [CVE-2019-13751](#)
- [CVE-2019-14615](#)

- [CVE-2019-17594](#)
- [CVE-2019-17595](#)
- [CVE-2019-18218](#)
- [CVE-2019-19603](#)
- [CVE-2019-20838](#)
- [CVE-2020-0427](#)
- [CVE-2020-10001](#)
- [CVE-2020-12762](#)
- [CVE-2020-13435](#)
- [CVE-2020-14145](#)
- [CVE-2020-14155](#)
- [CVE-2020-16135](#)
- [CVE-2020-17541](#)
- [CVE-2020-24370](#)
- [CVE-2020-24502](#)
- [CVE-2020-24503](#)
- [CVE-2020-24504](#)
- [CVE-2020-24586](#)
- [CVE-2020-24587](#)
- [CVE-2020-24588](#)
- [CVE-2020-26139](#)
- [CVE-2020-26140](#)
- [CVE-2020-26141](#)
- [CVE-2020-26143](#)
- [CVE-2020-26144](#)
- [CVE-2020-26145](#)
- [CVE-2020-26146](#)
- [CVE-2020-26147](#)
- [CVE-2020-27777](#)

- [CVE-2020-29368](#)
- [CVE-2020-29660](#)
- [CVE-2020-35448](#)
- [CVE-2020-35521](#)
- [CVE-2020-35522](#)
- [CVE-2020-35523](#)
- [CVE-2020-35524](#)
- [CVE-2020-36158](#)
- [CVE-2020-36312](#)
- [CVE-2020-36330](#)
- [CVE-2020-36331](#)
- [CVE-2020-36332](#)
- [CVE-2020-36386](#)
- [CVE-2021-0129](#)
- [CVE-2021-3200](#)
- [CVE-2021-3348](#)
- [CVE-2021-3426](#)
- [CVE-2021-3445](#)
- [CVE-2021-3481](#)
- [CVE-2021-3487](#)
- [CVE-2021-3489](#)
- [CVE-2021-3564](#)
- [CVE-2021-3572](#)
- [CVE-2021-3573](#)
- [CVE-2021-3580](#)
- [CVE-2021-3600](#)
- [CVE-2021-3635](#)
- [CVE-2021-3659](#)
- [CVE-2021-3679](#)

- [CVE-2021-3732](#)
- [CVE-2021-3778](#)
- [CVE-2021-3796](#)
- [CVE-2021-3800](#)
- [CVE-2021-20194](#)
- [CVE-2021-20197](#)
- [CVE-2021-20231](#)
- [CVE-2021-20232](#)
- [CVE-2021-20239](#)
- [CVE-2021-20266](#)
- [CVE-2021-20284](#)
- [CVE-2021-22876](#)
- [CVE-2021-22898](#)
- [CVE-2021-22925](#)
- [CVE-2021-23133](#)
- [CVE-2021-23840](#)
- [CVE-2021-23841](#)
- [CVE-2021-27645](#)
- [CVE-2021-28153](#)
- [CVE-2021-28950](#)
- [CVE-2021-28971](#)
- [CVE-2021-29155](#)
- [ICVE-2021-29646](#)
- [CVE-2021-29650](#)
- [CVE-2021-31440](#)
- [CVE-2021-31535](#)
- [CVE-2021-31829](#)
- [CVE-2021-31916](#)
- [CVE-2021-33033](#)

- [CVE-2021-33194](#)
- [CVE-2021-33200](#)
- [CVE-2021-33560](#)
- [CVE-2021-33574](#)
- [CVE-2021-35942](#)
- [CVE-2021-36084](#)
- [CVE-2021-36085](#)
- [CVE-2021-36086](#)
- [CVE-2021-36087](#)
- [CVE-2021-42574](#)

1.39. LOGGING 5.2.13

This release includes [RHSA-2022:5909-OpenShift Logging Bug Fix Release 5.2.13](#).

1.39.1. Bug fixes

- [BZ-2100495](#)

1.39.2. CVEs

Example 1.15. Click to expand CVEs

- [CVE-2021-38561](#)
- [CVE-2021-40528](#)
- [CVE-2022-1271](#)
- [CVE-2022-1621](#)
- [CVE-2022-1629](#)
- [CVE-2022-21540](#)
- [CVE-2022-21541](#)
- [CVE-2022-22576](#)
- [CVE-2022-25313](#)
- [CVE-2022-25314](#)
- [CVE-2022-27774](#)
- [CVE-2022-27776](#)

- [CVE-2022-27782](#)
- [CVE-2022-29824](#)
- [CVE-2022-34169](#)

1.40. LOGGING 5.2.12

This release includes [RHBA-2022:5558-OpenShift Logging Bug Fix Release 5.2.12](#).

1.40.1. Bug fixes

None.

1.40.2. CVEs

Example 1.16. Click to expand CVEs

- [CVE-2020-28915](#)
- [CVE-2021-40528](#)
- [CVE-2022-1271](#)
- [CVE-2022-1621](#)
- [CVE-2022-1629](#)
- [CVE-2022-22576](#)
- [CVE-2022-25313](#)
- [CVE-2022-25314](#)
- [CVE-2022-26691](#)
- [CVE-2022-27666](#)
- [CVE-2022-27774](#)
- [CVE-2022-27776](#)
- [CVE-2022-27782](#)
- [CVE-2022-29824](#)

1.41. LOGGING 5.2.11

This release includes [RHBA-2022:5012-OpenShift Logging Bug Fix Release 5.2.11](#)

1.41.1. Bug fixes

- Before this update, clusters configured to perform CloudWatch forwarding wrote rejected log files to temporary storage, causing cluster instability over time. With this update, chunk backup for CloudWatch has been disabled, resolving the issue. ([LOG-2635](#))

1.41.2. CVEs

Example 1.17. Click to expand CVEs

- [CVE-2018-25032](#)
- [CVE-2020-0404](#)
- [CVE-2020-4788](#)
- [CVE-2020-13974](#)
- [CVE-2020-19131](#)
- [CVE-2020-27820](#)
- [CVE-2021-0941](#)
- [CVE-2021-3612](#)
- [CVE-2021-3634](#)
- [CVE-2021-3669](#)
- [CVE-2021-3737](#)
- [CVE-2021-3743](#)
- [CVE-2021-3744](#)
- [CVE-2021-3752](#)
- [CVE-2021-3759](#)
- [CVE-2021-3764](#)
- [CVE-2021-3772](#)
- [CVE-2021-3773](#)
- [CVE-2021-4002](#)
- [CVE-2021-4037](#)
- [CVE-2021-4083](#)
- [CVE-2021-4157](#)
- [CVE-2021-4189](#)
- [CVE-2021-4197](#)
- [CVE-2021-4203](#)

- [CVE-2021-20322](#)
- [CVE-2021-21781](#)
- [CVE-2021-23222](#)
- [CVE-2021-26401](#)
- [CVE-2021-29154](#)
- [CVE-2021-37159](#)
- [CVE-2021-41617](#)
- [CVE-2021-41864](#)
- [CVE-2021-42739](#)
- [CVE-2021-43056](#)
- [CVE-2021-43389](#)
- [CVE-2021-43976](#)
- [CVE-2021-44733](#)
- [CVE-2021-45485](#)
- [CVE-2021-45486](#)
- [CVE-2022-0001](#)
- [CVE-2022-0002](#)
- [CVE-2022-0286](#)
- [CVE-2022-0322](#)
- [CVE-2022-1011](#)
- [CVE-2022-1271](#)

1.42. OPENSIFT LOGGING 5.2.10

This release includes [OpenShift Logging Bug Fix Release 5.2.10](#)

1.42.1. Bug fixes

- Before this update some log forwarder outputs could re-order logs with the same time-stamp. With this update, a sequence number has been added to the log record to order entries that have matching timestamps.([LOG-2335](#))
- Before this update, clusters with a large number of namespaces caused Elasticsearch to stop serving requests because the list of namespaces reached the maximum header size limit. With this update, headers only include a list of namespace names, resolving the issue. ([LOG-2475](#))

- Before this update, **system:serviceaccount:openshift-monitoring:prometheus-k8s** had cluster level privileges as a **clusterrole** and **clusterrolebinding**. This update restricts the **serviceaccount** to the **openshift-logging** namespace with a role and rolebinding. ([LOG-2480](#))
- Before this update, the **cluster-logging-operator** utilized cluster scoped roles and bindings to establish permissions for the Prometheus service account to scrape metrics. These permissions were only created when deploying the Operator using the console interface and were missing when the Operator was deployed from the command line. This fixes the issue by making this role and binding namespace scoped. ([LOG-1972](#))

1.42.2. CVEs

Example 1.18. Click to expand CVEs

- [CVE-2018-25032](#)
- [CVE-2021-4028](#)
- [CVE-2021-37136](#)
- [CVE-2021-37137](#)
- [CVE-2021-43797](#)
- [CVE-2022-0778](#)
- [CVE-2022-1154](#)
- [CVE-2022-1271](#)
- [CVE-2022-21426](#)
- [CVE-2022-21434](#)
- [CVE-2022-21443](#)
- [CVE-2022-21476](#)
- [CVE-2022-21496](#)
- [CVE-2022-21698](#)
- [CVE-2022-25636](#)

1.43. OPENSIFT LOGGING 5.2.9

This release includes [RHBA-2022:1375 OpenShift Logging Bug Fix Release 5.2.9](#)]

1.43.1. Bug fixes

- Before this update, defining a toleration with no key and the existing Operator caused the Operator to be unable to complete an upgrade. With this update, this toleration no longer blocks the upgrade from completing. ([LOG-2304](#))

1.44. OPENSIFT LOGGING 5.2.8

This release includes [RHSA-2022:0728 OpenShift Logging Bug Fix Release 5.2.8](#)

1.44.1. Bug fixes

- Before this update, if you removed OpenShift Logging from OpenShift Container Platform, the web console continued displaying a link to the **Logging** page. With this update, removing or uninstalling OpenShift Logging also removes that link. ([LOG-2180](#))

1.44.2. CVEs

Example 1.19. Click to expand CVEs

- [CVE-2020-28491](#)
 - [BZ-1930423](#)
- [CVE-2022-0552](#)
 - [BG-2052539](#)

1.45. OPENSIFT LOGGING 5.2.7

This release includes [RHBA-2022:0478 OpenShift Logging Bug Fix Release 5.2.7](#)

1.45.1. Bug fixes

- Before this update, Elasticsearch pods with FIPS enabled failed to start after updating. With this update, Elasticsearch pods start successfully. ([LOG-2000](#))
- Before this update, if a persistent volume claim (PVC) already existed, Elasticsearch generated an error, "Unable to create PersistentVolumeClaim due to forbidden: exceeded quota: infra-storage-quota." With this update, Elasticsearch checks for existing PVCs, resolving the issue. ([LOG-2118](#))

1.45.2. CVEs

Example 1.20. Click to expand CVEs

- [CVE-2021-3521](#)
- [CVE-2021-3872](#)
- [CVE-2021-3984](#)
- [CVE-2021-4019](#)
- [CVE-2021-4122](#)
- [CVE-2021-4155](#)
- [CVE-2021-4192](#)

- [CVE-2021-4193](#)
- [CVE-2022-0185](#)

1.46. OPENSIFT LOGGING 5.2.6

This release includes [RHSA-2022:0230 OpenShift Logging Bug Fix Release 5.2.6](#)

1.46.1. Bug fixes

- Before this update, the release did not include a filter change which caused Fluentd to crash. With this update, the missing filter has been corrected. ([LOG-2104](#))
- This update changes the log4j dependency to 2.17.1 to resolve [CVE-2021-44832](#).([LOG-2101](#))

1.46.2. CVEs

Example 1.21. Click to expand CVEs

- [CVE-2021-27292](#)
 - [BZ-1940613](#)
- [CVE-2021-44832](#)
 - [BZ-2035951](#)

1.47. OPENSIFT LOGGING 5.2.5

This release includes [RHSA-2022:0043 OpenShift Logging Bug Fix Release 5.2.5](#)

1.47.1. Bug fixes

- Before this update, Elasticsearch rejected logs from the Event Router due to a parsing error. This update changes the data model to resolve the parsing error. However, as a result, previous indices might cause warnings or errors within Kibana. The **kubernetes.event.metadata.resourceVersion** field causes errors until existing indices are removed or reindexed. If this field is not used in Kibana, you can ignore the error messages. If you have a retention policy that deletes old indices, the policy eventually removes the old indices and stops the error messages. Otherwise, manually reindex to stop the error messages. ([LOG-2087](#))

1.47.2. CVEs

Example 1.22. Click to expand CVEs

- [CVE-2021-3712](#)
- [CVE-2021-20321](#)
- [CVE-2021-42574](#)

- [CVE-2021-45105](#)

1.48. OPENSIFT LOGGING 5.2.4

This release includes [RHSA-2021:5127 OpenShift Logging Bug Fix Release 5.2.4](#)

1.48.1. Bug fixes

- Before this update, records shipped via syslog would serialize a ruby hash encoding key/value pairs to contain a '⇒' character, as well as replace tabs with "#11". This update serializes the message correctly as proper JSON. ([LOG-1775](#))
- Before this update, the Elasticsearch Prometheus exporter plugin compiled index-level metrics using a high-cost query that impacted the Elasticsearch node performance. This update implements a lower-cost query that improves performance. ([LOG-1970](#))
- Before this update, Elasticsearch sometimes rejected messages when Log Forwarding was configured with multiple outputs. This happened because configuring one of the outputs modified message content to be a single message. With this update, Log Forwarding duplicates the messages for each output so that output-specific processing does not affect the other outputs. ([LOG-1824](#))

1.48.2. CVEs

Example 1.23. Click to expand CVEs

- [CVE-2018-25009](#)
- [CVE-2018-25010](#)
- [CVE-2018-25012](#)
- [CVE-2018-25013](#)
- [CVE-2018-25014](#)
- [CVE-2019-5827](#)
- [CVE-2019-13750](#)
- [CVE-2019-13751](#)
- [CVE-2019-17594](#)
- [CVE-2019-17595](#)
- [CVE-2019-18218](#)
- [CVE-2019-19603](#)
- [CVE-2019-20838](#)
- [CVE-2020-12762](#)
- [CVE-2020-13435](#)

- [CVE-2020-14145](#)
- [CVE-2020-14155](#)
- [CVE-2020-16135](#)
- [CVE-2020-17541](#)
- [CVE-2020-24370](#)
- [CVE-2020-35521](#)
- [CVE-2020-35522](#)
- [CVE-2020-35523](#)
- [CVE-2020-35524](#)
- [CVE-2020-36330](#)
- [CVE-2020-36331](#)
- [CVE-2020-36332](#)
- [CVE-2021-3200](#)
- [CVE-2021-3426](#)
- [CVE-2021-3445](#)
- [CVE-2021-3481](#)
- [CVE-2021-3572](#)
- [CVE-2021-3580](#)
- [CVE-2021-3712](#)
- [CVE-2021-3800](#)
- [CVE-2021-20231](#)
- [CVE-2021-20232](#)
- [CVE-2021-20266](#)
- [CVE-2021-20317](#)
- [CVE-2021-21409](#)
- [CVE-2021-22876](#)
- [CVE-2021-22898](#)
- [CVE-2021-22925](#)
- [CVE-2021-27645](#)

- [CVE-2021-28153](#)
- [CVE-2021-31535](#)
- [CVE-2021-33560](#)
- [CVE-2021-33574](#)
- [CVE-2021-35942](#)
- [CVE-2021-36084](#)
- [CVE-2021-36085](#)
- [CVE-2021-36086](#)
- [CVE-2021-36087](#)
- [CVE-2021-37136](#)
- [CVE-2021-37137](#)
- [CVE-2021-42574](#)
- [CVE-2021-43267](#)
- [CVE-2021-43527](#)
- [CVE-2021-44228](#)
- [CVE-2021-45046](#)

1.49. OPENSIFT LOGGING 5.2.3

This release includes [RHSA-2021:4032 OpenShift Logging Bug Fix Release 5.2.3](#)

1.49.1. Bug fixes

- Before this update, some alerts did not include a namespace label. This omission does not comply with the OpenShift Monitoring Team’s guidelines for writing alerting rules in OpenShift Container Platform. With this update, all the alerts in Elasticsearch Operator include a namespace label and follow all the guidelines for writing alerting rules in OpenShift Container Platform. ([LOG-1857](#))
- Before this update, a regression introduced in a prior release intentionally disabled JSON message parsing. This update re-enables JSON parsing. It also sets the log entry **level** based on the **level** field in parsed JSON message or by using regex to extract a match from a message field. ([LOG-1759](#))

1.49.2. CVEs

Example 1.24. Click to expand CVEs

- [CVE-2021-23369](#)

- BZ-1948761
- CVE-2021-23383
 - BZ-1956688
- CVE-2018-20673
- CVE-2019-5827
- CVE-2019-13750
- CVE-2019-13751
- CVE-2019-17594
- CVE-2019-17595
- CVE-2019-18218
- CVE-2019-19603
- CVE-2019-20838
- CVE-2020-12762
- CVE-2020-13435
- CVE-2020-14155
- CVE-2020-16135
- CVE-2020-24370
- CVE-2021-3200
- CVE-2021-3426
- CVE-2021-3445
- CVE-2021-3572
- CVE-2021-3580
- CVE-2021-3778
- CVE-2021-3796
- CVE-2021-3800
- CVE-2021-20231
- CVE-2021-20232
- CVE-2021-20266
- CVE-2021-22876

- [CVE-2021-22898](#)
- [CVE-2021-22925](#)
- [CVE-2021-23840](#)
- [CVE-2021-23841](#)
- [CVE-2021-27645](#)
- [CVE-2021-28153](#)
- [CVE-2021-33560](#)
- [CVE-2021-33574](#)
- [CVE-2021-35942](#)
- [CVE-2021-36084](#)
- [CVE-2021-36085](#)
- [CVE-2021-36086](#)
- [CVE-2021-36087](#)

1.50. OPENSIFT LOGGING 5.2.2

This release includes [RHBA-2021:3747 OpenShift Logging Bug Fix Release 5.2.2](#)

1.50.1. Bug fixes

- Before this update, the **ClusterLogging** custom resource (CR) applied the value of the **totalLimitSize** field to the Fluentd **total_limit_size** field, even if the required buffer space was not available. With this update, the CR applies the lesser of the two **totalLimitSize** or 'default' values to the Fluentd **total_limit_size** field, resolving the issue. ([LOG-1738](#))
- Before this update, a regression introduced in a prior release configuration caused the collector to flush its buffered messages before shutdown, creating a delay to the termination and restart of collector pods. With this update, Fluentd no longer flushes buffers at shutdown, resolving the issue. ([LOG-1739](#))
- Before this update, an issue in the bundle manifests prevented installation of the Elasticsearch Operator through OLM on OpenShift Container Platform 4.9. With this update, a correction to bundle manifests re-enables installation and upgrade in 4.9. ([LOG-1780](#))

1.50.2. CVEs

Example 1.25. Click to expand CVEs

- [CVE-2020-25648](#)
- [CVE-2021-22922](#)
- [CVE-2021-22923](#)

- [CVE-2021-22924](#)
- [CVE-2021-36222](#)
- [CVE-2021-37576](#)
- [CVE-2021-37750](#)
- [CVE-2021-38201](#)

1.51. OPENSIFT LOGGING 5.2.1

This release includes [RHBA-2021:3550 OpenShift Logging Bug Fix Release 5.2.1](#)

1.51.1. Bug fixes

- Before this update, due to an issue in the release pipeline scripts, the value of the **olm.skipRange** field remained unchanged at **5.2.0** instead of reflecting the current release number. This update fixes the pipeline scripts to update the value of this field when the release numbers change. ([LOG-1743](#))

1.51.2. CVEs

(None)

1.52. OPENSIFT LOGGING 5.2.0

This release includes [RHBA-2021:3393 OpenShift Logging Bug Fix Release 5.2.0](#)

1.52.1. New features and enhancements

- With this update, you can forward log data to Amazon CloudWatch, which provides application and infrastructure monitoring. For more information, see [Forwarding logs to Amazon CloudWatch](#). ([LOG-1173](#))
- With this update, you can forward log data to Loki, a horizontally scalable, highly available, multi-tenant log aggregation system. For more information, see [Forwarding logs to Loki](#). ([LOG-684](#))
- With this update, if you use the Fluentd forward protocol to forward log data over a TLS-encrypted connection, now you can use a password-encrypted private key file and specify the passphrase in the Cluster Log Forwarder configuration. For more information, see [Forwarding logs using the Fluentd forward protocol](#). ([LOG-1525](#))
- This enhancement enables you to use a username and password to authenticate a log forwarding connection to an external Elasticsearch instance. For example, if you cannot use mutual TLS (mTLS) because a third-party operates the Elasticsearch instance, you can use HTTP or HTTPS and set a secret that contains the username and password. For more information, see [Forwarding logs to an external Elasticsearch instance](#). ([LOG-1022](#))
- With this update, you can collect OVN network policy audit logs for forwarding to a logging server. ([LOG-1526](#))

- By default, the data model introduced in OpenShift Container Platform 4.5 gave logs from different namespaces a single index in common. This change made it harder to see which namespaces produced the most logs.
The current release adds namespace metrics to the **Logging** dashboard in the OpenShift Container Platform console. With these metrics, you can see which namespaces produce logs and how many logs each namespace produces for a given timestamp.

To see these metrics, open the **Administrator** perspective in the OpenShift Container Platform web console, and navigate to **Observe** → **Dashboards** → **Logging/Elasticsearch**. ([LOG-1680](#))

- The current release, OpenShift Logging 5.2, enables two new metrics: For a given timestamp or duration, you can see the total logs produced or logged by individual containers, and the total logs collected by the collector. These metrics are labeled by namespace, pod, and container name so that you can see how many logs each namespace and pod collects and produces. ([LOG-1213](#))

1.52.2. Bug fixes

- Before this update, when the OpenShift Elasticsearch Operator created index management cronjobs, it added the **POLICY_MAPPING** environment variable twice, which caused the apiserver to report the duplication. This update fixes the issue so that the **POLICY_MAPPING** environment variable is set only once per cronjob, and there is no duplication for the apiserver to report. ([LOG-1130](#))
- Before this update, suspending an Elasticsearch cluster to zero nodes did not suspend the index-management cronjobs, which put these cronjobs into maximum backoff. Then, after unsuspending the Elasticsearch cluster, these cronjobs stayed halted due to maximum backoff reached. This update resolves the issue by suspending the cronjobs and the cluster. ([LOG-1268](#))
- Before this update, in the **Logging** dashboard in the OpenShift Container Platform console, the list of top 10 log-producing containers was missing the "chart namespace" label and provided the incorrect metric name, **fluentd_input_status_total_bytes_logged**. With this update, the chart shows the namespace label and the correct metric name, **log_logged_bytes_total**. ([LOG-1271](#))
- Before this update, if an index management cronjob terminated with an error, it did not report the error exit code: instead, its job status was "complete." This update resolves the issue by reporting the error exit codes of index management cronjobs that terminate with errors. ([LOG-1273](#))
- The **priorityclasses.v1beta1.scheduling.k8s.io** was removed in 1.22 and replaced by **priorityclasses.v1.scheduling.k8s.io** (**v1beta1** was replaced by **v1**). Before this update, **APIRemovedInNextReleaseInUse** alerts were generated for **priorityclasses** because **v1beta1** was still present. This update resolves the issue by replacing **v1beta1** with **v1**. The alert is no longer generated. ([LOG-1385](#))
- Previously, the OpenShift Elasticsearch Operator and Red Hat OpenShift Logging Operator did not have the annotation that was required for them to appear in the OpenShift Container Platform web console list of Operators that can run in a disconnected environment. This update adds the **operators.openshift.io/infrastructure-features: ["Disconnected"]** annotation to these two Operators so that they appear in the list of Operators that run in disconnected environments. ([LOG-1420](#))
- Before this update, Red Hat OpenShift Logging Operator pods were scheduled on CPU cores that were reserved for customer workloads on performance-optimized single-node clusters. With this update, cluster logging Operator pods are scheduled on the correct CPU cores. ([LOG-1440](#))

- Before this update, some log entries had unrecognized UTF-8 bytes, which caused Elasticsearch to reject the messages and block the entire buffered payload. With this update, rejected payloads drop the invalid log entries and resubmit the remaining entries to resolve the issue. ([LOG-1499](#))
- Before this update, the **kibana-proxy** pod sometimes entered the **CrashLoopBackoff** state and logged the following message **Invalid configuration: cookie_secret must be 16, 24, or 32 bytes to create an AES cipher when pass_access_token == true or cookie_refresh != 0, but is 29 bytes**. The exact actual number of bytes could vary. With this update, the generation of the Kibana session secret has been corrected, and the kibana-proxy pod no longer enters a **CrashLoopBackoff** state due to this error. ([LOG-1446](#))
- Before this update, the AWS CloudWatch Fluentd plugin logged its AWS API calls to the Fluentd log at all log levels, consuming additional OpenShift Container Platform node resources. With this update, the AWS CloudWatch Fluentd plugin logs AWS API calls only at the "debug" and "trace" log levels. This way, at the default "warn" log level, Fluentd does not consume extra node resources. ([LOG-1071](#))
- Before this update, the Elasticsearch OpenDistro security plugin caused user index migrations to fail. This update resolves the issue by providing a newer version of the plugin. Now, index migrations proceed without errors. ([LOG-1276](#))
- Before this update, in the **Logging** dashboard in the OpenShift Container Platform console, the list of top 10 log-producing containers lacked data points. This update resolves the issue, and the dashboard displays all data points. ([LOG-1353](#))
- Before this update, if you were tuning the performance of the Fluentd log forwarder by adjusting the **chunkLimitSize** and **totalLimitSize** values, the **Setting queued_chunks_limit_size for each buffer to** message reported values that were too low. The current update fixes this issue so that this message reports the correct values. ([LOG-1411](#))
- Before this update, the Kibana OpenDistro security plugin caused user index migrations to fail. This update resolves the issue by providing a newer version of the plugin. Now, index migrations proceed without errors. ([LOG-1558](#))
- Before this update, using a namespace input filter prevented logs in that namespace from appearing in other inputs. With this update, logs are sent to all inputs that can accept them. ([LOG-1570](#))
- Before this update, a missing license file for the **viaq/logerr** dependency caused license scanners to abort without success. With this update, the **viaq/logerr** dependency is licensed under Apache 2.0 and the license scanners run successfully. ([LOG-1590](#))
- Before this update, an incorrect brew tag for **curator5** within the **elasticsearch-operator-bundle** build pipeline caused the pull of an image pinned to a dummy SHA1. With this update, the build pipeline uses the **logging-curator5-rhel8** reference for **curator5**, enabling index management cronjobs to pull the correct image from **registry.redhat.io**. ([LOG-1624](#))
- Before this update, an issue with the **ServiceAccount** permissions caused errors such as **no permissions for [indices:admin/aliases/get]**. With this update, a permission fix resolves the issue. ([LOG-1657](#))
- Before this update, the Custom Resource Definition (CRD) for the Red Hat OpenShift Logging Operator was missing the Loki output type, which caused the admission controller to reject the **ClusterLogForwarder** custom resource object. With this update, the CRD includes Loki as an output type so that administrators can configure **ClusterLogForwarder** to send logs to a Loki server. ([LOG-1683](#))

- Before this update, OpenShift Elasticsearch Operator reconciliation of the **ServiceAccounts** overwrote third-party-owned fields that contained secrets. This issue caused memory and CPU spikes due to frequent recreation of secrets. This update resolves the issue. Now, the OpenShift Elasticsearch Operator does not overwrite third-party-owned fields. ([LOG-1714](#))
- Before this update, in the **ClusterLogging** custom resource (CR) definition, if you specified a **flush_interval** value but did not set **flush_mode** to **interval**, the Red Hat OpenShift Logging Operator generated a Fluentd configuration. However, the Fluentd collector generated an error at runtime. With this update, the Red Hat OpenShift Logging Operator validates the **ClusterLogging** CR definition and only generates the Fluentd configuration if both fields are specified. ([LOG-1723](#))

1.52.3. Known issues

- If you forward logs to an external Elasticsearch server and then change a configured value in the pipeline secret, such as the username and password, the Fluentd forwarder loads the new secret but uses the old value to connect to an external Elasticsearch server. This issue happens because the Red Hat OpenShift Logging Operator does not currently monitor secrets for content changes. ([LOG-1652](#))

As a workaround, if you change the secret, you can force the Fluentd pods to redeploy by entering:

```
$ oc delete pod -l component=collector
```

1.52.4. Deprecated and removed features

Some features available in previous releases have been deprecated or removed.

Deprecated functionality is still included in OpenShift Logging and continues to be supported; however, it will be removed in a future release of this product and is not recommended for new deployments.

1.52.5. Forwarding logs using the legacy Fluentd and legacy syslog methods have been deprecated

From OpenShift Container Platform 4.6 to the present, forwarding logs by using the following legacy methods have been deprecated and will be removed in a future release:

- Forwarding logs using the legacy Fluentd method
- Forwarding logs using the legacy syslog method

Instead, use the following non-legacy methods:

- [Forwarding logs using the Fluentd forward protocol](https://www.redhat.com/security/data/cve/CVE-2021-22922.html)
- [Forwarding logs using the syslog protocol](#)

1.52.6. CVEs

Example 1.26. Click to expand CVEs

- [CVE-2021-22922](#)

- [CVE-2021-22923](#)
- [CVE-2021-22924](#)
- [CVE-2021-32740](#)
- [CVE-2021-36222](#)
- [CVE-2021-37750](#)

CHAPTER 2. UNDERSTANDING THE LOGGING SUBSYSTEM FOR RED HAT OPENS SHIFT

As a cluster administrator, you can deploy the logging subsystem to aggregate all the logs from your OpenShift Container Platform cluster, such as node system audit logs, application container logs, and infrastructure logs. The logging subsystem aggregates these logs from throughout your cluster and stores them in a default log store. You can [use the Kibana web console to visualize log data](#) .

The logging subsystem aggregates the following types of logs:

- **application** - Container logs generated by user applications running in the cluster, except infrastructure container applications.
- **infrastructure** - Logs generated by infrastructure components running in the cluster and OpenShift Container Platform nodes, such as journal logs. Infrastructure components are pods that run in the **openshift***, **kube***, or **default** projects.
- **audit** - Logs generated by auditd, the node audit system, which are stored in the `/var/log/audit/audit.log` file, and the audit logs from the Kubernetes apiserver and the OpenShift apiserver.



NOTE

Because the internal OpenShift Container Platform Elasticsearch log store does not provide secure storage for audit logs, audit logs are not stored in the internal Elasticsearch instance by default. If you want to send the audit logs to the default internal Elasticsearch log store, for example to view the audit logs in Kibana, you must use the Log Forwarding API as described in [Forward audit logs to the log store](#) .

2.1. GLOSSARY OF COMMON TERMS FOR OPENS SHIFT CONTAINER PLATFORM LOGGING

This glossary defines common terms that are used in the OpenShift Container Platform Logging content.

annotation

You can use annotations to attach metadata to objects.

Cluster Logging Operator (CLO)

The Cluster Logging Operator provides a set of APIs to control the collection and forwarding of application, infrastructure, and audit logs.

Custom Resource (CR)

A CR is an extension of the Kubernetes API. To configure OpenShift Container Platform Logging and log forwarding, you can customize the **ClusterLogging** and the **ClusterLogForwarder** custom resources.

event router

The event router is a pod that watches OpenShift Container Platform events. It collects logs by using OpenShift Container Platform Logging.

Fluentd

Fluentd is a log collector that resides on each OpenShift Container Platform node. It gathers application, infrastructure, and audit logs and forwards them to different outputs.

garbage collection

Garbage collection is the process of cleaning up cluster resources, such as terminated containers and images that are not referenced by any running pods.

Elasticsearch

Elasticsearch is a distributed search and analytics engine. OpenShift Container Platform uses Elasticsearch as a default log store for OpenShift Container Platform Logging.

Elasticsearch Operator

Elasticsearch operator is used to run Elasticsearch cluster on top of OpenShift Container Platform. The Elasticsearch Operator provides self-service for the Elasticsearch cluster operations and is used by OpenShift Container Platform Logging.

indexing

Indexing is a data structure technique that is used to quickly locate and access data. Indexing optimizes the performance by minimizing the amount of disk access required when a query is processed.

JSON logging

OpenShift Container Platform Logging Log Forwarding API enables you to parse JSON logs into a structured object and forward them to either OpenShift Container Platform Logging-managed Elasticsearch or any other third-party system supported by the Log Forwarding API.

Kibana

Kibana is a browser-based console interface to query, discover, and visualize your Elasticsearch data through histograms, line graphs, and pie charts.

Kubernetes API server

Kubernetes API server validates and configures data for the API objects.

Labels

Labels are key-value pairs that you can use to organize and select subsets of objects, such as a pod.

Logging

With OpenShift Container Platform Logging you can aggregate application, infrastructure, and audit logs throughout your cluster. You can also store them to a default log store, forward them to third party systems, and query and visualize the stored logs in the default log store.

logging collector

A logging collector collects logs from the cluster, formats them, and forwards them to the log store or third party systems.

log store

A log store is used to store aggregated logs. You can use the default Elasticsearch log store or forward logs to external log stores. The default log store is optimized and tested for short-term storage.

log visualizer

Log visualizer is the user interface (UI) component you can use to view information such as logs, graphs, charts, and other metrics. The current implementation is Kibana.

node

A node is a worker machine in the OpenShift Container Platform cluster. A node is either a virtual machine (VM) or a physical machine.

Operators

Operators are the preferred method of packaging, deploying, and managing a Kubernetes application in an OpenShift Container Platform cluster. An Operator takes human operational knowledge and encodes it into software that is packaged and shared with customers.

pod

A pod is the smallest logical unit in Kubernetes. A pod consists of one or more containers and runs on a worker node..

Role-based access control (RBAC)

RBAC is a key security control to ensure that cluster users and workloads have access only to resources required to execute their roles.

shards

Elasticsearch organizes the log data from Fluentd into datastores, or indices, then subdivides each index into multiple pieces called shards.

taint

Taints ensure that pods are scheduled onto appropriate nodes. You can apply one or more taints on a node.

toleration

You can apply tolerations to pods. Tolerations allow the scheduler to schedule pods with matching taints.

web console

A user interface (UI) to manage OpenShift Container Platform.

2.2. ABOUT DEPLOYING THE LOGGING SUBSYSTEM FOR RED HAT OPENSIFT

OpenShift Container Platform cluster administrators can deploy the logging subsystem using the OpenShift Container Platform web console or CLI to install the OpenShift Elasticsearch Operator and Red Hat OpenShift Logging Operator. When the Operators are installed, you create a **ClusterLogging** custom resource (CR) to schedule logging subsystem pods and other resources necessary to support the logging subsystem. The Operators are responsible for deploying, upgrading, and maintaining the logging subsystem.

The **ClusterLogging** CR defines a complete logging subsystem environment that includes all the components of the logging stack to collect, store and visualize logs. The Red Hat OpenShift Logging Operator watches the logging subsystem CR and adjusts the logging deployment accordingly.

Administrators and application developers can view the logs of the projects for which they have view access.

For information, see [Configuring the log collector](#).

2.2.1. About JSON OpenShift Container Platform Logging

You can use JSON logging to configure the Log Forwarding API to parse JSON strings into a structured object. You can perform the following tasks:

- Parse JSON logs
- Configure JSON log data for Elasticsearch
- Forward JSON logs to the Elasticsearch log store

2.2.2. About collecting and storing Kubernetes events

The OpenShift Container Platform Event Router is a pod that watches Kubernetes events and logs them for collection by OpenShift Container Platform Logging. You must manually deploy the Event Router.

For information, see [About collecting and storing Kubernetes events](#).

2.2.3. About updating OpenShift Container Platform Logging

OpenShift Container Platform allows you to update OpenShift Container Platform logging. You must update the following operators while updating OpenShift Container Platform Logging:

- Elasticsearch Operator
- Cluster Logging Operator

For information, see [About updating OpenShift Container Platform Logging](#).

2.2.4. About viewing the cluster dashboard

The OpenShift Container Platform Logging dashboard contains charts that show details about your Elasticsearch instance at the cluster level. These charts help you diagnose and anticipate problems.

For information, see [About viewing the cluster dashboard](#).

2.2.5. About troubleshooting OpenShift Container Platform Logging

You can troubleshoot the logging issues by performing the following tasks:

- Viewing logging status
- Viewing the status of the log store
- Understanding logging alerts
- Collecting logging data for Red Hat Support
- Troubleshooting for critical alerts

2.2.6. About uninstalling OpenShift Container Platform Logging

You can stop log aggregation by deleting the ClusterLogging custom resource (CR). After deleting the CR, there are other cluster logging components that remain, which you can optionally remove.

For information, see [About uninstalling OpenShift Container Platform Logging](#).

2.2.7. About exporting fields

The logging system exports fields. Exported fields are present in the log records and are available for searching from Elasticsearch and Kibana.

For information, see [About exporting fields](#).

2.2.8. About logging subsystem components

The logging subsystem components include a collector deployed to each node in the OpenShift Container Platform cluster that collects all node and container logs and writes them to a log store. You can use a centralized web UI to create rich visualizations and dashboards with the aggregated data.

The major components of the logging subsystem are:

- collection - This is the component that collects logs from the cluster, formats them, and forwards them to the log store. The current implementation is Fluentd.
- log store - This is where the logs are stored. The default implementation is Elasticsearch. You can use the default Elasticsearch log store or forward logs to external log stores. The default log store is optimized and tested for short-term storage.
- visualization - This is the UI component you can use to view logs, graphs, charts, and so forth. The current implementation is Kibana.

This document might refer to log store or Elasticsearch, visualization or Kibana, collection or Fluentd, interchangeably, except where noted.

2.2.9. About the logging collector

The logging subsystem for Red Hat OpenShift collects container and node logs.

By default, the log collector uses the following sources:

- journald for all system logs
- `/var/log/containers/*.log` for all container logs

If you configure the log collector to collect audit logs, it gets them from `/var/log/audit/audit.log`.

The logging collector is a daemon set that deploys pods to each OpenShift Container Platform node. System and infrastructure logs are generated by journald log messages from the operating system, the container runtime, and OpenShift Container Platform. Application logs are generated by the CRI-O container engine. Fluentd collects the logs from these sources and forwards them internally or externally as you configure in OpenShift Container Platform.

The container runtimes provide minimal information to identify the source of log messages: project, pod name, and container ID. This information is not sufficient to uniquely identify the source of the logs. If a pod with a given name and project is deleted before the log collector begins processing its logs, information from the API server, such as labels and annotations, might not be available. There might not be a way to distinguish the log messages from a similarly named pod and project or trace the logs to their source. This limitation means that log collection and normalization are considered **best effort**.



IMPORTANT

The available container runtimes provide minimal information to identify the source of log messages and do not guarantee unique individual log messages or that these messages can be traced to their source.

For information, see [Configuring the log collector](#).

2.2.10. About the log store

By default, OpenShift Container Platform uses [Elasticsearch \(ES\)](#) to store log data. Optionally you can use the Log Forwarder API to forward logs to an external store. Several types of store are supported, including fluentd, rsyslog, kafka and others.

The logging subsystem Elasticsearch instance is optimized and tested for short term storage, approximately seven days. If you want to retain your logs over a longer term, it is recommended you move the data to a third-party storage system.

Elasticsearch organizes the log data from Fluentd into *datastores*, or *indices*, then subdivides each index into multiple pieces called *shards*, which it spreads across a set of Elasticsearch nodes in an Elasticsearch cluster. You can configure Elasticsearch to make copies of the shards, called *replicas*, which Elasticsearch also spreads across the Elasticsearch nodes. The **ClusterLogging** custom resource (CR) allows you to specify how the shards are replicated to provide data redundancy and resilience to failure. You can also specify how long the different types of logs are retained using a retention policy in the **ClusterLogging** CR.



NOTE

The number of primary shards for the index templates is equal to the number of Elasticsearch data nodes.

The Red Hat OpenShift Logging Operator and companion OpenShift Elasticsearch Operator ensure that each Elasticsearch node is deployed using a unique deployment that includes its own storage volume. You can use a **ClusterLogging** custom resource (CR) to increase the number of Elasticsearch nodes, as needed. See the [Elasticsearch documentation](#) for considerations involved in configuring storage.



NOTE

A highly-available Elasticsearch environment requires at least three Elasticsearch nodes, each on a different host.

Role-based access control (RBAC) applied on the Elasticsearch indices enables the controlled access of the logs to the developers. Administrators can access all logs and developers can access only the logs in their projects.

For information, see [Configuring the log store](#).

2.2.11. About logging visualization

OpenShift Container Platform uses Kibana to display the log data collected by Fluentd and indexed by Elasticsearch.

Kibana is a browser-based console interface to query, discover, and visualize your Elasticsearch data through histograms, line graphs, pie charts, and other visualizations.

For information, see [Configuring the log visualizer](#).

2.2.12. About event routing

The Event Router is a pod that watches OpenShift Container Platform events so they can be collected by the logging subsystem for Red Hat OpenShift. The Event Router collects events from all projects and writes them to **STDOUT**. Fluentd collects those events and forwards them into the OpenShift Container Platform Elasticsearch instance. Elasticsearch indexes the events to the **infra** index.

You must manually deploy the Event Router.

For information, see [Collecting and storing Kubernetes events](#).

2.2.13. About log forwarding

By default, the logging subsystem for Red Hat OpenShift sends logs to the default internal Elasticsearch log store, defined in the **ClusterLogging** custom resource (CR). If you want to forward logs to other log aggregators, you can use the log forwarding features to send logs to specific endpoints within or outside your cluster.

For information, see [Forwarding logs to third-party systems](#).

CHAPTER 3. INSTALLING THE LOGGING SUBSYSTEM FOR RED HAT OPENSIFT

You can install the logging subsystem for Red Hat OpenShift by deploying the OpenShift Elasticsearch and Red Hat OpenShift Logging Operators. The OpenShift Elasticsearch Operator creates and manages the Elasticsearch cluster used by OpenShift Logging. The logging subsystem Operator creates and manages the components of the logging stack.

The process for deploying the logging subsystem to OpenShift Container Platform involves:

- Reviewing the [Logging subsystem storage considerations](#).
- Installing the OpenShift Elasticsearch Operator and Red Hat OpenShift Logging Operator using the OpenShift Container Platform [web console](#) or [CLI](#).

3.1. INSTALLING THE LOGGING SUBSYSTEM FOR RED HAT OPENSIFT USING THE WEB CONSOLE

You can use the OpenShift Container Platform web console to install the OpenShift Elasticsearch and Red Hat OpenShift Logging Operators.



NOTE

If you do not want to use the default Elasticsearch log store, you can remove the internal Elasticsearch **logStore** and Kibana **visualization** components from the **ClusterLogging** custom resource (CR). Removing these components is optional but saves resources. For more information, see [Removing unused components if you do not use the default Elasticsearch log store](#).

Prerequisites

- Ensure that you have the necessary persistent storage for Elasticsearch. Note that each Elasticsearch node requires its own storage volume.



NOTE

If you use a local volume for persistent storage, do not use a raw block volume, which is described with **volumeMode: block** in the **LocalVolume** object. Elasticsearch cannot use raw block volumes.

Elasticsearch is a memory-intensive application. By default, OpenShift Container Platform installs three Elasticsearch nodes with memory requests and limits of 16 GB. This initial set of three OpenShift Container Platform nodes might not have enough memory to run Elasticsearch within your cluster. If you experience memory issues that are related to Elasticsearch, add more Elasticsearch nodes to your cluster rather than increasing the memory on existing nodes.

Procedure

To install the OpenShift Elasticsearch Operator and Red Hat OpenShift Logging Operator using the OpenShift Container Platform web console:

1. Install the OpenShift Elasticsearch Operator:
 - a. In the OpenShift Container Platform web console, click **Operators** → **OperatorHub**.

- b. Choose **OpenShift Elasticsearch Operator** from the list of available Operators, and click **Install**.
 - c. Ensure that the **All namespaces on the cluster** is selected under **Installation Mode**.
 - d. Ensure that **openshift-operators-redhat** is selected under **Installed Namespace**.
You must specify the **openshift-operators-redhat** namespace. The **openshift-operators** namespace might contain Community Operators, which are untrusted and could publish a metric with the same name as an OpenShift Container Platform metric, which would cause conflicts.
 - e. Select **Enable operator recommended cluster monitoring on this namespace**
This option sets the **openshift.io/cluster-monitoring: "true"** label in the Namespace object. You must select this option to ensure that cluster monitoring scrapes the **openshift-operators-redhat** namespace.
 - f. Select **stable-5.x** as the **Update Channel**.
 - g. Select an **Approval Strategy**.
 - The **Automatic** strategy allows Operator Lifecycle Manager (OLM) to automatically update the Operator when a new version is available.
 - The **Manual** strategy requires a user with appropriate credentials to approve the Operator update.
 - h. Click **Install**.
 - i. Verify that the OpenShift Elasticsearch Operator installed by switching to the **Operators → Installed Operators** page.
 - j. Ensure that **OpenShift Elasticsearch Operator** is listed in all projects with a **Status** of **Succeeded**.
2. Install the Red Hat OpenShift Logging Operator:
 - a. In the OpenShift Container Platform web console, click **Operators → OperatorHub**.
 - b. Choose **Red Hat OpenShift Logging** from the list of available Operators, and click **Install**.
 - c. Ensure that the **A specific namespace on the cluster** is selected under **Installation Mode**.
 - d. Ensure that **Operator recommended namespace** is **openshift-logging** under **Installed Namespace**.
 - e. Select **Enable operator recommended cluster monitoring on this namespace**
This option sets the **openshift.io/cluster-monitoring: "true"** label in the Namespace object. You must select this option to ensure that cluster monitoring scrapes the **openshift-logging** namespace.
 - f. Select **stable-5.x** as the **Update Channel**.
 - g. Select an **Approval Strategy**.
 - The **Automatic** strategy allows Operator Lifecycle Manager (OLM) to automatically update the Operator when a new version is available.

- The **Manual** strategy requires a user with appropriate credentials to approve the Operator update.
- h. Click **Install**.
 - i. Verify that the Red Hat OpenShift Logging Operator installed by switching to the **Operators → Installed Operators** page.
 - j. Ensure that **Red Hat OpenShift Logging** is listed in the **openshift-logging** project with a **Status** of **Succeeded**.
If the Operator does not appear as installed, to troubleshoot further:
 - Switch to the **Operators → Installed Operators** page and inspect the **Status** column for any errors or failures.
 - Switch to the **Workloads → Pods** page and check the logs in any pods in the **openshift-logging** project that are reporting issues.
3. Create an OpenShift Logging instance:
 - a. Switch to the **Administration → Custom Resource Definitions** page.
 - b. On the **Custom Resource Definitions** page, click **ClusterLogging**.
 - c. On the **Custom Resource Definition details** page, select **View Instances** from the **Actions** menu.
 - d. On the **ClusterLoggings** page, click **Create ClusterLogging**.
You might have to refresh the page to load the data.
 - e. In the YAML field, replace the code with the following:



NOTE

This default OpenShift Logging configuration should support a wide array of environments. Review the topics on tuning and configuring logging subsystem components for information on modifications you can make to your OpenShift Logging cluster.

```

apiVersion: "logging.openshift.io/v1"
kind: "ClusterLogging"
metadata:
  name: "instance" 1
  namespace: "openshift-logging"
spec:
  managementState: "Managed" 2
  logStore:
    type: "elasticsearch" 3
    retentionPolicy: 4
      application:
        maxAge: 1d
      infra:
        maxAge: 7d
      audit:
        maxAge: 7d

```

```

elasticsearch:
  nodeCount: 3 5
  storage:
    storageClassName: "<storage_class_name>" 6
    size: 200G
  resources: 7
    limits:
      memory: "16Gi"
    requests:
      memory: "16Gi"
  proxy: 8
    resources:
      limits:
        memory: 256Mi
      requests:
        memory: 256Mi
    redundancyPolicy: "SingleRedundancy"
  visualization:
    type: "kibana" 9
    kibana:
      replicas: 1
  collection:
    logs:
      type: "fluentd" 10
      fluentd: {}

```

- 1** The name must be **instance**.
- 2** The OpenShift Logging management state. In some cases, if you change the OpenShift Logging defaults, you must set this to **Unmanaged**. However, an unmanaged deployment does not receive updates until OpenShift Logging is placed back into a managed state.
- 3** Settings for configuring Elasticsearch. Using the CR, you can configure shard replication policy and persistent storage.
- 4** Specify the length of time that Elasticsearch should retain each log source. Enter an integer and a time designation: weeks(w), hours(h/H), minutes(m) and seconds(s). For example, **7d** for seven days. Logs older than the **maxAge** are deleted. You must specify a retention policy for each log source or the Elasticsearch indices will not be created for that source.
- 5** Specify the number of Elasticsearch nodes. See the note that follows this list.
- 6** Enter the name of an existing storage class for Elasticsearch storage. For best performance, specify a storage class that allocates block storage. If you do not specify a storage class, OpenShift Logging uses ephemeral storage.
- 7** Specify the CPU and memory requests for Elasticsearch as needed. If you leave these values blank, the OpenShift Elasticsearch Operator sets default values that should be sufficient for most deployments. The default values are **16Gi** for the memory request and **1** for the CPU request.
- 8** Specify the CPU and memory requests for the Elasticsearch proxy as needed. If you leave these values blank, the OpenShift Elasticsearch Operator sets default values that should be sufficient for most deployments. The default values are **256Mi** for the

memory request and **100m** for the CPU request.

- 9 Settings for configuring Kibana. Using the CR, you can scale Kibana for redundancy and configure the CPU and memory for your Kibana nodes. For more information, see **Configuring the log visualizer**.
- 10 Settings for configuring Fluentd. Using the CR, you can configure Fluentd CPU and memory limits. For more information, see **Configuring Fluentd**.



NOTE

The maximum number of Elasticsearch control plane nodes is three. If you specify a **nodeCount** greater than **3**, OpenShift Container Platform creates three Elasticsearch nodes that are Master-eligible nodes, with the master, client, and data roles. The additional Elasticsearch nodes are created as Data-only nodes, using client and data roles. Control plane nodes perform cluster-wide actions such as creating or deleting an index, shard allocation, and tracking nodes. Data nodes hold the shards and perform data-related operations such as CRUD, search, and aggregations. Data-related operations are I/O-, memory-, and CPU-intensive. It is important to monitor these resources and to add more Data nodes if the current nodes are overloaded.

For example, if **nodeCount=4**, the following nodes are created:

```
$ oc get deployment
```

Example output

```
cluster-logging-operator      1/1    1      1      18h
elasticsearch-cd-x6kdekli-1   0/1    1      0      6m54s
elasticsearch-cdm-x6kdekli-1  1/1    1      1      18h
elasticsearch-cdm-x6kdekli-2  0/1    1      0      6m49s
elasticsearch-cdm-x6kdekli-3  0/1    1      0      6m44s
```

The number of primary shards for the index templates is equal to the number of Elasticsearch data nodes.

- f. Click **Create**. This creates the logging subsystem components, the **Elasticsearch** custom resource and components, and the Kibana interface.
4. Verify the install:
 - a. Switch to the **Workloads → Pods** page.
 - b. Select the **openshift-logging** project.
You should see several pods for OpenShift Logging, Elasticsearch, Fluentd, and Kibana similar to the following list:
 - cluster-logging-operator-cb795f8dc-xkckc
 - elasticsearch-cdm-b3nqzchd-1-5c6797-67kfz
 - elasticsearch-cdm-b3nqzchd-2-6657f4-wtprv

- elasticsearch-cdm-b3nqzchd-3-588c65-clg7g
- fluentd-2c7dg
- fluentd-9z7kk
- fluentd-br7r2
- fluentd-fn2sb
- fluentd-pb2f8
- fluentd-zqqqx
- kibana-7fb4fd4cc9-bvt4p

Additional resources

- [Installing Operators from the OperatorHub](#)

3.2. POST-INSTALLATION TASKS

If you plan to use Kibana, you must [manually create your Kibana index patterns and visualizations](#) to explore and visualize data in Kibana.


If your cluster network provider enforces network isolation, [allow network traffic between the projects that contain the logging subsystem Operators](#).

3.3. INSTALLING THE LOGGING SUBSYSTEM FOR RED HAT OPENSIFT USING THE CLI

You can use the OpenShift Container Platform CLI to install the OpenShift Elasticsearch and Red Hat OpenShift Logging Operators.

Prerequisites

- Ensure that you have the necessary persistent storage for Elasticsearch. Note that each Elasticsearch node requires its own storage volume.



NOTE

If you use a local volume for persistent storage, do not use a raw block volume, which is described with **volumeMode: block** in the **LocalVolume** object. Elasticsearch cannot use raw block volumes.

Elasticsearch is a memory-intensive application. By default, OpenShift Container Platform installs three Elasticsearch nodes with memory requests and limits of 16 GB. This initial set of three OpenShift Container Platform nodes might not have enough memory to run Elasticsearch within your cluster. If you experience memory issues that are related to Elasticsearch, add more Elasticsearch nodes to your cluster rather than increasing the memory on existing nodes.

Procedure

To install the OpenShift Elasticsearch Operator and Red Hat OpenShift Logging Operator using the CLI:

1. Create a namespace for the OpenShift Elasticsearch Operator.
 - a. Create a namespace object YAML file (for example, **eo-namespace.yaml**) for the OpenShift Elasticsearch Operator:

```
apiVersion: v1
kind: Namespace
metadata:
  name: openshift-operators-redhat 1
  annotations:
    openshift.io/node-selector: ""
  labels:
    openshift.io/cluster-monitoring: "true" 2
```

- 1 You must specify the **openshift-operators-redhat** namespace. To prevent possible conflicts with metrics, you should configure the Prometheus Cluster Monitoring stack to scrape metrics from the **openshift-operators-redhat** namespace and not the **openshift-operators** namespace. The **openshift-operators** namespace might contain community Operators, which are untrusted and could publish a metric with the same name as an OpenShift Container Platform metric, which would cause conflicts.
- 2 String. You must specify this label as shown to ensure that cluster monitoring scrapes the **openshift-operators-redhat** namespace.

- b. Create the namespace:

```
$ oc create -f <file-name>.yaml
```

For example:

```
$ oc create -f eo-namespace.yaml
```

2. Create a namespace for the Red Hat OpenShift Logging Operator:
 - a. Create a namespace object YAML file (for example, **olo-namespace.yaml**) for the Red Hat OpenShift Logging Operator:

```
apiVersion: v1
kind: Namespace
metadata:
  name: openshift-logging
  annotations:
    openshift.io/node-selector: ""
  labels:
    openshift.io/cluster-monitoring: "true"
```

- b. Create the namespace:

```
$ oc create -f <file-name>.yaml
```

For example:

```
$ oc create -f olo-namespace.yaml
```

3. Install the OpenShift Elasticsearch Operator by creating the following objects:

- a. Create an Operator Group object YAML file (for example, **eo-og.yaml**) for the OpenShift Elasticsearch Operator:

```
apiVersion: operators.coreos.com/v1
kind: OperatorGroup
metadata:
  name: openshift-operators-redhat
  namespace: openshift-operators-redhat 1
spec: {}
```

- 1** You must specify the **openshift-operators-redhat** namespace.

- b. Create an Operator Group object:

```
$ oc create -f <file-name>.yaml
```

For example:

```
$ oc create -f eo-og.yaml
```

- c. Create a Subscription object YAML file (for example, **eo-sub.yaml**) to subscribe a namespace to the OpenShift Elasticsearch Operator.

Example Subscription

```
apiVersion: operators.coreos.com/v1alpha1
kind: Subscription
metadata:
  name: "elasticsearch-operator"
  namespace: "openshift-operators-redhat" 1
spec:
  channel: "stable-5.1" 2
  installPlanApproval: "Automatic" 3
  source: "redhat-operators" 4
  sourceNamespace: "openshift-marketplace"
  name: "elasticsearch-operator"
```

- 1** You must specify the **openshift-operators-redhat** namespace.
- 2** Specify **stable**, or **stable-5.<x>** as the channel. See the following note.
- 3** **Automatic** allows the Operator Lifecycle Manager (OLM) to automatically update the Operator when a new version is available. **Manual** requires a user with appropriate credentials to approve the Operator update.
- 4** Specify **redhat-operators**. If your OpenShift Container Platform cluster is installed on a restricted network, also known as a disconnected cluster, specify the name of the CatalogSource object created when you configured the Operator Lifecycle Manager (OLM).

**NOTE**

Specifying **stable** installs the current version of the latest stable release. Using **stable** with **installPlanApproval: "Automatic"**, will automatically upgrade your operators to the latest stable major and minor release.

Specifying **stable-5.<x>** installs the current minor version of a specific major release. Using **stable-5.<x>** with **installPlanApproval: "Automatic"**, will automatically upgrade your operators to the latest stable minor release within the major release you specify with **x**.

- d. Create the Subscription object:

```
$ oc create -f <file-name>.yaml
```

For example:

```
$ oc create -f eo-sub.yaml
```

The OpenShift Elasticsearch Operator is installed to the **openshift-operators-redhat** namespace and copied to each project in the cluster.

- e. Verify the Operator installation:

```
$ oc get csv --all-namespaces
```

Example output

```

NAMESPACE                                NAME                                DISPLAY
VERSION      REPLACES  PHASE
default                                elasticsearch-operator.5.1.0-202007012112.p0
OpenShift Elasticsearch Operator 5.1.0-202007012112.p0      Succeeded
kube-node-lease                                elasticsearch-operator.5.1.0-202007012112.p0
OpenShift Elasticsearch Operator 5.1.0-202007012112.p0      Succeeded
kube-public                                elasticsearch-operator.5.1.0-202007012112.p0
OpenShift Elasticsearch Operator 5.1.0-202007012112.p0      Succeeded
kube-system                                elasticsearch-operator.5.1.0-202007012112.p0
OpenShift Elasticsearch Operator 5.1.0-202007012112.p0      Succeeded
openshift-apiserver-operator              elasticsearch-operator.5.1.0-
202007012112.p0  OpenShift Elasticsearch Operator 5.1.0-202007012112.p0
Succeeded
openshift-apiserver                      elasticsearch-operator.5.1.0-202007012112.p0
OpenShift Elasticsearch Operator 5.1.0-202007012112.p0      Succeeded
openshift-authentication-operator        elasticsearch-operator.5.1.0-
202007012112.p0  OpenShift Elasticsearch Operator 5.1.0-202007012112.p0
Succeeded
openshift-authentication                  elasticsearch-operator.5.1.0-
202007012112.p0  OpenShift Elasticsearch Operator 5.1.0-202007012112.p0
Succeeded
...
```

There should be an OpenShift Elasticsearch Operator in each namespace. The version number might be different than shown.

4. Install the Red Hat OpenShift Logging Operator by creating the following objects:

- a. Create an Operator Group object YAML file (for example, **olo-og.yaml**) for the Red Hat OpenShift Logging Operator:

```
apiVersion: operators.coreos.com/v1
kind: OperatorGroup
metadata:
  name: cluster-logging
  namespace: openshift-logging ❶
spec:
  targetNamespaces:
    - openshift-logging ❷
```

❶ ❷ You must specify the **openshift-logging** namespace.

- b. Create the OperatorGroup object:

```
$ oc create -f <file-name>.yaml
```

For example:

```
$ oc create -f olo-og.yaml
```

- c. Create a Subscription object YAML file (for example, **olo-sub.yaml**) to subscribe a namespace to the Red Hat OpenShift Logging Operator.

```
apiVersion: operators.coreos.com/v1alpha1
kind: Subscription
metadata:
  name: cluster-logging
  namespace: openshift-logging ❶
spec:
  channel: "stable" ❷
  name: cluster-logging
  source: redhat-operators ❸
  sourceNamespace: openshift-marketplace
```

❶ You must specify the **openshift-logging** namespace.

❷ Specify **stable**, or **stable-5.<x>** as the channel.

❸ Specify **redhat-operators**. If your OpenShift Container Platform cluster is installed on a restricted network, also known as a disconnected cluster, specify the name of the CatalogSource object you created when you configured the Operator Lifecycle Manager (OLM).

```
$ oc create -f <file-name>.yaml
```

For example:

```
$ oc create -f olo-sub.yaml
```


The Red Hat OpenShift Logging Operator is installed to the **openshift-logging** namespace.

- d. Verify the Operator installation.

There should be a Red Hat OpenShift Logging Operator in the **openshift-logging** namespace. The Version number might be different than shown.

```
$ oc get csv -n openshift-logging
```

Example output

```

NAMESPACE                               NAME                               DISPLAY
VERSION      REPLACES  PHASE
...
openshift-logging          clusterlogging.5.1.0-202007012112.p0
OpenShift Logging        5.1.0-202007012112.p0          Succeeded
...

```

5. Create an OpenShift Logging instance:

- a. Create an instance object YAML file (for example, **olo-instance.yaml**) for the Red Hat OpenShift Logging Operator:



NOTE

This default OpenShift Logging configuration should support a wide array of environments. Review the topics on tuning and configuring logging subsystem components for information on modifications you can make to your OpenShift Logging cluster.

```

apiVersion: "logging.openshift.io/v1"
kind: "ClusterLogging"
metadata:
  name: "instance" 1
  namespace: "openshift-logging"
spec:
  managementState: "Managed" 2
  logStore:
    type: "elasticsearch" 3
    retentionPolicy: 4
      application:
        maxAge: 1d
      infra:
        maxAge: 7d
      audit:
        maxAge: 7d
    elasticsearch:
      nodeCount: 3 5
      storage:
        storageClassName: "<storage-class-name>" 6
        size: 200G
      resources: 7
        limits:
          memory: "16Gi"

```

```

    requests:
      memory: "16Gi"
    proxy: 8
    resources:
      limits:
        memory: 256Mi
      requests:
        memory: 256Mi
    redundancyPolicy: "SingleRedundancy"
  visualization:
    type: "kibana" 9
    kibana:
      replicas: 1
  collection:
    logs:
      type: "fluentd" 10
      fluentd: {}

```

- 1 The name must be **instance**.
- 2 The OpenShift Logging management state. In some cases, if you change the OpenShift Logging defaults, you must set this to **Unmanaged**. However, an unmanaged deployment does not receive updates until OpenShift Logging is placed back into a managed state. Placing a deployment back into a managed state might revert any modifications you made.
- 3 Settings for configuring Elasticsearch. Using the custom resource (CR), you can configure shard replication policy and persistent storage.
- 4 Specify the length of time that Elasticsearch should retain each log source. Enter an integer and a time designation: weeks(w), hours(h/H), minutes(m) and seconds(s). For example, **7d** for seven days. Logs older than the **maxAge** are deleted. You must specify a retention policy for each log source or the Elasticsearch indices will not be created for that source.
- 5 Specify the number of Elasticsearch nodes. See the note that follows this list.
- 6 Enter the name of an existing storage class for Elasticsearch storage. For best performance, specify a storage class that allocates block storage. If you do not specify a storage class, OpenShift Container Platform deploys OpenShift Logging with ephemeral storage only.
- 7 Specify the CPU and memory requests for Elasticsearch as needed. If you leave these values blank, the OpenShift Elasticsearch Operator sets default values that are sufficient for most deployments. The default values are **16Gi** for the memory request and **1** for the CPU request.
- 8 Specify the CPU and memory requests for the Elasticsearch proxy as needed. If you leave these values blank, the OpenShift Elasticsearch Operator sets default values that should be sufficient for most deployments. The default values are **256Mi** for the memory request and **100m** for the CPU request.
- 9 Settings for configuring Kibana. Using the CR, you can scale Kibana for redundancy and configure the CPU and memory for your Kibana pods. For more information, see **Configuring the log visualizer**.

- 10 Settings for configuring Fluentd. Using the CR, you can configure Fluentd CPU and memory limits. For more information, see **Configuring Fluentd**.



NOTE

The maximum number of Elasticsearch control plane nodes is three. If you specify a **nodeCount** greater than **3**, OpenShift Container Platform creates three Elasticsearch nodes that are Master-eligible nodes, with the master, client, and data roles. The additional Elasticsearch nodes are created as Data-only nodes, using client and data roles. Control plane nodes perform cluster-wide actions such as creating or deleting an index, shard allocation, and tracking nodes. Data nodes hold the shards and perform data-related operations such as CRUD, search, and aggregations. Data-related operations are I/O-, memory-, and CPU-intensive. It is important to monitor these resources and to add more Data nodes if the current nodes are overloaded.

For example, if **nodeCount=4**, the following nodes are created:

```
$ oc get deployment
```

Example output

```
cluster-logging-operator      1/1    1      1      18h
elasticsearch-cd-x6kdekli-1  1/1    1      0      6m54s
elasticsearch-cdm-x6kdekli-1  1/1    1      1      18h
elasticsearch-cdm-x6kdekli-2  1/1    1      0      6m49s
elasticsearch-cdm-x6kdekli-3  1/1    1      0      6m44s
```

The number of primary shards for the index templates is equal to the number of Elasticsearch data nodes.

- b. Create the instance:

```
$ oc create -f <file-name>.yaml
```

For example:

```
$ oc create -f olo-instance.yaml
```

This creates the logging subsystem components, the **Elasticsearch** custom resource and components, and the Kibana interface.

6. Verify the installation by listing the pods in the **openshift-logging** project. You should see several pods for components of the Logging subsystem, similar to the following list:

```
$ oc get pods -n openshift-logging
```

Example output

```
NAME                                READY STATUS RESTARTS AGE
cluster-logging-operator-66f77fccb-ppzbg  1/1 Running 0      7m
```

elasticsearch-cdm-ftuhduuw-1-ffc4b9566-q6bhp	2/2	Running	0	2m40s
elasticsearch-cdm-ftuhduuw-2-7b4994dbfc-rd2gc	2/2	Running	0	2m36s
elasticsearch-cdm-ftuhduuw-3-84b5ff7ff8-gqnm2	2/2	Running	0	2m4s
collector-587vb	1/1	Running	0	2m26s
collector-7mpb9	1/1	Running	0	2m30s
collector-flm6j	1/1	Running	0	2m33s
collector-gn4rn	1/1	Running	0	2m26s
collector-nlgb6	1/1	Running	0	2m30s
collector-snpkt	1/1	Running	0	2m28s
kibana-d6d5668c5-rppqm	2/2	Running	0	2m39s

3.4. POST-INSTALLATION TASKS

If you plan to use Kibana, you must [manually create your Kibana index patterns and visualizations](#) to explore and visualize data in Kibana.

If your cluster network provider enforces network isolation, [allow network traffic between the projects that contain the logging subsystem Operators](#).

3.4.1. Defining Kibana index patterns

An index pattern defines the Elasticsearch indices that you want to visualize. To explore and visualize data in Kibana, you must create an index pattern.

Prerequisites

- A user must have the **cluster-admin** role, the **cluster-reader** role, or both roles to view the **infra** and **audit** indices in Kibana. The default **kubeadmin** user has proper permissions to view these indices.

If you can view the pods and logs in the **default**, **kube-** and **openshift-** projects, you should be able to access these indices. You can use the following command to check if the current user has appropriate permissions:

```
$ oc auth can-i get pods/log -n <project>
```

Example output

```
yes
```




NOTE

The audit logs are not stored in the internal OpenShift Container Platform Elasticsearch instance by default. To view the audit logs in Kibana, you must use the Log Forwarding API to configure a pipeline that uses the **default** output for audit logs.

- Elasticsearch documents must be indexed before you can create index patterns. This is done automatically, but it might take a few minutes in a new or updated cluster.

Procedure

To define index patterns and create visualizations in Kibana:

1. In the OpenShift Container Platform console, click the Application Launcher  and select **Logging**.
2. Create your Kibana index patterns by clicking **Management** → **Index Patterns** → **Create index pattern**:
 - Each user must manually create index patterns when logging into Kibana the first time to see logs for their projects. Users must create an index pattern named **app** and use the **@timestamp** time field to view their container logs.
 - Each admin user must create index patterns when logged into Kibana the first time for the **app**, **infra**, and **audit** indices using the **@timestamp** time field.
3. Create Kibana Visualizations from the new index patterns.

3.4.2. Allowing traffic between projects when network isolation is enabled

Your cluster network provider might enforce network isolation. If so, you must allow network traffic between the projects that contain the operators deployed by OpenShift Logging.

Network isolation blocks network traffic between pods or services that are in different projects. The logging subsystem installs the *OpenShift Elasticsearch Operator* in the **openshift-operators-redhat** project and the *Red Hat OpenShift Logging Operator* in the **openshift-logging** project. Therefore, you must allow traffic between these two projects.

OpenShift Container Platform offers two supported choices for the default Container Network Interface (CNI) network provider, OpenShift SDN and OVN-Kubernetes. These two providers implement various network isolation policies.

OpenShift SDN has three modes:

network policy

This is the default mode. If no policy is defined, it allows all traffic. However, if a user defines a policy, they typically start by denying all traffic and then adding exceptions. This process might break applications that are running in different projects. Therefore, explicitly configure the policy to allow traffic to egress from one logging-related project to the other.

multitenant

This mode enforces network isolation. You must join the two logging-related projects to allow traffic between them.

subnet

This mode allows all traffic. It does not enforce network isolation. No action is needed.

OVN-Kubernetes always uses a **network policy**. Therefore, as with OpenShift SDN, you must configure the policy to allow traffic to egress from one logging-related project to the other.

Procedure

- If you are using OpenShift SDN in **multitenant** mode, join the two projects. For example:

```
$ oc adm pod-network join-projects --to=openshift-operators-redhat openshift-logging
```

- Otherwise, for OpenShift SDN in **network policy** mode and OVN-Kubernetes, perform the following actions:

- a. Set a label on the **openshift-operators-redhat** namespace. For example:

```
$ oc label namespace openshift-operators-redhat project=openshift-operators-redhat
```

- b. Create a network policy object in the **openshift-logging** namespace that allows ingress from the **openshift-operators-redhat**, **openshift-monitoring** and **openshift-ingress** projects to the openshift-logging project. For example:

```
apiVersion: networking.k8s.io/v1
kind: NetworkPolicy
metadata:
  name: allow-from-openshift-monitoring-ingress-operators-redhat
spec:
  ingress:
  - from:
    - podSelector: {}
  - from:
    - namespaceSelector:
        matchLabels:
          project: "openshift-operators-redhat"
  - from:
    - namespaceSelector:
        matchLabels:
          name: "openshift-monitoring"
  - from:
    - namespaceSelector:
        matchLabels:
          network.openshift.io/policy-group: ingress
  podSelector: {}
  policyTypes:
  - Ingress
```

Additional resources

- [About network policy](#)
- [About the OpenShift SDN default CNI network provider](#)
- [About the OVN-Kubernetes default Container Network Interface \(CNI\) network provider](#)

CHAPTER 4. CONFIGURING YOUR LOGGING DEPLOYMENT

4.1. ABOUT THE CLUSTER LOGGING CUSTOM RESOURCE

To configure logging subsystem for Red Hat OpenShift you customize the **ClusterLogging** custom resource (CR).

4.1.1. About the ClusterLogging custom resource

To make changes to your logging subsystem environment, create and modify the **ClusterLogging** custom resource (CR).

Instructions for creating or modifying a CR are provided in this documentation as appropriate.

The following example shows a typical custom resource for the logging subsystem.

Sample ClusterLogging custom resource (CR)

```

apiVersion: "logging.openshift.io/v1"
kind: "ClusterLogging"
metadata:
  name: "instance" ❶
  namespace: "openshift-logging" ❷
spec:
  managementState: "Managed" ❸
  logStore:
    type: "elasticsearch" ❹
    retentionPolicy:
      application:
        maxAge: 1d
      infra:
        maxAge: 7d
      audit:
        maxAge: 7d
    elasticsearch:
      nodeCount: 3
      resources:
        limits:
          memory: 16Gi
        requests:
          cpu: 500m
          memory: 16Gi
      storage:
        storageClassName: "gp2"
        size: "200G"
      redundancyPolicy: "SingleRedundancy"
  visualization: ❺
    type: "kibana"
    kibana:
      resources:
        limits:
          memory: 736Mi
        requests:
          cpu: 100m

```

```

    memory: 736Mi
    replicas: 1
collection: 6
logs:
  type: "fluentd"
  fluentd:
    resources:
      limits:
        memory: 736Mi
      requests:
        cpu: 100m
        memory: 736Mi

```

- 1 The CR name must be **instance**.
- 2 The CR must be installed to the **openshift-logging** namespace.
- 3 The Red Hat OpenShift Logging Operator management state. When set to **unmanaged** the operator is in an unsupported state and will not get updates.
- 4 Settings for the log store, including retention policy, the number of nodes, the resource requests and limits, and the storage class.
- 5 Settings for the visualizer, including the resource requests and limits, and the number of pod replicas.
- 6 Settings for the log collector, including the resource requests and limits.

4.2. CONFIGURING THE LOGGING COLLECTOR

Logging subsystem for Red Hat OpenShift collects operations and application logs from your cluster and enriches the data with Kubernetes pod and project metadata.

You can configure the CPU and memory limits for the log collector and [move the log collector pods to specific nodes](#). All supported modifications to the log collector can be performed through the **spec.collection.log.fluentd** stanza in the **ClusterLogging** custom resource (CR).

4.2.1. About unsupported configurations

The supported way of configuring the logging subsystem for Red Hat OpenShift is by configuring it using the options described in this documentation. Do not use other configurations, as they are unsupported. Configuration paradigms might change across OpenShift Container Platform releases, and such cases can only be handled gracefully if all configuration possibilities are controlled. If you use configurations other than those described in this documentation, your changes will disappear because the OpenShift Elasticsearch Operator and Red Hat OpenShift Logging Operator reconcile any differences. The Operators reverse everything to the defined state by default and by design.



NOTE

If you *must* perform configurations not described in the OpenShift Container Platform documentation, you *must* set your Red Hat OpenShift Logging Operator or OpenShift Elasticsearch Operator to **Unmanaged**. An unmanaged OpenShift Logging environment is *not supported* and does not receive updates until you return OpenShift Logging to **Managed**.

4.2.2. Viewing logging collector pods

You can view the Fluentd logging collector pods and the corresponding nodes that they are running on. The Fluentd logging collector pods run only in the **openshift-logging** project.

Procedure

- Run the following command in the **openshift-logging** project to view the Fluentd logging collector pods and their details:

```
$ oc get pods --selector component=collector -o wide -n openshift-logging
```

Example output

```

NAME           READY STATUS   RESTARTS  AGE   IP           NODE           NOMINATED
NODE READINESS GATES
fluentd-8d69v 1/1   Running 0       134m  10.130.2.30  master1.example.com <none>
<none>
fluentd-bd225 1/1   Running 0       134m  10.131.1.11  master2.example.com <none>
<none>
fluentd-cvrzs 1/1   Running 0       134m  10.130.0.21  master3.example.com <none>
<none>
fluentd-gpqg2 1/1   Running 0       134m  10.128.2.27  worker1.example.com <none>
<none>
fluentd-l9j7j 1/1   Running 0       134m  10.129.2.31  worker2.example.com <none>
<none>

```

4.2.3. Configure log collector CPU and memory limits

The log collector allows for adjustments to both the CPU and memory limits.

Procedure

1. Edit the **ClusterLogging** custom resource (CR) in the **openshift-logging** project:

```
$ oc -n openshift-logging edit ClusterLogging instance
```

```

apiVersion: "logging.openshift.io/v1"
kind: "ClusterLogging"
metadata:
  name: "instance"
  namespace: openshift-logging
...
spec:
  collection:
    logs:
      fluentd:
        resources:
          limits: 1
            memory: 736Mi

```

```
requests:  
cpu: 100m  
memory: 736Mi
```



Specify the CPU and memory limits and requests as needed. The values shown are the default values.

4.2.4. Advanced configuration for the log forwarder

The logging subsystem for Red Hat OpenShift includes multiple Fluentd parameters that you can use for tuning the performance of the Fluentd log forwarder. With these parameters, you can change the following Fluentd behaviors:

- Chunk and chunk buffer sizes
- Chunk flushing behavior
- Chunk forwarding retry behavior

Fluentd collects log data in a single blob called a *chunk*. When Fluentd creates a chunk, the chunk is considered to be in the *stage*, where the chunk gets filled with data. When the chunk is full, Fluentd moves the chunk to the *queue*, where chunks are held before being flushed, or written out to their destination. Fluentd can fail to flush a chunk for a number of reasons, such as network issues or capacity issues at the destination. If a chunk cannot be flushed, Fluentd retries flushing as configured.

By default in OpenShift Container Platform, Fluentd uses the *exponential backoff* method to retry flushing, where Fluentd doubles the time it waits between attempts to retry flushing again, which helps reduce connection requests to the destination. You can disable exponential backoff and use the *periodic* retry method instead, which retries flushing the chunks at a specified interval.

These parameters can help you determine the trade-offs between latency and throughput.

- To optimize Fluentd for throughput, you could use these parameters to reduce network packet count by configuring larger buffers and queues, delaying flushes, and setting longer times between retries. Be aware that larger buffers require more space on the node file system.
- To optimize for low latency, you could use the parameters to send data as soon as possible, avoid the build-up of batches, have shorter queues and buffers, and use more frequent flush and retries.

You can configure the chunking and flushing behavior using the following parameters in the **ClusterLogging** custom resource (CR). The parameters are then automatically added to the Fluentd config map for use by Fluentd.

**NOTE**

These parameters are:

- Not relevant to most users. The default settings should give good general performance.
- Only for advanced users with detailed knowledge of Fluentd configuration and performance.
- Only for performance tuning. They have no effect on functional aspects of logging.

Table 4.1. Advanced Fluentd Configuration Parameters

Parameter	Description	Default
chunkLimitSize	The maximum size of each chunk. Fluentd stops writing data to a chunk when it reaches this size. Then, Fluentd sends the chunk to the queue and opens a new chunk.	8m
totalLimitSize	The maximum size of the buffer, which is the total size of the stage and the queue. If the buffer size exceeds this value, Fluentd stops adding data to chunks and fails with an error. All data not in chunks is lost.	8G
flushInterval	The interval between chunk flushes. You can use s (seconds), m (minutes), h (hours), or d (days).	1s
flushMode	The method to perform flushes: <ul style="list-style-type: none"> • lazy: Flush chunks based on the timekey parameter. You cannot modify the timekey parameter. • interval: Flush chunks based on the flushInterval parameter. • immediate: Flush chunks immediately after data is added to a chunk. 	interval

Parameter	Description	Default
flushThreadCount	The number of threads that perform chunk flushing. Increasing the number of threads improves the flush throughput, which hides network latency.	2
overflowAction	The chunking behavior when the queue is full: <ul style="list-style-type: none"> ● throw_exception: Raise an exception to show in the log. ● block: Stop data chunking until the full buffer issue is resolved. ● drop_oldest_chunk: Drop the oldest chunk to accept new incoming chunks. Older chunks have less value than newer chunks. 	block
retryMaxInterval	The maximum time in seconds for the exponential_backoff retry method.	300s
retryType	The retry method when flushing fails: <ul style="list-style-type: none"> ● exponential_backoff: Increase the time between flush retries. Fluentd doubles the time it waits until the next retry until the retry_max_interval parameter is reached. ● periodic: Retries flushes periodically, based on the retryWait parameter. 	exponential_backoff
retryTimeout	The maximum time interval to attempt retries before the record is discarded.	60m
retryWait	The time in seconds before the next chunk flush.	1s

For more information on the Fluentd chunk lifecycle, see [Buffer Plugins](#) in the Fluentd documentation.

Procedure

1. Edit the **ClusterLogging** custom resource (CR) in the **openshift-logging** project:

```
$ oc edit ClusterLogging instance
```

2. Add or modify any of the following parameters:

```
apiVersion: logging.openshift.io/v1
kind: ClusterLogging
metadata:
  name: instance
  namespace: openshift-logging
spec:
  forwarder:
    fluentd:
      buffer:
        chunkLimitSize: 8m 1
        flushInterval: 5s 2
        flushMode: interval 3
        flushThreadCount: 3 4
        overflowAction: throw_exception 5
        retryMaxInterval: "300s" 6
        retryType: periodic 7
        retryWait: 1s 8
        totalLimitSize: 32m 9
  ...
```

- 1 Specify the maximum size of each chunk before it is queued for flushing.
- 2 Specify the interval between chunk flushes.
- 3 Specify the method to perform chunk flushes: **lazy**, **interval**, or **immediate**.
- 4 Specify the number of threads to use for chunk flushes.
- 5 Specify the chunking behavior when the queue is full: **throw_exception**, **block**, or **drop_oldest_chunk**.
- 6 Specify the maximum interval in seconds for the **exponential_backoff** chunk flushing method.
- 7 Specify the retry type when chunk flushing fails: **exponential_backoff** or **periodic**.
- 8 Specify the time in seconds before the next chunk flush.
- 9 Specify the maximum size of the chunk buffer.

3. Verify that the Fluentd pods are redeployed:

```
$ oc get pods -l component=collector -n openshift-logging
```

4. Check that the new values are in the **fluentd** config map:

```
$ oc extract configmap/fluentd --confirm
```

Example fluentd.conf

```
<buffer>
  @type file
  path '/var/lib/fluentd/default'
  flush_mode interval
  flush_interval 5s
  flush_thread_count 3
  retry_type periodic
  retry_wait 1s
  retry_max_interval 300s
  retry_timeout 60m
  queued_chunks_limit_size "#{ENV['BUFFER_QUEUE_LIMIT'] || '32'}"
  total_limit_size 32m
  chunk_limit_size 8m
  overflow_action throw_exception
</buffer>
```

4.2.5. Removing unused components if you do not use the default Elasticsearch log store

As an administrator, in the rare case that you forward logs to a third-party log store and do not use the default Elasticsearch log store, you can remove several unused components from your logging cluster.

In other words, if you do not use the default Elasticsearch log store, you can remove the internal Elasticsearch **logStore** and Kibana **visualization** components from the **ClusterLogging** custom resource (CR). Removing these components is optional but saves resources.

Prerequisites

- Verify that your log forwarder does not send log data to the default internal Elasticsearch cluster. Inspect the **ClusterLogForwarder** CR YAML file that you used to configure log forwarding. Verify that it *does not* have an **outputRefs** element that specifies **default**. For example:

```
outputRefs:
- default
```



WARNING

Suppose the **ClusterLogForwarder** CR forwards log data to the internal Elasticsearch cluster, and you remove the **logStore** component from the **ClusterLogging** CR. In that case, the internal Elasticsearch cluster will not be present to store the log data. This absence can cause data loss.

Procedure

1. Edit the **ClusterLogging** custom resource (CR) in the **openshift-logging** project:

```
$ oc edit ClusterLogging instance
```

2. If they are present, remove the **logStore** and **visualization** stanzas from the **ClusterLogging** CR.
3. Preserve the **collection** stanza of the **ClusterLogging** CR. The result should look similar to the following example:

```
apiVersion: "logging.openshift.io/v1"
kind: "ClusterLogging"
metadata:
  name: "instance"
  namespace: "openshift-logging"
spec:
  managementState: "Managed"
  collection:
    logs:
      type: "fluentd"
      fluentd: {}
```

4. Verify that the collector pods are redeployed:

```
$ oc get pods -l component=collector -n openshift-logging
```

Additional resources

- [Forwarding logs to third-party systems](#)

4.3. CONFIGURING THE LOG STORE

Logging subsystem for Red Hat OpenShift uses Elasticsearch 6 (ES) to store and organize the log data.

You can make modifications to your log store, including:

- storage for your Elasticsearch cluster
- shard replication across data nodes in the cluster, from full replication to no replication
- external access to Elasticsearch data


Elasticsearch is a memory-intensive application. Each Elasticsearch node needs at least 16G of memory for both memory requests and limits, unless you specify otherwise in the **ClusterLogging** custom resource. The initial set of OpenShift Container Platform nodes might not be large enough to support the Elasticsearch cluster. You must add additional nodes to the OpenShift Container Platform cluster to run with the recommended or higher memory, up to a maximum of 64G for each Elasticsearch node.

Each Elasticsearch node can operate with a lower memory setting, though this is not recommended for production environments.

4.3.1. Forwarding audit logs to the log store

By default, OpenShift Logging does not store audit logs in the internal OpenShift Container Platform Elasticsearch log store. You can send audit logs to this log store so, for example, you can view them in Kibana.

To send the audit logs to the default internal Elasticsearch log store, for example to view the audit logs in Kibana, you must use the Log Forwarding API.



IMPORTANT

The internal OpenShift Container Platform Elasticsearch log store does not provide secure storage for audit logs. Verify that the system to which you forward audit logs complies with your organizational and governmental regulations and is properly secured. The logging subsystem for Red Hat OpenShift does not comply with those regulations.

Procedure

To use the Log Forward API to forward audit logs to the internal Elasticsearch instance:

1. Create or edit a YAML file that defines the **ClusterLogForwarder** CR object:
 - Create a CR to send all log types to the internal Elasticsearch instance. You can use the following example without making any changes:

```
apiVersion: logging.openshift.io/v1
kind: ClusterLogForwarder
metadata:
  name: instance
  namespace: openshift-logging
spec:
  pipelines: 1
  - name: all-to-default
    inputRefs:
      - infrastructure
      - application
      - audit
    outputRefs:
      - default
```

- 1** A pipeline defines the type of logs to forward using the specified output. The default output forwards logs to the internal Elasticsearch instance.



NOTE

You must specify all three types of logs in the pipeline: application, infrastructure, and audit. If you do not specify a log type, those logs are not stored and will be lost.

- If you have an existing **ClusterLogForwarder** CR, add a pipeline to the default output for the audit logs. You do not need to define the default output. For example:

```
apiVersion: "logging.openshift.io/v1"
kind: ClusterLogForwarder
metadata:
  name: instance
```



```

namespace: openshift-logging
spec:
  outputs:
    - name: elasticsearch-insecure
      type: "elasticsearch"
      url: http://elasticsearch-insecure.messaging.svc.cluster.local
      insecure: true
    - name: elasticsearch-secure
      type: "elasticsearch"
      url: https://elasticsearch-secure.messaging.svc.cluster.local
      secret:
        name: es-audit
    - name: secureforward-offcluster
      type: "fluentdForward"
      url: https://secureforward.offcluster.com:24224
      secret:
        name: secureforward
  pipelines:
    - name: container-logs
      inputRefs:
        - application
      outputRefs:
        - secureforward-offcluster
    - name: infra-logs
      inputRefs:
        - infrastructure
      outputRefs:
        - elasticsearch-insecure
    - name: audit-logs
      inputRefs:
        - audit
      outputRefs:
        - elasticsearch-secure
        - default 1

```

- 1** This pipeline sends the audit logs to the internal Elasticsearch instance in addition to an external instance.

Additional resources

- For more information on the Log Forwarding API, see [Forwarding logs using the Log Forwarding API](#).

4.3.2. Configuring log retention time

You can configure a *retention policy* that specifies how long the default Elasticsearch log store keeps indices for each of the three log sources: infrastructure logs, application logs, and audit logs.

To configure the retention policy, you set a **maxAge** parameter for each log source in the **ClusterLogging** custom resource (CR). The CR applies these values to the Elasticsearch rollover schedule, which determines when Elasticsearch deletes the rolled-over indices.

Elasticsearch rolls over an index, moving the current index and creating a new index, when an index matches any of the following conditions:

- The index is older than the **rollover.maxAge** value in the **Elasticsearch** CR.
- The index size is greater than 40 GB × the number of primary shards.
- The index doc count is greater than 40960 KB × the number of primary shards.

Elasticsearch deletes the rolled-over indices based on the retention policy you configure. If you do not create a retention policy for any log sources, logs are deleted after seven days by default.

Prerequisites

- The logging subsystem for Red Hat OpenShift and the OpenShift Elasticsearch Operator must be installed.

Procedure

To configure the log retention time:

1. Edit the **ClusterLogging** CR to add or modify the **retentionPolicy** parameter:

```
apiVersion: "logging.openshift.io/v1"
kind: "ClusterLogging"
...
spec:
  managementState: "Managed"
  logStore:
    type: "elasticsearch"
    retentionPolicy: 1
    application:
      maxAge: 1d
    infra:
      maxAge: 7d
    audit:
      maxAge: 7d
  elasticsearch:
    nodeCount: 3
...
```

- 1 Specify the time that Elasticsearch should retain each log source. Enter an integer and a time designation: weeks(w), hours(h/H), minutes(m) and seconds(s). For example, **1d** for one day. Logs older than the **maxAge** are deleted. By default, logs are retained for seven days.

2. You can verify the settings in the **Elasticsearch** custom resource (CR). For example, the Red Hat OpenShift Logging Operator updated the following **Elasticsearch** CR to configure a retention policy that includes settings to roll over active indices for the infrastructure logs every eight hours and the rolled-over indices are deleted seven days after rollover. OpenShift Container Platform checks every 15 minutes to determine if the indices need to be rolled over.

```
apiVersion: "logging.openshift.io/v1"
kind: "Elasticsearch"
metadata:
  name: "elasticsearch"
spec:
```

```

...
indexManagement:
  policies: 1
    - name: infra-policy
      phases:
        delete:
          minAge: 7d 2
        hot:
          actions:
            rollover:
              maxAge: 8h 3
      pollInterval: 15m 4
...

```

- 1 For each log source, the retention policy indicates when to delete and roll over logs for that source.
- 2 When OpenShift Container Platform deletes the rolled-over indices. This setting is the **maxAge** you set in the **ClusterLogging** CR.
- 3 The index age for OpenShift Container Platform to consider when rolling over the indices. This value is determined from the **maxAge** you set in the **ClusterLogging** CR.
- 4 When OpenShift Container Platform checks if the indices should be rolled over. This setting is the default and cannot be changed.



NOTE

Modifying the **Elasticsearch** CR is not supported. All changes to the retention policies must be made in the **ClusterLogging** CR.

The OpenShift Elasticsearch Operator deploys a cron job to roll over indices for each mapping using the defined policy, scheduled using the **pollInterval**.


```
$ oc get cronjob
```

Example output

NAME	SCHEDULE	SUSPEND	ACTIVE	LAST SCHEDULE	AGE
elasticsearch-im-app	*/15 * * * *	False	0	<none>	4s
elasticsearch-im-audit	*/15 * * * *	False	0	<none>	4s
elasticsearch-im-infra	*/15 * * * *	False	0	<none>	4s

4.3.3. Configuring CPU and memory requests for the log store

Each component specification allows for adjustments to both the CPU and memory requests. You should not have to manually adjust these values as the OpenShift Elasticsearch Operator sets values sufficient for your environment.



NOTE

In large-scale clusters, the default memory limit for the Elasticsearch proxy container might not be sufficient, causing the proxy container to be OOMKilled. If you experience this issue, increase the memory requests and limits for the Elasticsearch proxy.

Each Elasticsearch node can operate with a lower memory setting though this is **not** recommended for production deployments. For production use, you should have no less than the default 16Gi allocated to each pod. Preferably you should allocate as much as possible, up to 64Gi per pod.

Prerequisites

- The Red Hat OpenShift Logging and Elasticsearch Operators must be installed.

Procedure

1. Edit the **ClusterLogging** custom resource (CR) in the **openshift-logging** project:

```
$ oc edit ClusterLogging instance

apiVersion: "logging.openshift.io/v1"
kind: "ClusterLogging"
metadata:
  name: "instance"
....
spec:
  logStore:
    type: "elasticsearch"
    elasticsearch: 1
    resources:
      limits: 2
        memory: "32Gi"
      requests: 3
        cpu: "1"
        memory: "16Gi"
    proxy: 4
    resources:
      limits:
        memory: 100Mi
      requests:
        memory: 100Mi
```

- 1** Specify the CPU and memory requests for Elasticsearch as needed. If you leave these values blank, the OpenShift Elasticsearch Operator sets default values that should be sufficient for most deployments. The default values are **16Gi** for the memory request and **1** for the CPU request.
- 2** The maximum amount of resources a pod can use.
- 3** The minimum resources required to schedule a pod.
- 4** Specify the CPU and memory requests for the Elasticsearch proxy as needed. If you leave these values blank, the OpenShift Elasticsearch Operator sets default values that are sufficient for most deployments. The default values are **256Mi** for the memory request and

100m for the CPU request.

When adjusting the amount of Elasticsearch memory, the same value should be used for both **requests** and **limits**.

For example:

```
resources:
  limits: 1
    memory: "32Gi"
  requests: 2
    cpu: "8"
    memory: "32Gi"
```

- 1** The maximum amount of the resource.
- 2** The minimum amount required.

Kubernetes generally adheres the node configuration and does not allow Elasticsearch to use the specified limits. Setting the same value for the **requests** and **limits** ensures that Elasticsearch can use the memory you want, assuming the node has the memory available.

4.3.4. Configuring replication policy for the log store

You can define how Elasticsearch shards are replicated across data nodes in the cluster.

Prerequisites

- The Red Hat OpenShift Logging and Elasticsearch Operators must be installed.

Procedure

1. Edit the **ClusterLogging** custom resource (CR) in the **openshift-logging** project:

```
$ oc edit clusterlogging instance

apiVersion: "logging.openshift.io/v1"
kind: "ClusterLogging"
metadata:
  name: "instance"
....
spec:
  logStore:
    type: "elasticsearch"
    elasticsearch:
      redundancyPolicy: "SingleRedundancy" 1
```

- 1** Specify a redundancy policy for the shards. The change is applied upon saving the changes.

- **FullRedundancy.** Elasticsearch fully replicates the primary shards for each index to every data node. This provides the highest safety, but at the cost of the highest amount of disk required and the poorest performance.
- **MultipleRedundancy.** Elasticsearch fully replicates the primary shards for each index to half of the data nodes. This provides a good tradeoff between safety and performance.
- **SingleRedundancy.** Elasticsearch makes one copy of the primary shards for each index. Logs are always available and recoverable as long as at least two data nodes exist. Better performance than MultipleRedundancy, when using 5 or more nodes. You cannot apply this policy on deployments of single Elasticsearch node.
- **ZeroRedundancy.** Elasticsearch does not make copies of the primary shards. Logs might be unavailable or lost in the event a node is down or fails. Use this mode when you are more concerned with performance than safety, or have implemented your own disk/PVC backup/restore strategy.




NOTE

The number of primary shards for the index templates is equal to the number of Elasticsearch data nodes.

4.3.5. Scaling down Elasticsearch pods

Reducing the number of Elasticsearch pods in your cluster can result in data loss or Elasticsearch performance degradation.

If you scale down, you should scale down by one pod at a time and allow the cluster to re-balance the shards and replicas. After the Elasticsearch health status returns to **green**, you can scale down by another pod.



NOTE

If your Elasticsearch cluster is set to **ZeroRedundancy**, you should not scale down your Elasticsearch pods.

4.3.6. Configuring persistent storage for the log store

Elasticsearch requires persistent storage. The faster the storage, the faster the Elasticsearch performance.



WARNING

Using NFS storage as a volume or a persistent volume (or via NAS such as Gluster) is not supported for Elasticsearch storage, as Lucene relies on file system behavior that NFS does not supply. Data corruption and other problems can occur.

Prerequisites

- The Red Hat OpenShift Logging and Elasticsearch Operators must be installed.

Procedure

1. Edit the **ClusterLogging** CR to specify that each data node in the cluster is bound to a Persistent Volume Claim.

```
apiVersion: "logging.openshift.io/v1"
kind: "ClusterLogging"
metadata:
  name: "instance"
# ...
spec:
  logStore:
    type: "elasticsearch"
    elasticsearch:
      nodeCount: 3
      storage:
        storageClassName: "gp2"
        size: "200G"
```

This example specifies each data node in the cluster is bound to a Persistent Volume Claim that requests "200G" of AWS General Purpose SSD (gp2) storage.



NOTE

If you use a local volume for persistent storage, do not use a raw block volume, which is described with **volumeMode: block** in the **LocalVolume** object. Elasticsearch cannot use raw block volumes.

4.3.7. Configuring the log store for emptyDir storage

You can use emptyDir with your log store, which creates an ephemeral deployment in which all of a pod's data is lost upon restart.



NOTE

When using emptyDir, if log storage is restarted or redeployed, you will lose data.

Prerequisites

- The Red Hat OpenShift Logging and Elasticsearch Operators must be installed.

Procedure

1. Edit the **ClusterLogging** CR to specify emptyDir:

```
spec:
  logStore:
    type: "elasticsearch"
    elasticsearch:
      nodeCount: 3
      storage: {}
```

4.3.8. Performing an Elasticsearch rolling cluster restart

Perform a rolling restart when you change the **elasticsearch** config map or any of the **elasticsearch-*** deployment configurations.

Also, a rolling restart is recommended if the nodes on which an Elasticsearch pod runs requires a reboot.

Prerequisites

- The Red Hat OpenShift Logging and Elasticsearch Operators must be installed.

Procedure

To perform a rolling cluster restart:

1. Change to the **openshift-logging** project:

```
$ oc project openshift-logging
```

2. Get the names of the Elasticsearch pods:

```
$ oc get pods -l component=elasticsearch-
```

3. Scale down the collector pods so they stop sending new logs to Elasticsearch:

```
$ oc -n openshift-logging patch daemonset/collector -p '{"spec":{"template":{"spec":{"nodeSelector":{"logging-infra-collector": "false"}}}}}'
```

4. Perform a shard synced flush using the OpenShift Container Platform [es_util](#) tool to ensure there are no pending operations waiting to be written to disk prior to shutting down:

```
$ oc exec <any_es_pod_in_the_cluster> -c elasticsearch -- es_util --query="_flush/synced" -XPOST
```

For example:

```
$ oc exec -c elasticsearch-cdm-5ceex6ts-1-dcd6c4c7c-jpw6 -c elasticsearch -- es_util --query="_flush/synced" -XPOST
```

Example output

```
{"_shards":{"total":4,"successful":4,"failed":0},".security":{"total":2,"successful":2,"failed":0},".kibana_1":{"total":2,"successful":2,"failed":0}}
```

5. Prevent shard balancing when purposely bringing down nodes using the OpenShift Container Platform [es_util](#) tool:

```
$ oc exec <any_es_pod_in_the_cluster> -c elasticsearch -- es_util --query="_cluster/settings" -XPUT -d '{"persistent": {"cluster.routing.allocation.enable": "primaries"} }'
```

For example:


```
$ oc exec elasticsearch-cdm-5ceex6ts-1-dcd6c4c7c-jpw6 -c elasticsearch -- es_util --
query="_cluster/settings" -XPUT -d '{ "persistent": { "cluster.routing.allocation.enable" :
"primaries" } }'
```

Example output

```
{"acknowledged":true,"persistent":{"cluster":{"routing":{"allocation":
{"enable":"primaries"}}},"transient":
```

6. After the command is complete, for each deployment you have for an ES cluster:
 - a. By default, the OpenShift Container Platform Elasticsearch cluster blocks rollouts to their nodes. Use the following command to allow rollouts and allow the pod to pick up the changes:

```
$ oc rollout resume deployment/<deployment-name>
```

For example:

```
$ oc rollout resume deployment/elasticsearch-cdm-0-1
```

Example output

```
deployment.extensions/elasticsearch-cdm-0-1 resumed
```

A new pod is deployed. After the pod has a ready container, you can move on to the next deployment.

```
$ oc get pods -l component=elasticsearch-
```

Example output

NAME	READY	STATUS	RESTARTS	AGE
elasticsearch-cdm-5ceex6ts-1-dcd6c4c7c-jpw6k	2/2	Running	0	22h
elasticsearch-cdm-5ceex6ts-2-f799564cb-l9mj7	2/2	Running	0	22h
elasticsearch-cdm-5ceex6ts-3-585968dc68-k7kjr	2/2	Running	0	22h

- b. After the deployments are complete, reset the pod to disallow rollouts:

```
$ oc rollout pause deployment/<deployment-name>
```

For example:

```
$ oc rollout pause deployment/elasticsearch-cdm-0-1
```

Example output

```
deployment.extensions/elasticsearch-cdm-0-1 paused
```

- c. Check that the Elasticsearch cluster is in a **green** or **yellow** state:

```
$ oc exec <any_es_pod_in_the_cluster> -c elasticsearch -- es_util --
query=_cluster/health?pretty=true
```

**NOTE**

If you performed a rollout on the Elasticsearch pod you used in the previous commands, the pod no longer exists and you need a new pod name here.

For example:

```
$ oc exec elasticsearch-cdm-5ceex6ts-1-dcd6c4c7c-jpw6 -c elasticsearch -- es_util --
query=_cluster/health?pretty=true
```

```
{
  "cluster_name" : "elasticsearch",
  "status" : "yellow", 1
  "timed_out" : false,
  "number_of_nodes" : 3,
  "number_of_data_nodes" : 3,
  "active_primary_shards" : 8,
  "active_shards" : 16,
  "relocating_shards" : 0,
  "initializing_shards" : 0,
  "unassigned_shards" : 1,
  "delayed_unassigned_shards" : 0,
  "number_of_pending_tasks" : 0,
  "number_of_in_flight_fetch" : 0,
  "task_max_waiting_in_queue_millis" : 0,
  "active_shards_percent_as_number" : 100.0
}
```

1 Make sure this parameter value is **green** or **yellow** before proceeding.

7. If you changed the Elasticsearch configuration map, repeat these steps for each Elasticsearch pod.
8. After all the deployments for the cluster have been rolled out, re-enable shard balancing:

```
$ oc exec <any_es_pod_in_the_cluster> -c elasticsearch -- es_util --
query="_cluster/settings" -XPUT -d '{"persistent": {"cluster.routing.allocation.enable" : "all" }
}'
```

For example:

```
$ oc exec elasticsearch-cdm-5ceex6ts-1-dcd6c4c7c-jpw6 -c elasticsearch -- es_util --
query="_cluster/settings" -XPUT -d '{"persistent": {"cluster.routing.allocation.enable" : "all" }
}'
```

Example output

```
{
  "acknowledged" : true,
```

```

    "persistent" : {},
    "transient" : {
      "cluster" : {
        "routing" : {
          "allocation" : {
            "enable" : "all"
          }
        }
      }
    }
  }
}
}
}
}
}

```

- Scale up the collector pods so they send new logs to Elasticsearch.

```

$ oc -n openshift-logging patch daemonset/collector -p '{"spec":{"template":{"spec":{"nodeSelector":{"logging-infra-collector": "true"}}}}}'

```

4.3.9. Exposing the log store service as a route

By default, the log store that is deployed with the logging subsystem for Red Hat OpenShift is not accessible from outside the logging cluster. You can enable a route with re-encryption termination for external access to the log store service for those tools that access its data.

Externally, you can access the log store by creating a reencrypt route, your OpenShift Container Platform token and the installed log store CA certificate. Then, access a node that hosts the log store service with a cURL request that contains:

- The **Authorization: Bearer \${token}**
- The Elasticsearch reencrypt route and an [Elasticsearch API request](#).

Internally, you can access the log store service using the log store cluster IP, which you can get by using either of the following commands:

```

$ oc get service elasticsearch -o jsonpath={.spec.clusterIP} -n openshift-logging

```

Example output

```

172.30.183.229

```

```

$ oc get service elasticsearch -n openshift-logging

```

Example output

```

NAME          TYPE          CLUSTER-IP    EXTERNAL-IP  PORT(S)  AGE
elasticsearch ClusterIP     172.30.183.229 <none>      9200/TCP  22h

```

You can check the cluster IP address with a command similar to the following:

```

$ oc exec elasticsearch-cdm-oplnhinv-1-5746475887-fj2f8 -n openshift-logging -- curl -tlsv1.2 --insecure -H "Authorization: Bearer ${token}" "https://172.30.183.229:9200/_cat/health"

```

Example output

```
% Total  % Received % Xferd  Average Speed   Time    Time     Time  Current
           Dload  Upload   Total   Spent    Left  Speed
100  29 100  29  0  0  108  0  --:--:-- --:--:-- --:--:--  108
```

Prerequisites

- The Red Hat OpenShift Logging and Elasticsearch Operators must be installed.
- You must have access to the project to be able to access to the logs.

Procedure

To expose the log store externally:

1. Change to the **openshift-logging** project:

```
$ oc project openshift-logging
```

2. Extract the CA certificate from the log store and write to the **admin-ca** file:

```
$ oc extract secret/elasticsearch --to=. --keys=admin-ca
```

Example output

```
admin-ca
```

3. Create the route for the log store service as a YAML file:
 - a. Create a YAML file with the following:

```
apiVersion: route.openshift.io/v1
kind: Route
metadata:
  name: elasticsearch
  namespace: openshift-logging
spec:
  host:
  to:
    kind: Service
    name: elasticsearch
  tls:
    termination: reencrypt
    destinationCACertificate: | 1
```

- 1** Add the log store CA certificate or use the command in the next step. You do not have to set the **spec.tls.key**, **spec.tls.certificate**, and **spec.tls.caCertificate** parameters required by some reencrypt routes.

- b. Run the following command to add the log store CA certificate to the route YAML you created in the previous step:

```
$ cat ./admin-ca | sed -e "s/^/ /" >> <file-name>.yaml
```

- c. Create the route:

```
$ oc create -f <file-name>.yaml
```

Example output

```
route.route.openshift.io/elasticsearch created
```

4. Check that the Elasticsearch service is exposed:

- a. Get the token of this service account to be used in the request:

```
$ token=$(oc whoami -t)
```

- b. Set the **elasticsearch** route you created as an environment variable.

```
$ routeES=`oc get route elasticsearch -o jsonpath={.spec.host}`
```

- c. To verify the route was successfully created, run the following command that accesses Elasticsearch through the exposed route:

```
curl -tlsv1.2 --insecure -H "Authorization: Bearer ${token}" "https://${routeES}"
```

The response appears similar to the following:

Example output

```
{
  "name" : "elasticsearch-cdm-i40ktba0-1",
  "cluster_name" : "elasticsearch",
  "cluster_uuid" : "0eY-tJzcR3KOdpgeMJo-MQ",
  "version" : {
    "number" : "6.8.1",
    "build_flavor" : "oss",
    "build_type" : "zip",
    "build_hash" : "Unknown",
    "build_date" : "Unknown",
    "build_snapshot" : true,
    "lucene_version" : "7.7.0",
    "minimum_wire_compatibility_version" : "5.6.0",
    "minimum_index_compatibility_version" : "5.0.0"
  },
  "<tagline>" : "<for search>"
}
```

4.4. CONFIGURING THE LOG VISUALIZER

OpenShift Container Platform uses Kibana to display the log data collected by the logging subsystem.

You can scale Kibana for redundancy and configure the CPU and memory for your Kibana nodes.

4.4.1. Configuring CPU and memory limits

The logging subsystem components allow for adjustments to both the CPU and memory limits.

Procedure

1. Edit the **ClusterLogging** custom resource (CR) in the **openshift-logging** project:

```
$ oc -n openshift-logging edit ClusterLogging instance
```

```
apiVersion: "logging.openshift.io/v1"
kind: "ClusterLogging"
metadata:
  name: "instance"
  namespace: openshift-logging
...
spec:
  managementState: "Managed"
  logStore:
    type: "elasticsearch"
    elasticsearch:
      nodeCount: 3
      resources: 1
      limits:
        memory: 16Gi
      requests:
        cpu: 200m
        memory: 16Gi
    storage:
      storageClassName: "gp2"
      size: "200G"
      redundancyPolicy: "SingleRedundancy"
  visualization:
    type: "kibana"
    kibana:
      resources: 2
      limits:
        memory: 1Gi
      requests:
        cpu: 500m
        memory: 1Gi
    proxy:
      resources: 3
      limits:
        memory: 100Mi
      requests:
        cpu: 100m
        memory: 100Mi
    replicas: 2
  collection:
    logs:
      type: "fluentd"
      fluentd:
```

```
resources: 4
  limits:
    memory: 736Mi
  requests:
    cpu: 200m
    memory: 736Mi
```

- 1 Specify the CPU and memory limits and requests for the log store as needed. For Elasticsearch, you must adjust both the request value and the limit value.
- 2 3 Specify the CPU and memory limits and requests for the log visualizer as needed.
- 4 Specify the CPU and memory limits and requests for the log collector as needed.

4.4.2. Scaling redundancy for the log visualizer nodes

You can scale the pod that hosts the log visualizer for redundancy.

Procedure

1. Edit the **ClusterLogging** custom resource (CR) in the **openshift-logging** project:

```
$ oc edit ClusterLogging instance

$ oc edit ClusterLogging instance

apiVersion: "logging.openshift.io/v1"
kind: "ClusterLogging"
metadata:
  name: "instance"

...

spec:
  visualization:
    type: "kibana"
  kibana:
    replicas: 1 1
```


- 1 Specify the number of Kibana nodes.

4.5. CONFIGURING LOGGING SUBSYSTEM STORAGE

Elasticsearch is a memory-intensive application. The default logging subsystem installation deploys 16G of memory for both memory requests and memory limits. The initial set of OpenShift Container Platform nodes might not be large enough to support the Elasticsearch cluster. You must add additional nodes to the OpenShift Container Platform cluster to run with the recommended or higher memory. Each Elasticsearch node can operate with a lower memory setting, though this is not recommended for production environments.

4.5.1. Storage considerations for the logging subsystem for Red Hat OpenShift

A persistent volume is required for each Elasticsearch deployment configuration. On OpenShift Container Platform this is achieved using persistent volume claims.



NOTE

If you use a local volume for persistent storage, do not use a raw block volume, which is described with **volumeMode: block** in the **LocalVolume** object. Elasticsearch cannot use raw block volumes.

The OpenShift Elasticsearch Operator names the PVCs using the Elasticsearch resource name.

Fluentd ships any logs from **systemd journal** and **/var/log/containers/** to Elasticsearch.

Elasticsearch requires sufficient memory to perform large merge operations. If it does not have enough memory, it becomes unresponsive. To avoid this problem, evaluate how much application log data you need, and allocate approximately double that amount of free storage capacity.

By default, when storage capacity is 85% full, Elasticsearch stops allocating new data to the node. At 90%, Elasticsearch attempts to relocate existing shards from that node to other nodes if possible. But if no nodes have a free capacity below 85%, Elasticsearch effectively rejects creating new indices and becomes RED.



NOTE

These low and high watermark values are Elasticsearch defaults in the current release. You can modify these default values. Although the alerts use the same default values, you cannot change these values in the alerts.

4.5.2. Additional resources

- [Configuring persistent storage for the log store](#)

4.6. CONFIGURING CPU AND MEMORY LIMITS FOR LOGGING SUBSYSTEM COMPONENTS

You can configure both the CPU and memory limits for each of the logging subsystem components as needed.

4.6.1. Configuring CPU and memory limits

The logging subsystem components allow for adjustments to both the CPU and memory limits.

Procedure

1. Edit the **ClusterLogging** custom resource (CR) in the **openshift-logging** project:

```
$ oc -n openshift-logging edit ClusterLogging instance
```

```
apiVersion: "logging.openshift.io/v1"
kind: "ClusterLogging"
metadata:
  name: "instance"
  namespace: openshift-logging
```



```

...
spec:
  managementState: "Managed"
  logStore:
    type: "elasticsearch"
    elasticsearch:
      nodeCount: 3
      resources: ①
      limits:
        memory: 16Gi
      requests:
        cpu: 200m
        memory: 16Gi
    storage:
      storageClassName: "gp2"
      size: "200G"
      redundancyPolicy: "SingleRedundancy"
  visualization:
    type: "kibana"
    kibana:
      resources: ②
      limits:
        memory: 1Gi
      requests:
        cpu: 500m
        memory: 1Gi
    proxy:
      resources: ③
      limits:
        memory: 100Mi
      requests:
        cpu: 100m
        memory: 100Mi
  replicas: 2
  collection:
    logs:
      type: "fluentd"
      fluentd:
        resources: ④
        limits:
          memory: 736Mi
        requests:
          cpu: 200m
          memory: 736Mi

```

- ① Specify the CPU and memory limits and requests for the log store as needed. For Elasticsearch, you must adjust both the request value and the limit value.
- ② ③ Specify the CPU and memory limits and requests for the log visualizer as needed.
- ④ Specify the CPU and memory limits and requests for the log collector as needed.

4.7. USING TOLERATIONS TO CONTROL OPENSIFT LOGGING POD PLACEMENT

You can use taints and tolerations to ensure that logging subsystem pods run on specific nodes and that no other workload can run on those nodes.

Taints and tolerations are simple **key:value** pair. A taint on a node instructs the node to repel all pods that do not tolerate the taint.

The **key** is any string, up to 253 characters and the **value** is any string up to 63 characters. The string must begin with a letter or number, and may contain letters, numbers, hyphens, dots, and underscores.

Sample logging subsystem CR with tolerations

```

apiVersion: "logging.openshift.io/v1"
kind: "ClusterLogging"
metadata:
  name: "instance"
  namespace: openshift-logging
...
spec:
  managementState: "Managed"
  logStore:
    type: "elasticsearch"
    elasticsearch:
      nodeCount: 3
      tolerations: 1
      - key: "logging"
        operator: "Exists"
        effect: "NoExecute"
        tolerationSeconds: 6000
    resources:
      limits:
        memory: 16Gi
      requests:
        cpu: 200m
        memory: 16Gi
    storage: {}
    redundancyPolicy: "ZeroRedundancy"
  visualization:
    type: "kibana"
    kibana:
      tolerations: 2
      - key: "logging"
        operator: "Exists"
        effect: "NoExecute"
        tolerationSeconds: 6000
    resources:
      limits:
        memory: 2Gi
      requests:
        cpu: 100m
        memory: 1Gi

```

```

replicas: 1
collection:
logs:
  type: "fluentd"
  fluentd:
    tolerations: 3
    - key: "logging"
      operator: "Exists"
      effect: "NoExecute"
      tolerationSeconds: 6000
resources:
  limits:
    memory: 2Gi
  requests:
    cpu: 100m
    memory: 1Gi

```

- 1 This toleration is added to the Elasticsearch pods.
- 2 This toleration is added to the Kibana pod.
- 3 This toleration is added to the logging collector pods.

4.7.1. Using tolerations to control the log store pod placement

You can control which nodes the log store pods runs on and prevent other workloads from using those nodes by using tolerations on the pods.

You apply tolerations to the log store pods through the **ClusterLogging** custom resource (CR) and apply taints to a node through the node specification. A taint on a node is a **key:value pair** that instructs the node to repel all pods that do not tolerate the taint. Using a specific **key:value** pair that is not on other pods ensures only the log store pods can run on that node.

By default, the log store pods have the following toleration:

```

tolerations:
- effect: "NoExecute"
  key: "node.kubernetes.io/disk-pressure"
  operator: "Exists"

```

Prerequisites

- The Red Hat OpenShift Logging and Elasticsearch Operators must be installed.

Procedure

1. Use the following command to add a taint to a node where you want to schedule the OpenShift Logging pods:

```
$ oc adm taint nodes <node-name> <key>=<value>:<effect>
```

For example:

```
$ oc adm taint nodes node1 elasticsearch=node:NoExecute
```

This example places a taint on **node1** that has key **elasticsearch**, value **node**, and taint effect **NoExecute**. Nodes with the **NoExecute** effect schedule only pods that match the taint and remove existing pods that do not match.

2. Edit the **logstore** section of the **ClusterLogging** CR to configure a toleration for the Elasticsearch pods:

```
logStore:
  type: "elasticsearch"
  elasticsearch:
    nodeCount: 1
    tolerations:
      - key: "elasticsearch" 1
        operator: "Exists" 2
        effect: "NoExecute" 3
        tolerationSeconds: 6000 4
```

- 1 Specify the key that you added to the node.
- 2 Specify the **Exists** operator to require a taint with the key **elasticsearch** to be present on the Node.
- 3 Specify the **NoExecute** effect.
- 4 Optionally, specify the **tolerationSeconds** parameter to set how long a pod can remain bound to a node before being evicted.

This toleration matches the taint created by the **oc adm taint** command. A pod with this toleration could be scheduled onto **node1**.

4.7.2. Using tolerations to control the log visualizer pod placement

You can control the node where the log visualizer pod runs and prevent other workloads from using those nodes by using tolerations on the pods.

You apply tolerations to the log visualizer pod through the **ClusterLogging** custom resource (CR) and apply taints to a node through the node specification. A taint on a node is a **key:value** pair that instructs the node to repel all pods that do not tolerate the taint. Using a specific **key:value** pair that is not on other pods ensures only the Kibana pod can run on that node.

Prerequisites

- The Red Hat OpenShift Logging and Elasticsearch Operators must be installed.

Procedure

1. Use the following command to add a taint to a node where you want to schedule the log visualizer pod:

```
$ oc adm taint nodes <node-name> <key>=<value>:<effect>
```

For example:

```
$ oc adm taint nodes node1 kibana=node:NoExecute
```

This example places a taint on **node1** that has key **kibana**, value **node**, and taint effect **NoExecute**. You must use the **NoExecute** taint effect. **NoExecute** schedules only pods that match the taint and remove existing pods that do not match.

2. Edit the **visualization** section of the **ClusterLogging** CR to configure a toleration for the Kibana pod:

```
visualization:
  type: "kibana"
  kibana:
    tolerations:
      - key: "kibana" 1
        operator: "Exists" 2
        effect: "NoExecute" 3
        tolerationSeconds: 6000 4
```

- 1 Specify the key that you added to the node.
- 2 Specify the **Exists** operator to require the **key/value/effect** parameters to match.
- 3 Specify the **NoExecute** effect.
- 4 Optionally, specify the **tolerationSeconds** parameter to set how long a pod can remain bound to a node before being evicted.

This toleration matches the taint created by the **oc adm taint** command. A pod with this toleration would be able to schedule onto **node1**.

4.7.3. Using tolerations to control the log collector pod placement

You can ensure which nodes the logging collector pods run on and prevent other workloads from using those nodes by using tolerations on the pods.

You apply tolerations to logging collector pods through the **ClusterLogging** custom resource (CR) and apply taints to a node through the node specification. You can use taints and tolerations to ensure the pod does not get evicted for things like memory and CPU issues.

By default, the logging collector pods have the following toleration:

```
tolerations:
- key: "node-role.kubernetes.io/master"
  operator: "Exists"
  effect: "NoExecute"
```

Prerequisites

- The Red Hat OpenShift Logging and Elasticsearch Operators must be installed.

Procedure

1. Use the following command to add a taint to a node where you want logging collector pods to schedule logging collector pods:

```
$ oc adm taint nodes <node-name> <key>=<value>:<effect>
```

For example:

```
$ oc adm taint nodes node1 collector=node:NoExecute
```

This example places a taint on **node1** that has key **collector**, value **node**, and taint effect **NoExecute**. You must use the **NoExecute** taint effect. **NoExecute** schedules only pods that match the taint and removes existing pods that do not match.

2. Edit the **collection** stanza of the **ClusterLogging** custom resource (CR) to configure a toleration for the logging collector pods:

```
collection:
  logs:
    type: "fluentd"
    fluentd:
      tolerations:
        - key: "collector" 1
          operator: "Exists" 2
          effect: "NoExecute" 3
          tolerationSeconds: 6000 4
```

- 1 Specify the key that you added to the node.
- 2 Specify the **Exists** operator to require the **key/value/effect** parameters to match.
- 3 Specify the **NoExecute** effect.
- 4 Optionally, specify the **tolerationSeconds** parameter to set how long a pod can remain bound to a node before being evicted.

This toleration matches the taint created by the **oc adm taint** command. A pod with this toleration would be able to schedule onto **node1**.

4.7.4. Additional resources

- [Controlling pod placement using node taints](#) .

4.8. MOVING LOGGING SUBSYSTEM RESOURCES WITH NODE SELECTORS

You can use node selectors to deploy the Elasticsearch and Kibana pods to different nodes.

4.8.1. Moving OpenShift Logging resources

You can configure the Cluster Logging Operator to deploy the pods for logging subsystem components, such as Elasticsearch and Kibana, to different nodes. You cannot move the Cluster Logging Operator pod from its installed location.

For example, you can move the Elasticsearch pods to a separate node because of high CPU, memory, and disk requirements.

Prerequisites

- The Red Hat OpenShift Logging and Elasticsearch Operators must be installed. These features are not installed by default.

Procedure

1. Edit the **ClusterLogging** custom resource (CR) in the **openshift-logging** project:

```
$ oc edit ClusterLogging instance

apiVersion: logging.openshift.io/v1
kind: ClusterLogging

...

spec:
  collection:
    logs:
      fluentd:
        resources: null
        type: fluentd
  logStore:
    elasticsearch:
      nodeCount: 3
      nodeSelector: 1
        node-role.kubernetes.io/infra: "
    tolerations:
      - effect: NoSchedule
        key: node-role.kubernetes.io/infra
        value: reserved
      - effect: NoExecute
        key: node-role.kubernetes.io/infra
        value: reserved
    redundancyPolicy: SingleRedundancy
    resources:
      limits:
        cpu: 500m
        memory: 16Gi
      requests:
        cpu: 500m
        memory: 16Gi
      storage: {}
    type: elasticsearch
  managementState: Managed
  visualization:
    kibana:
      nodeSelector: 2
        node-role.kubernetes.io/infra: "
      tolerations:
        - effect: NoSchedule
          key: node-role.kubernetes.io/infra
```

```

value: reserved
- effect: NoExecute
  key: node-role.kubernetes.io/infra
  value: reserved
proxy:
  resources: null
replicas: 1
resources: null
type: kibana
...

```

- 1 2 Add a **nodeSelector** parameter with the appropriate value to the component you want to move. You can use a **nodeSelector** in the format shown or use **<key>: <value>** pairs, based on the value specified for the node. If you added a taint to the infrastructure node, also add a matching toleration.

Verification

To verify that a component has moved, you can use the **oc get pod -o wide** command.

For example:

- You want to move the Kibana pod from the **ip-10-0-147-79.us-east-2.compute.internal** node:

```
$ oc get pod kibana-5b8bdf44f9-ccpq9 -o wide
```

Example output

```

NAME                                READY STATUS RESTARTS AGE IP          NODE
NOMINATED NODE READINESS GATES
kibana-5b8bdf44f9-ccpq9 2/2   Running 0      27s 10.129.2.18 ip-10-0-147-79.us-
east-2.compute.internal <none>    <none>

```

- You want to move the Kibana pod to the **ip-10-0-139-48.us-east-2.compute.internal** node, a dedicated infrastructure node:

```
$ oc get nodes
```

Example output

```

NAME                                STATUS ROLES    AGE  VERSION
ip-10-0-133-216.us-east-2.compute.internal Ready master    60m  v1.22.1
ip-10-0-139-146.us-east-2.compute.internal Ready master    60m  v1.22.1
ip-10-0-139-192.us-east-2.compute.internal Ready worker    51m  v1.22.1
ip-10-0-139-241.us-east-2.compute.internal Ready worker    51m  v1.22.1
ip-10-0-147-79.us-east-2.compute.internal Ready worker    51m  v1.22.1
ip-10-0-152-241.us-east-2.compute.internal Ready master    60m  v1.22.1
ip-10-0-139-48.us-east-2.compute.internal Ready infra     51m  v1.22.1

```

Note that the node has a **node-role.kubernetes.io/infra: "** label:

```
$ oc get node ip-10-0-139-48.us-east-2.compute.internal -o yaml
```


Example output

```

kind: Node
apiVersion: v1
metadata:
  name: ip-10-0-139-48.us-east-2.compute.internal
  selfLink: /api/v1/nodes/ip-10-0-139-48.us-east-2.compute.internal
  uid: 62038aa9-661f-41d7-ba93-b5f1b6ef8751
  resourceVersion: '39083'
  creationTimestamp: '2020-04-13T19:07:55Z'
  labels:
    node-role.kubernetes.io/infra: "
...

```

- To move the Kibana pod, edit the **ClusterLogging** CR to add a node selector:

```

apiVersion: logging.openshift.io/v1
kind: ClusterLogging

...

spec:

...

visualization:
  kibana:
    nodeSelector: ❶
    node-role.kubernetes.io/infra: "
    proxy:
      resources: null
    replicas: 1
    resources: null
    type: kibana

```

- ❶ Add a node selector to match the label in the node specification.

- After you save the CR, the current Kibana pod is terminated and new pod is deployed:

```
$ oc get pods
```

Example output

NAME	READY	STATUS	RESTARTS	AGE
cluster-logging-operator-84d98649c4-zb9g7	1/1	Running	0	29m
elasticsearch-cdm-hwv01pf7-1-56588f554f-kpmlg	2/2	Running	0	28m
elasticsearch-cdm-hwv01pf7-2-84c877d75d-75wqj	2/2	Running	0	28m
elasticsearch-cdm-hwv01pf7-3-f5d95b87b-4nx78	2/2	Running	0	28m
fluentd-42dzz	1/1	Running	0	28m
fluentd-d74rq	1/1	Running	0	28m
fluentd-m5vr9	1/1	Running	0	28m
fluentd-nkx17	1/1	Running	0	28m
fluentd-pdvqb	1/1	Running	0	28m

```

fluentd-tflh6                1/1  Running    0    28m
kibana-5b8bdf44f9-ccpq9      2/2  Terminating 0    4m11s
kibana-7d85dcffc8-bfpfp      2/2  Running    0    33s

```

- The new pod is on the **ip-10-0-139-48.us-east-2.compute.internal** node:

```
$ oc get pod kibana-7d85dcffc8-bfpfp -o wide
```

Example output

```

NAME                READY  STATUS   RESTARTS  AGE  IP            NODE
NOMINATED NODE     READINESS GATES
kibana-7d85dcffc8-bfpfp 2/2    Running    0         43s  10.131.0.22  ip-10-0-139-48.us-east-2.compute.internal <none>

```

- After a few moments, the original Kibana pod is removed.

```
$ oc get pods
```

Example output

```

NAME                READY  STATUS   RESTARTS  AGE
cluster-logging-operator-84d98649c4-zb9g7 1/1    Running    0         30m
elasticsearch-cdm-hwv01pf7-1-56588f554f-kpmlg 2/2    Running    0         29m
elasticsearch-cdm-hwv01pf7-2-84c877d75d-75wqj 2/2    Running    0         29m
elasticsearch-cdm-hwv01pf7-3-f5d95b87b-4nx78 2/2    Running    0         29m
fluentd-42dzz                1/1    Running    0         29m
fluentd-d74rq                1/1    Running    0         29m
fluentd-m5vr9                1/1    Running    0         29m
fluentd-nkxl7                1/1    Running    0         29m
fluentd-pdvqb                1/1    Running    0         29m
fluentd-tflh6                1/1    Running    0         29m
kibana-7d85dcffc8-bfpfp      2/2    Running    0         62s

```

4.9. CONFIGURING SYSTEMD-JOURNALD AND FLUENTD

Because Fluentd reads from the journal, and the journal default settings are very low, journal entries can be lost because the journal cannot keep up with the logging rate from system services.

We recommend setting **RateLimitIntervalSec=30s** and **RateLimitBurst=10000** (or even higher if necessary) to prevent the journal from losing entries.

4.9.1. Configuring systemd-journald for OpenShift Logging


As you scale up your project, the default logging environment might need some adjustments.

For example, if you are missing logs, you might have to increase the rate limits for journald. You can adjust the number of messages to retain for a specified period of time to ensure that OpenShift Logging does not use excessive resources without dropping logs.

You can also determine if you want the logs compressed, how long to retain logs, how or if the logs are stored, and other settings.

Procedure

1. Create a Butane config file, **40-worker-custom-journald.bu**, that includes an **/etc/systemd/journald.conf** file with the required settings.



NOTE

See "Creating machine configs with Butane" for information about Butane.

```
variant: openshift
version: 4.9.0
metadata:
  name: 40-worker-custom-journald
  labels:
    machineconfiguration.openshift.io/role: "worker"
storage:
  files:
    - path: /etc/systemd/journald.conf
      mode: 0644 1
      overwrite: true
      contents:
        inline: |
          Compress=yes 2
          ForwardToConsole=no 3
          ForwardToSyslog=no
          MaxRetentionSec=1month 4
          RateLimitBurst=10000 5
          RateLimitIntervalSec=30s
          Storage=persistent 6
          SyncIntervalSec=1s 7
          SystemMaxUse=8G 8
          SystemKeepFree=20% 9
          SystemMaxFileSize=10M 10
```

- 1 Set the permissions for the **journal.conf** file. It is recommended to set **0644** permissions.
- 2 Specify whether you want logs compressed before they are written to the file system. Specify **yes** to compress the message or **no** to not compress. The default is **yes**.
- 3 Configure whether to forward log messages. Defaults to **no** for each. Specify:
 - **ForwardToConsole** to forward logs to the system console.
 - **ForwardToKsmg** to forward logs to the kernel log buffer.
 - **ForwardToSyslog** to forward to a syslog daemon.
 - **ForwardToWall** to forward messages as wall messages to all logged-in users.
- 4 Specify the maximum time to store journal entries. Enter a number to specify seconds. Or include a unit: "year", "month", "week", "day", "h" or "m". Enter **0** to disable. The default is **1month**.
- 5

Configure rate limiting. If more logs are received than what is specified in **RateLimitBurst** during the time interval defined by **RateLimitIntervalSec**, all further messages within the

- 6 Specify how logs are stored. The default is **persistent**:
 - **volatile** to store logs in memory in `/var/log/journal/`.
 - **persistent** to store logs to disk in `/var/log/journal/`. `systemd` creates the directory if it does not exist.
 - **auto** to store logs in `/var/log/journal/` if the directory exists. If it does not exist, `systemd` temporarily stores logs in `/run/systemd/journal`.
 - **none** to not store logs. `systemd` drops all logs.
- 7 Specify the timeout before synchronizing journal files to disk for **ERR**, **WARNING**, **NOTICE**, **INFO**, and **DEBUG** logs. `systemd` immediately syncs after receiving a **CRIT**, **ALERT**, or **EMERG** log. The default is **1s**.
- 8 Specify the maximum size the journal can use. The default is **8G**.
- 9 Specify how much disk space `systemd` must leave free. The default is **20%**.
- 10 Specify the maximum size for individual journal files stored persistently in `/var/log/journal`. The default is **10M**.



NOTE

If you are removing the rate limit, you might see increased CPU utilization on the system logging daemons as it processes any messages that would have previously been throttled.

For more information on `systemd` settings, see <https://www.freedesktop.org/software/systemd/man/journald.conf.html>. The default settings listed on that page might not apply to OpenShift Container Platform.

2. Use Butane to generate a **MachineConfig** object file, **40-worker-custom-journald.yaml**, containing the configuration to be delivered to the nodes:

```
$ butane 40-worker-custom-journald.bu -o 40-worker-custom-journald.yaml
```

3. Apply the machine config. For example:

```
$ oc apply -f 40-worker-custom-journald.yaml
```

The controller detects the new **MachineConfig** object and generates a new **rendered-worker-`<hash>`** version.

4. Monitor the status of the rollout of the new rendered configuration to each node:

```
$ oc describe machineconfigpool/worker
```

Example output

```

Name:      worker
Namespace:
Labels:    machineconfiguration.openshift.io/mco-built-in=
Annotations: <none>
API Version: machineconfiguration.openshift.io/v1
Kind:      MachineConfigPool

...

Conditions:
Message:
Reason:      All nodes are updating to rendered-worker-
913514517bcea7c93bd446f4830bc64e

```

4.10. MAINTENANCE AND SUPPORT

4.10.1. About unsupported configurations

The supported way of configuring the logging subsystem for Red Hat OpenShift is by configuring it using the options described in this documentation. Do not use other configurations, as they are unsupported. Configuration paradigms might change across OpenShift Container Platform releases, and such cases can only be handled gracefully if all configuration possibilities are controlled. If you use configurations other than those described in this documentation, your changes will disappear because the OpenShift Elasticsearch Operator and Red Hat OpenShift Logging Operator reconcile any differences. The Operators reverse everything to the defined state by default and by design.



NOTE

If you *must* perform configurations not described in the OpenShift Container Platform documentation, you *must* set your Red Hat OpenShift Logging Operator or OpenShift Elasticsearch Operator to **Unmanaged**. An unmanaged OpenShift Logging environment is *not supported* and does not receive updates until you return OpenShift Logging to **Managed**.

4.10.2. Unsupported configurations

You must set the Red Hat OpenShift Logging Operator to the unmanaged state to modify the following components:

- The **Elasticsearch** CR
- The Kibana deployment
- The **fluent.conf** file
- The Fluentd daemon set

You must set the OpenShift Elasticsearch Operator to the unmanaged state to modify the following component:

- the Elasticsearch deployment files.

Explicitly unsupported cases include:

- **Configuring default log rotation** You cannot modify the default log rotation configuration.
- **Configuring the collected log location** You cannot change the location of the log collector output file, which by default is `/var/log/fluentd/fluentd.log`.
- **Throttling log collection.** You cannot throttle down the rate at which the logs are read in by the log collector.
- **Configuring the logging collector using environment variables** You cannot use environment variables to modify the log collector.
- **Configuring how the log collector normalizes logs** You cannot modify default log normalization.

4.10.3. Support policy for unmanaged Operators

The *management state* of an Operator determines whether an Operator is actively managing the resources for its related component in the cluster as designed. If an Operator is set to an *unmanaged* state, it does not respond to changes in configuration nor does it receive updates.

While this can be helpful in non-production clusters or during debugging, Operators in an unmanaged state are unsupported and the cluster administrator assumes full control of the individual component configurations and upgrades.

An Operator can be set to an unmanaged state using the following methods:

- **Individual Operator configuration**
Individual Operators have a **managementState** parameter in their configuration. This can be accessed in different ways, depending on the Operator. For example, the Red Hat OpenShift Logging Operator accomplishes this by modifying a custom resource (CR) that it manages, while the Cluster Samples Operator uses a cluster-wide configuration resource.

Changing the **managementState** parameter to **Unmanaged** means that the Operator is not actively managing its resources and will take no action related to the related component. Some Operators might not support this management state as it might damage the cluster and require manual recovery.



WARNING

Changing individual Operators to the **Unmanaged** state renders that particular component and functionality unsupported. Reported issues must be reproduced in **Managed** state for support to proceed.

- **Cluster Version Operator (CVO) overrides**
The **spec.overrides** parameter can be added to the CVO's configuration to allow administrators to provide a list of overrides to the CVO's behavior for a component. Setting the **spec.overrides[].unmanaged** parameter to **true** for a component blocks cluster upgrades and alerts the administrator after a CVO override has been set:

Disabling ownership via cluster version overrides prevents upgrades. Please remove overrides before continuing.

■

**WARNING**

Setting a CVO override puts the entire cluster in an unsupported state. Reported issues must be reproduced after removing any overrides for support to proceed.

CHAPTER 5. VIEWING LOGS FOR A RESOURCE

You can view the logs for various resources, such as builds, deployments, and pods by using the OpenShift CLI (`oc`) and the web console.



NOTE

Resource logs are a default feature that provides limited log viewing capability. To enhance your log retrieving and viewing experience, it is recommended that you install [OpenShift Logging](#). The logging subsystem aggregates all the logs from your OpenShift Container Platform cluster, such as node system audit logs, application container logs, and infrastructure logs, into a dedicated log store. You can then query, discover, and visualize your log data through the [Kibana interface](#). Resource logs do not access the logging subsystem log store.

5.1. VIEWING RESOURCE LOGS


You can view the log for various resources in the OpenShift CLI (`oc`) and web console. Logs read from the tail, or end, of the log.

Prerequisites

- Access to the OpenShift CLI (`oc`).

Procedure (UI)

1. In the OpenShift Container Platform console, navigate to **Workloads** → **Pods** or navigate to the pod through the resource you want to investigate.



NOTE

Some resources, such as builds, do not have pods to query directly. In such instances, you can locate the **Logs** link on the **Details** page for the resource.

2. Select a project from the drop-down menu.
3. Click the name of the pod you want to investigate.
4. Click **Logs**.

Procedure (CLI)

- View the log for a specific pod:

```
$ oc logs -f <pod_name> -c <container_name>
```

where:

-f

Optional: Specifies that the output follows what is being written into the logs.

<pod_name>

Specifies the name of the pod.

<container_name>

Optional: Specifies the name of a container. When a pod has more than one container, you must specify the container name.

For example:

```
$ oc logs ruby-58cd97df55-mww7r
```

```
$ oc logs -f ruby-57f7f4855b-znl92 -c ruby
```

The contents of log files are printed out.

- View the log for a specific resource:

```
$ oc logs <object_type>/<resource_name> 1
```

- 1** Specifies the resource type and name.

For example:

```
$ oc logs deployment/ruby
```

The contents of log files are printed out.

CHAPTER 6. VIEWING CLUSTER LOGS BY USING KIBANA

The logging subsystem includes a web console for visualizing collected log data. Currently, OpenShift Container Platform deploys the Kibana console for visualization.

Using the log visualizer, you can do the following with your data:

- search and browse the data using the **Discover** tab.
- chart and map the data using the **Visualize** tab.
- create and view custom dashboards using the **Dashboard** tab.

Use and configuration of the Kibana interface is beyond the scope of this documentation. For more information, on using the interface, see the [Kibana documentation](#).



NOTE

The audit logs are not stored in the internal OpenShift Container Platform Elasticsearch instance by default. To view the audit logs in Kibana, you must use the [Log Forwarding API](#) to configure a pipeline that uses the **default** output for audit logs.

6.1. DEFINING KIBANA INDEX PATTERNS

An index pattern defines the Elasticsearch indices that you want to visualize. To explore and visualize data in Kibana, you must create an index pattern.

Prerequisites


- A user must have the **cluster-admin** role, the **cluster-reader** role, or both roles to view the **infra** and **audit** indices in Kibana. The default **kubeadmin** user has proper permissions to view these indices.

If you can view the pods and logs in the **default**, **kube-** and **openshift-** projects, you should be able to access these indices. You can use the following command to check if the current user has appropriate permissions:

```
$ oc auth can-i get pods/log -n <project>
```

Example output

```
yes
```




NOTE

The audit logs are not stored in the internal OpenShift Container Platform Elasticsearch instance by default. To view the audit logs in Kibana, you must use the Log Forwarding API to configure a pipeline that uses the **default** output for audit logs.

- Elasticsearch documents must be indexed before you can create index patterns. This is done automatically, but it might take a few minutes in a new or updated cluster.

Procedure

To define index patterns and create visualizations in Kibana:

1. In the OpenShift Container Platform console, click the Application Launcher  and select **Logging**.
2. Create your Kibana index patterns by clicking **Management** → **Index Patterns** → **Create index pattern**:
 - Each user must manually create index patterns when logging into Kibana the first time to see logs for their projects. Users must create an index pattern named **app** and use the **@timestamp** time field to view their container logs.
 - Each admin user must create index patterns when logged into Kibana the first time for the **app**, **infra**, and **audit** indices using the **@timestamp** time field.
3. Create Kibana Visualizations from the new index patterns.

6.2. VIEWING CLUSTER LOGS IN KIBANA

You view cluster logs in the Kibana web console. The methods for viewing and visualizing your data in Kibana that are beyond the scope of this documentation. For more information, refer to the [Kibana documentation](#).

Prerequisites

- The Red Hat OpenShift Logging and Elasticsearch Operators must be installed.
- Kibana index patterns must exist.
- A user must have the **cluster-admin** role, the **cluster-reader** role, or both roles to view the **infra** and **audit** indices in Kibana. The default **kubeadmin** user has proper permissions to view these indices.

If you can view the pods and logs in the **default**, **kube-** and **openshift-** projects, you should be able to access these indices. You can use the following command to check if the current user has appropriate permissions:

```
$ oc auth can-i get pods/log -n <project>
```

Example output

```
yes
```




NOTE

The audit logs are not stored in the internal OpenShift Container Platform Elasticsearch instance by default. To view the audit logs in Kibana, you must use the Log Forwarding API to configure a pipeline that uses the **default** output for audit logs.

Procedure

To view logs in Kibana:

1. In the OpenShift Container Platform console, click the Application Launcher  and select **Logging**.
2. Log in using the same credentials you use to log in to the OpenShift Container Platform console.
The Kibana interface launches.
3. In Kibana, click **Discover**.
4. Select the index pattern you created from the drop-down menu in the top-left corner: **app**, **audit**, or **infra**.
The log data displays as time-stamped documents.
5. Expand one of the time-stamped documents.
6. Click the **JSON** tab to display the log entry for that document.

Example 6.1. Sample infrastructure log entry in Kibana

```
{
  "_index": "infra-000001",
  "_type": "_doc",
  "_id": "YmJmYTBINDkZTRmLTliMGQtMjE3NmFiOGUyOWM3",
  "_version": 1,
  "_score": null,
  "_source": {
    "docker": {
      "container_id": "f85fa55bbef7bb783f041066be1e7c267a6b88c4603dfce213e32c1"
    },
    "kubernetes": {
      "container_name": "registry-server",
      "namespace_name": "openshift-marketplace",
      "pod_name": "redhat-marketplace-n64gc",
      "container_image": "registry.redhat.io/redhat/redhat-marketplace-index:v4.7",
      "container_image_id": "registry.redhat.io/redhat/redhat-marketplace-
index@sha256:65fc0c45aabb95809e376feb065771ecda9e5e59cc8b3024c4545c168f",
      "pod_id": "8f594ea2-c866-4b5c-a1c8-a50756704b2a",
      "host": "ip-10-0-182-28.us-east-2.compute.internal",
      "master_url": "https://kubernetes.default.svc",
      "namespace_id": "3abab127-7669-4eb3-b9ef-44c04ad68d38",
      "namespace_labels": {
        "openshift_io/cluster-monitoring": "true"
      },
      "flat_labels": [
        "catalogsource_operators_coreos_com/update=redhat-marketplace"
      ]
    },
    "message": "time=\"2020-09-23T20:47:03Z\" level=info msg=\"serving registry\"
database=/database/index.db port=50051",
    "level": "unknown",
    "hostname": "ip-10-0-182-28.internal",
    "pipeline_metadata": {
      "collector": {
        "ipaddr4": "10.0.182.28",
        "inputname": "fluent-plugin-systemd",
        "name": "fluentd",

```

```
"received_at": "2020-09-23T20:47:15.007583+00:00",
"version": "1.7.4 1.6.0"
}
},
"@timestamp": "2020-09-23T20:47:03.422465+00:00",
"viaq_msg_id": "YmJmYTBINDktMDMGQtMjE3NmFiOGUyOWM3",
"openshift": {
  "labels": {
    "logging": "infra"
  }
}
},
"fields": {
  "@timestamp": [
    "2020-09-23T20:47:03.422Z"
  ],
  "pipeline_metadata.collector.received_at": [
    "2020-09-23T20:47:15.007Z"
  ]
},
"sort": [
  1600894023422
]
}
```

CHAPTER 7. FORWARDING LOGS TO EXTERNAL THIRD-PARTY LOGGING SYSTEMS

By default, the logging subsystem sends container and infrastructure logs to the default internal Elasticsearch log store defined in the **ClusterLogging** custom resource. However, it does not send audit logs to the internal store because it does not provide secure storage. If this default configuration meets your needs, you do not need to configure the Cluster Log Forwarder.

To send logs to other log aggregators, you use the OpenShift Container Platform Cluster Log Forwarder. This API enables you to send container, infrastructure, and audit logs to specific endpoints within or outside your cluster. In addition, you can send different types of logs to various systems so that various individuals can access each type. You can also enable Transport Layer Security (TLS) support to send logs securely, as required by your organization.



NOTE

To send audit logs to the default internal Elasticsearch log store, use the Cluster Log Forwarder as described in [Forward audit logs to the log store](#).

When you forward logs externally, the logging subsystem creates or modifies a Fluentd config map to send logs using your desired protocols. You are responsible for configuring the protocol on the external log aggregator.



IMPORTANT

You cannot use the config map methods and the Cluster Log Forwarder in the same cluster.

7.1. ABOUT FORWARDING LOGS TO THIRD-PARTY SYSTEMS

To send logs to specific endpoints inside and outside your OpenShift Container Platform cluster, you specify a combination of *outputs* and *pipelines* in a **ClusterLogForwarder** custom resource (CR). You can also use *inputs* to forward the application logs associated with a specific project to an endpoint. Authentication is provided by a Kubernetes *Secret* object.

output

The destination for log data that you define, or where you want the logs sent. An output can be one of the following types:

- **elasticsearch**. An external Elasticsearch instance. The **elasticsearch** output can use a TLS connection.
- **fluentdForward**. An external log aggregation solution that supports Fluentd. This option uses the Fluentd **forward** protocols. The **fluentdForward** output can use a TCP or TLS connection and supports shared-key authentication by providing a **shared_key** field in a secret. Shared-key authentication can be used with or without TLS.
- **syslog**. An external log aggregation solution that supports the syslog [RFC3164](#) or [RFC5424](#) protocols. The **syslog** output can use a UDP, TCP, or TLS connection.
- **cloudwatch**. Amazon CloudWatch, a monitoring and log storage service hosted by Amazon Web Services (AWS).
- **loki**. Loki, a horizontally scalable, highly available, multi-tenant log aggregation system.

- **kafka**. A Kafka broker. The **kafka** output can use a TCP or TLS connection.
- **default**. The internal OpenShift Container Platform Elasticsearch instance. You are not required to configure the default output. If you do configure a **default** output, you receive an error message because the **default** output is reserved for the Red Hat OpenShift Logging Operator.

pipeline

Defines simple routing from one log type to one or more outputs, or which logs you want to send. The log types are one of the following:

- **application**. Container logs generated by user applications running in the cluster, except infrastructure container applications.
- **infrastructure**. Container logs from pods that run in the **openshift***, **kube***, or **default** projects and journal logs sourced from node file system.
- **audit**. Audit logs generated by the node audit system, **auditd**, Kubernetes API server, OpenShift API server, and OVN network.

You can add labels to outbound log messages by using **key:value** pairs in the pipeline. For example, you might add a label to messages that are forwarded to other data centers or label the logs by type. Labels that are added to objects are also forwarded with the log message.

input

Forwards the application logs associated with a specific project to a pipeline.

In the pipeline, you define which log types to forward using an **inputRef** parameter and where to forward the logs to using an **outputRef** parameter.

Secret

A **key:value map** that contains confidential data such as user credentials.

Note the following:

- If a **ClusterLogForwarder** CR object exists, logs are not forwarded to the default Elasticsearch instance, unless there is a pipeline with the **default** output.
- By default, the logging subsystem sends container and infrastructure logs to the default internal Elasticsearch log store defined in the **ClusterLogging** custom resource. However, it does not send audit logs to the internal store because it does not provide secure storage. If this default configuration meets your needs, do not configure the Log Forwarding API.
- If you do not define a pipeline for a log type, the logs of the undefined types are dropped. For example, if you specify a pipeline for the **application** and **audit** types, but do not specify a pipeline for the **infrastructure** type, **infrastructure** logs are dropped.
- You can use multiple types of outputs in the **ClusterLogForwarder** custom resource (CR) to send logs to servers that support different protocols.
- The internal OpenShift Container Platform Elasticsearch instance does not provide secure storage for audit logs. We recommend you ensure that the system to which you forward audit logs is compliant with your organizational and governmental regulations and is properly secured. The logging subsystem does not comply with those regulations.

The following example forwards the audit logs to a secure external Elasticsearch instance, the infrastructure logs to an insecure external Elasticsearch instance, the application logs to a Kafka broker, and the application logs from the **my-apps-logs** project to the internal Elasticsearch instance.

Sample log forwarding outputs and pipelines

```

apiVersion: "logging.openshift.io/v1"
kind: ClusterLogForwarder
metadata:
  name: instance 1
  namespace: openshift-logging 2
spec:
  outputs:
    - name: elasticsearch-secure 3
      type: "elasticsearch"
      url: https://elasticsearch.secure.com:9200
      secret:
        name: elasticsearch
    - name: elasticsearch-insecure 4
      type: "elasticsearch"
      url: http://elasticsearch.insecure.com:9200
    - name: kafka-app 5
      type: "kafka"
      url: tls://kafka.secure.com:9093/app-topic
  inputs: 6
    - name: my-app-logs
      application:
        namespaces:
          - my-project
  pipelines:
    - name: audit-logs 7
      inputRefs:
        - audit
      outputRefs:
        - elasticsearch-secure
        - default
      parse: json 8
      labels:
        secure: "true" 9
        datacenter: "east"
    - name: infrastructure-logs 10
      inputRefs:
        - infrastructure
      outputRefs:
        - elasticsearch-insecure
      labels:
        datacenter: "west"
    - name: my-app 11
      inputRefs:
        - my-app-logs
      outputRefs:
        - default
    - inputRefs: 12
      - application

```



```
outputRefs:  
- kafka-app  
labels:  
  datacenter: "south"
```

- 1 The name of the **ClusterLogForwarder** CR must be **instance**.
- 2 The namespace for the **ClusterLogForwarder** CR must be **openshift-logging**.
- 3 Configuration for an secure Elasticsearch output using a secret with a secure URL.
 - A name to describe the output.
 - The type of output: **elasticsearch**.
 - The secure URL and port of the Elasticsearch instance as a valid absolute URL, including the prefix.
 - The secret required by the endpoint for TLS communication. The secret must exist in the **openshift-logging** project.
- 4 Configuration for an insecure Elasticsearch output:
 - A name to describe the output.
 - The type of output: **elasticsearch**.
 - The insecure URL and port of the Elasticsearch instance as a valid absolute URL, including the prefix.
- 5 Configuration for a Kafka output using a client-authenticated TLS communication over a secure URL
 - A name to describe the output.
 - The type of output: **kafka**.
 - Specify the URL and port of the Kafka broker as a valid absolute URL, including the prefix.
- 6 Configuration for an input to filter application logs from the **my-project** namespace.
- 7 Configuration for a pipeline to send audit logs to the secure external Elasticsearch instance:
 - A name to describe the pipeline.
 - The **inputRefs** is the log type, in this example **audit**.
 - The **outputRefs** is the name of the output to use, in this example **elasticsearch-secure** to forward to the secure Elasticsearch instance and **default** to forward to the internal Elasticsearch instance.
 - Optional: Labels to add to the logs.
- 8 Optional: Specify whether to forward structured JSON log entries as JSON objects in the **structured** field. The log entry must contain valid structured JSON; otherwise, OpenShift Logging removes the **structured** field and instead sends the log entry to the default index, **app-00000x**.

- 9 Optional: String. One or more labels to add to the logs. Quote values like "true" so they are recognized as string values, not as a boolean.
- 10 Configuration for a pipeline to send infrastructure logs to the insecure external Elasticsearch instance.
- 11 Configuration for a pipeline to send logs from the **my-project** project to the internal Elasticsearch instance.
 - A name to describe the pipeline.
 - The **inputRefs** is a specific input: **my-app-logs**.
 - The **outputRefs** is **default**.
 - Optional: String. One or more labels to add to the logs.
- 12 Configuration for a pipeline to send logs to the Kafka broker, with no pipeline name:
 - The **inputRefs** is the log type, in this example **application**.
 - The **outputRefs** is the name of the output to use.
 - Optional: String. One or more labels to add to the logs.

Fluentd log handling when the external log aggregator is unavailable

If your external logging aggregator becomes unavailable and cannot receive logs, Fluentd continues to collect logs and stores them in a buffer. When the log aggregator becomes available, log forwarding resumes, including the buffered logs. If the buffer fills completely, Fluentd stops collecting logs. OpenShift Container Platform rotates the logs and deletes them. You cannot adjust the buffer size or add a persistent volume claim (PVC) to the Fluentd daemon set or pods.

Supported Authorization Keys

Common key types are provided here. Some output types support additional specialized keys, documented with the output-specific configuration field. All secret keys are optional. Enable the security features you want by setting the relevant keys. You are responsible for creating and maintaining any additional configurations that external destinations might require, such as keys and secrets, service accounts, port openings, or global proxy configuration. Open Shift Logging will not attempt to verify a mismatch between authorization combinations.

Transport Layer Security (TLS)

Using a TLS URL ('http://...' or 'ssl://...') without a Secret enables basic TLS server-side authentication. Additional TLS features are enabled by including a Secret and setting the following optional fields:

- **tls.crt**: (string) File name containing a client certificate. Enables mutual authentication. Requires **tls.key**.
- **tls.key**: (string) File name containing the private key to unlock the client certificate. Requires **tls.crt**.
- **passphrase**: (string) Passphrase to decode an encoded TLS private key. Requires **tls.key**.
- **ca-bundle.crt**: (string) File name of a customer CA for server authentication.

Username and Password

- **username:** (string) Authentication user name. Requires **password**.
- **password:** (string) Authentication password. Requires **username**.

Simple Authentication Security Layer (SASL)

- **sasl.enable** (boolean) Explicitly enable or disable SASL. If missing, SASL is automatically enabled when any of the other **sasl.** keys are set.
- **sasl.mechanisms:** (array) List of allowed SASL mechanism names. If missing or empty, the system defaults are used.
- **sasl.allow-insecure:** (boolean) Allow mechanisms that send clear-text passwords. Defaults to false.

7.1.1. Creating a Secret

You can create a secret in the directory that contains your certificate and key files by using the following command:

```
$ oc create secret generic -n openshift-logging <my-secret> \
--from-file=tls.key=<your_key_file>
--from-file=tls.crt=<your_cert_file>
--from-file=ca-bundle.crt=<your_bundle_file>
--from-literal=username=<your_username>
--from-literal=password=<your_password>
```



NOTE

Generic or opaque secrets are recommended for best results.

7.2. SUPPORTED LOG DATA OUTPUT TYPES IN OPENSIFT LOGGING 5.1

Red Hat OpenShift Logging 5.1 provides the following output types and protocols for sending log data to target log collectors.

Red Hat tests each of the combinations shown in the following table. However, you should be able to send log data to a wider range target log collectors that ingest these protocols.

Output types	Protocols	Tested with
elasticsearch	elasticsearch	Elasticsearch 6.8.1 Elasticsearch 6.8.4 Elasticsearch 7.12.2
fluentdForward	fluentd forward v1	fluentd 1.7.4 logstash 7.10.1

Output types	Protocols	Tested with
kafka	kafka 0.11	kafka 2.4.1 kafka 2.7.0
syslog	RFC-3164, RFC-5424	rsyslog-8.39.0

**NOTE**

Previously, the syslog output supported only RFC-3164. The current syslog output adds support for RFC-5424.

7.3. SUPPORTED LOG DATA OUTPUT TYPES IN OPENSIFT LOGGING 5.2

Red Hat OpenShift Logging 5.2 provides the following output types and protocols for sending log data to target log collectors.

Red Hat tests each of the combinations shown in the following table. However, you should be able to send log data to a wider range target log collectors that ingest these protocols.

Output types	Protocols	Tested with
Amazon CloudWatch	REST over HTTPS	The current version of Amazon CloudWatch
elasticsearch	elasticsearch	Elasticsearch 6.8.1 Elasticsearch 6.8.4 Elasticsearch 7.12.2
fluentdForward	fluentd forward v1	fluentd 1.7.4 logstash 7.10.1
Loki	REST over HTTP and HTTPS	Loki 2.3.0 deployed on OCP and Grafana labs
kafka	kafka 0.11	kafka 2.4.1 kafka 2.7.0
syslog	RFC-3164, RFC-5424	rsyslog-8.39.0

7.4. SUPPORTED LOG DATA OUTPUT TYPES IN OPENSIFT LOGGING 5.3

Red Hat OpenShift Logging 5.3 provides the following output types and protocols for sending log data to target log collectors.

Red Hat tests each of the combinations shown in the following table. However, you should be able to send log data to a wider range target log collectors that ingest these protocols.

Output types	Protocols	Tested with
Amazon CloudWatch	REST over HTTPS	The current version of Amazon CloudWatch
elasticsearch	elasticsearch	Elasticsearch 7.10.1
fluentdForward	fluentd forward v1	fluentd 1.7.4 logstash 7.10.1
Loki	REST over HTTP and HTTPS	Loki 2.2.1 deployed on OCP
kafka	kafka 0.11	kafka 2.7.0
syslog	RFC-3164, RFC-5424	rsyslog-8.39.0

7.5. SUPPORTED LOG DATA OUTPUT TYPES IN OPENSIFT LOGGING 5.4

Red Hat OpenShift Logging 5.4 provides the following output types and protocols for sending log data to target log collectors.

Red Hat tests each of the combinations shown in the following table. However, you should be able to send log data to a wider range target log collectors that ingest these protocols.

Output types	Protocols	Tested with
Amazon CloudWatch	REST over HTTPS	The current version of Amazon CloudWatch
elasticsearch	elasticsearch	Elasticsearch 7.10.1
fluentdForward	fluentd forward v1	fluentd 1.14.5 logstash 7.10.1
Loki	REST over HTTP and HTTPS	Loki 2.2.1 deployed on OCP
kafka	kafka 0.11	kafka 2.7.0
syslog	RFC-3164, RFC-5424	rsyslog-8.39.0

7.6. SUPPORTED LOG DATA OUTPUT TYPES IN OPENSIFT LOGGING 5.5

Red Hat OpenShift Logging 5.5 provides the following output types and protocols for sending log data to target log collectors.

Red Hat tests each of the combinations shown in the following table. However, you should be able to send log data to a wider range target log collectors that ingest these protocols.

Output types	Protocols	Tested with
Amazon CloudWatch	REST over HTTPS	The current version of Amazon CloudWatch
elasticsearch	elasticsearch	Elasticsearch 7.10.1
fluentdForward	fluentd forward v1	fluentd 1.14.6 logstash 7.10.1
Loki	REST over HTTP and HTTPS	Loki 2.5.0 deployed on OCP
kafka	kafka 0.11	kafka 2.7.0
syslog	RFC-3164, RFC-5424	rsyslog-8.39.0

7.7. SUPPORTED LOG DATA OUTPUT TYPES IN OPENSIFT LOGGING 5.6

Red Hat OpenShift Logging 5.6 provides the following output types and protocols for sending log data to target log collectors.

Red Hat tests each of the combinations shown in the following table. However, you should be able to send log data to a wider range target log collectors that ingest these protocols.

Output types	Protocols	Tested with
Amazon CloudWatch	REST over HTTPS	The current version of Amazon CloudWatch
elasticsearch	elasticsearch	Elasticsearch 6.8.23 Elasticsearch 7.10.1 Elasticsearch 8.6.1
fluentdForward	fluentd forward v1	fluentd 1.14.6 logstash 7.10.1

Output types	Protocols	Tested with
Loki	REST over HTTP and HTTPS	Loki 2.5.0 deployed on OCP
kafka	kafka 0.11	kafka 2.7.0
syslog	RFC-3164, RFC-5424	rsyslog-8.39.0



IMPORTANT

Fluentd doesn't support Elasticsearch 8 as of 5.6.2. Vector doesn't support fluentd/logstash/rsyslog before 5.7.0.

7.8. FORWARDING LOGS TO AN EXTERNAL ELASTICSEARCH INSTANCE

You can optionally forward logs to an external Elasticsearch instance in addition to, or instead of, the internal OpenShift Container Platform Elasticsearch instance. You are responsible for configuring the external log aggregator to receive log data from OpenShift Container Platform.

To configure log forwarding to an external Elasticsearch instance, you must create a **ClusterLogForwarder** custom resource (CR) with an output to that instance, and a pipeline that uses the output. The external Elasticsearch output can use the HTTP (insecure) or HTTPS (secure HTTP) connection.

To forward logs to both an external and the internal Elasticsearch instance, create outputs and pipelines to the external instance and a pipeline that uses the **default** output to forward logs to the internal instance. You do not need to create a **default** output. If you do configure a **default** output, you receive an error message because the **default** output is reserved for the Red Hat OpenShift Logging Operator.



NOTE

If you want to forward logs to **only** the internal OpenShift Container Platform Elasticsearch instance, you do not need to create a **ClusterLogForwarder** CR.

Prerequisites

- You must have a logging server that is configured to receive the logging data using the specified protocol or format.

Procedure

- Create or edit a YAML file that defines the **ClusterLogForwarder** CR object:

```
apiVersion: "logging.openshift.io/v1"
kind: ClusterLogForwarder
metadata:
  name: instance 1
  namespace: openshift-logging 2
spec:
  outputs:
```

```

- name: elasticsearch-insecure 3
  type: "elasticsearch" 4
  url: http://elasticsearch.insecure.com:9200 5
- name: elasticsearch-secure
  type: "elasticsearch"
  url: https://elasticsearch.secure.com:9200 6
  secret:
    name: es-secret 7
pipelines:
- name: application-logs 8
  inputRefs: 9
  - application
  - audit
  outputRefs:
  - elasticsearch-secure 10
  - default 11
  parse: json 12
  labels:
    myLabel: "myValue" 13
- name: infrastructure-audit-logs 14
  inputRefs:
  - infrastructure
  outputRefs:
  - elasticsearch-insecure
  labels:
    logs: "audit-infra"

```

- 1 The name of the **ClusterLogForwarder** CR must be **instance**.
- 2 The namespace for the **ClusterLogForwarder** CR must be **openshift-logging**.
- 3 Specify a name for the output.
- 4 Specify the **elasticsearch** type.
- 5 Specify the URL and port of the external Elasticsearch instance as a valid absolute URL. You can use the **http** (insecure) or **https** (secure HTTP) protocol. If the cluster-wide proxy using the CIDR annotation is enabled, the output must be a server name or FQDN, not an IP Address.
- 6 For a secure connection, you can specify an **https** or **http** URL that you authenticate by specifying a **secret**.
- 7 For an **https** prefix, specify the name of the secret required by the endpoint for TLS communication. The secret must exist in the **openshift-logging** project, and must have keys of: **tls.crt**, **tls.key**, and **ca-bundle.crt** that point to the respective certificates that they represent. Otherwise, for **http** and **https** prefixes, you can specify a secret that contains a username and password. For more information, see the following "Example: Setting secret that contains a username and password."
- 8 Optional: Specify a name for the pipeline.
- 9 Specify which log types to forward by using the pipeline: **application**, **infrastructure**, or **audit**.

- 10 Specify the name of the output to use when forwarding logs with this pipeline.
- 11 Optional: Specify the **default** output to send the logs to the internal Elasticsearch instance.
- 12 Optional: Specify whether to forward structured JSON log entries as JSON objects in the **structured** field. The log entry must contain valid structured JSON; otherwise, OpenShift Logging removes the **structured** field and instead sends the log entry to the default index, **app-00000x**.
- 13 Optional: String. One or more labels to add to the logs.
- 14 Optional: Configure multiple outputs to forward logs to other external log aggregators of any supported type:
 - A name to describe the pipeline.
 - The **inputRefs** is the log type to forward by using the pipeline: **application**, **infrastructure**, or **audit**.
 - The **outputRefs** is the name of the output to use.
 - Optional: String. One or more labels to add to the logs.

2. Create the CR object:

```
$ oc create -f <file-name>.yaml
```

Example: Setting a secret that contains a username and password

You can use a secret that contains a username and password to authenticate a secure connection to an external Elasticsearch instance.

For example, if you cannot use mutual TLS (mTLS) keys because a third party operates the Elasticsearch instance, you can use HTTP or HTTPS and set a secret that contains the username and password.

1. Create a **Secret** YAML file similar to the following example. Use base64-encoded values for the **username** and **password** fields. The secret type is opaque by default.

```
apiVersion: v1
kind: Secret
metadata:
  name: openshift-test-secret
data:
  username: dGVzdHVzZXJuYW1lCg==
  password: dGVzdHBhc3N3b3JkCg==
```

2. Create the secret:

```
$ oc create secret -n openshift-logging openshift-test-secret.yaml
```

3. Specify the name of the secret in the **ClusterLogForwarder** CR:

```
kind: ClusterLogForwarder
metadata:
```

```

name: instance
namespace: openshift-logging
spec:
  outputs:
  - name: elasticsearch
    type: "elasticsearch"
    url: https://elasticsearch.secure.com:9200
    secret:
      name: openshift-test-secret

```

**NOTE**

In the value of the **url** field, the prefix can be **http** or **https**.

4. Create the CR object:

```
$ oc create -f <file-name>.yaml
```

7.9. FORWARDING LOGS USING THE FLUENTD FORWARD PROTOCOL

You can use the Fluentd **forward** protocol to send a copy of your logs to an external log aggregator that is configured to accept the protocol instead of, or in addition to, the default Elasticsearch log store. You are responsible for configuring the external log aggregator to receive the logs from OpenShift Container Platform.

To configure log forwarding using the **forward** protocol, you must create a **ClusterLogForwarder** custom resource (CR) with one or more outputs to the Fluentd servers, and pipelines that use those outputs. The Fluentd output can use a TCP (insecure) or TLS (secure TCP) connection.

**NOTE**

Alternately, you can use a config map to forward logs using the **forward** protocols. However, this method is deprecated in OpenShift Container Platform and will be removed in a future release.

Prerequisites

- You must have a logging server that is configured to receive the logging data using the specified protocol or format.

Procedure

1. Create or edit a YAML file that defines the **ClusterLogForwarder** CR object:

```

apiVersion: logging.openshift.io/v1
kind: ClusterLogForwarder
metadata:
  name: instance 1
  namespace: openshift-logging 2
spec:
  outputs:

```

```

- name: fluentd-server-secure ❸
  type: fluentdForward ❹
  url: 'tls://fluentdserver.security.example.com:24224' ❺
  secret: ❻
    name: fluentd-secret
- name: fluentd-server-insecure
  type: fluentdForward
  url: 'tcp://fluentdserver.home.example.com:24224'
pipelines:
- name: forward-to-fluentd-secure ❷
  inputRefs: ❸
    - application
    - audit
  outputRefs:
    - fluentd-server-secure ❹
    - default ❺
  parse: json ❻
  labels:
    clusterId: "C1234" ❻
- name: forward-to-fluentd-insecure ❸
  inputRefs:
    - infrastructure
  outputRefs:
    - fluentd-server-insecure
  labels:
    clusterId: "C1234"

```

- ❶ The name of the **ClusterLogForwarder** CR must be **instance**.
- ❷ The namespace for the **ClusterLogForwarder** CR must be **openshift-logging**.
- ❸ Specify a name for the output.
- ❹ Specify the **fluentdForward** type.
- ❺ Specify the URL and port of the external Fluentd instance as a valid absolute URL. You can use the **tcp** (insecure) or **tls** (secure TCP) protocol. If the cluster-wide proxy using the CIDR annotation is enabled, the output must be a server name or FQDN, not an IP address.
- ❻ If using a **tls** prefix, you must specify the name of the secret required by the endpoint for TLS communication. The secret must exist in the **openshift-logging** project, and must have keys of: **tls.crt**, **tls.key**, and **ca-bundle.crt** that point to the respective certificates that they represent. Otherwise, for http and https prefixes, you can specify a secret that contains a username and password. For more information, see the following "Example: Setting secret that contains a username and password."
- ❼ Optional: Specify a name for the pipeline.
- ❽ Specify which log types to forward by using the pipeline: **application**, **infrastructure**, or **audit**.
- ❾ Specify the name of the output to use when forwarding logs with this pipeline.
- ❿ Optional: Specify the **default** output to forward logs to the internal Elasticsearch instance.

- 11 Optional: Specify whether to forward structured JSON log entries as JSON objects in the **structured** field. The log entry must contain valid structured JSON; otherwise, OpenShift
- 12 Optional: String. One or more labels to add to the logs.
- 13 Optional: Configure multiple outputs to forward logs to other external log aggregators of any supported type:
 - A name to describe the pipeline.
 - The **inputRefs** is the log type to forward by using the pipeline: **application**, **infrastructure**, or **audit**.
 - The **outputRefs** is the name of the output to use.
 - Optional: String. One or more labels to add to the logs.

2. Create the CR object:

```
$ oc create -f <file-name>.yaml
```

7.9.1. Enabling nanosecond precision for Logstash to ingest data from fluentd

For Logstash to ingest log data from fluentd, you must enable nanosecond precision in the Logstash configuration file.

Procedure

- In the Logstash configuration file, set **nanosecond_precision** to **true**.

Example Logstash configuration file

```
input { tcp { codec => fluent { nanosecond_precision => true } port => 24114 } }
filter { }
output { stdout { codec => rubydebug } }
```

7.10. FORWARDING LOGS USING THE SYSLOG PROTOCOL

You can use the **syslog** [RFC3164](#) or [RFC5424](#) protocol to send a copy of your logs to an external log aggregator that is configured to accept the protocol instead of, or in addition to, the default Elasticsearch log store. You are responsible for configuring the external log aggregator, such as a syslog server, to receive the logs from OpenShift Container Platform.

To configure log forwarding using the **syslog** protocol, you must create a **ClusterLogForwarder** custom resource (CR) with one or more outputs to the syslog servers, and pipelines that use those outputs. The syslog output can use a UDP, TCP, or TLS connection.



NOTE

Alternately, you can use a config map to forward logs using the **syslog** RFC3164 protocols. However, this method is deprecated in OpenShift Container Platform and will be removed in a future release.

Prerequisites

- You must have a logging server that is configured to receive the logging data using the specified protocol or format.

Procedure

- Create or edit a YAML file that defines the **ClusterLogForwarder** CR object:

```

apiVersion: logging.openshift.io/v1
kind: ClusterLogForwarder
metadata:
  name: instance 1
  namespace: openshift-logging 2
spec:
  outputs:
    - name: rsyslog-east 3
      type: syslog 4
      syslog: 5
        facility: local0
        rfc: RFC3164
        payloadKey: message
        severity: informational
      url: 'tls://rsyslogserver.east.example.com:514' 6
      secret: 7
        name: syslog-secret
    - name: rsyslog-west
      type: syslog
      syslog:
        appName: myapp
        facility: user
        msgID: mymsg
        proclD: myproc
        rfc: RFC5424
        severity: debug
      url: 'udp://rsyslogserver.west.example.com:514'
  pipelines:
    - name: syslog-east 8
      inputRefs: 9
        - audit
        - application
      outputRefs: 10
        - rsyslog-east
        - default 11
      parse: json 12
      labels:
        secure: "true" 13
        syslog: "east"
    - name: syslog-west 14
      inputRefs:
        - infrastructure
      outputRefs:
        - rsyslog-west

```

```

- default
labels:
  syslog: "west"

```

- 1 The name of the **ClusterLogForwarder** CR must be **instance**.
- 2 The namespace for the **ClusterLogForwarder** CR must be **openshift-logging**.
- 3 Specify a name for the output.
- 4 Specify the **syslog** type.
- 5 Optional: Specify the syslog parameters, listed below.
- 6 Specify the URL and port of the external syslog instance. You can use the **udp** (insecure), **tcp** (insecure) or **tls** (secure TCP) protocol. If the cluster-wide proxy using the CIDR annotation is enabled, the output must be a server name or FQDN, not an IP address.
- 7 If using a **tls** prefix, you must specify the name of the secret required by the endpoint for TLS communication. The secret must exist in the **openshift-logging** project, and must have keys of: **tls.crt**, **tls.key**, and **ca-bundle.crt** that point to the respective certificates that they represent.
- 8 Optional: Specify a name for the pipeline.
- 9 Specify which log types to forward by using the pipeline: **application**, **infrastructure**, or **audit**.
- 10 Specify the name of the output to use when forwarding logs with this pipeline.
- 11 Optional: Specify the **default** output to forward logs to the internal Elasticsearch instance.
- 12 Optional: Specify whether to forward structured JSON log entries as JSON objects in the **structured** field. The log entry must contain valid structured JSON; otherwise, OpenShift Logging removes the **structured** field and instead sends the log entry to the default index, **app-00000x**.
- 13 Optional: String. One or more labels to add to the logs. Quote values like "true" so they are recognized as string values, not as a boolean.
- 14 Optional: Configure multiple outputs to forward logs to other external log aggregators of any supported type:
 - A name to describe the pipeline.
 - The **inputRefs** is the log type to forward by using the pipeline: **application**, **infrastructure**, or **audit**.
 - The **outputRefs** is the name of the output to use.
 - Optional: String. One or more labels to add to the logs.

2. Create the CR object:

```
$ oc create -f <file-name>.yaml
```

7.10.1. Adding log source information to message output

You can add `namespace_name`, `pod_name`, and `container_name` elements to the `message` field of the record by adding the `AddLogSource` field to your `ClusterLogForwarder` custom resource (CR).

```
spec:
  outputs:
  - name: syslogout
    syslog:
      addLogSource: true
      facility: user
      payloadKey: message
      rfc: RFC3164
      severity: debug
      tag: mytag
      type: syslog
      url: tls://syslog-receiver.openshift-logging.svc:24224
    pipelines:
    - inputRefs:
      - application
      name: test-app
      outputRefs:
      - syslogout
```



NOTE

This configuration is compatible with both RFC3164 and RFC5424.

Example syslog message output without `AddLogSource`

```
<15>1 2020-11-15T17:06:14+00:00 fluentd-9hkb4 mytag - - - {"msgcontent"=>"Message Contents",
"timestamp"=>"2020-11-15 17:06:09", "tag_key"=>"rec_tag", "index"=>56}
```

Example syslog message output with `AddLogSource`

```
<15>1 2020-11-16T10:49:37+00:00 crc-j55b9-master-0 mytag - - - namespace_name=clo-test-
6327,pod_name=log-generator-ff9746c49-qxm7l,container_name=log-generator,message=
{"msgcontent":"My life is my message", "timestamp":"2020-11-16 10:49:36", "tag_key":"rec_tag",
"index":76}
```

7.10.2. Syslog parameters

You can configure the following for the `syslog` outputs. For more information, see the syslog [RFC3164](#) or [RFC5424](#) RFC.

- facility: The [syslog facility](#). The value can be a decimal integer or a case-insensitive keyword:
 - **0** or **kern** for kernel messages
 - **1** or **user** for user-level messages, the default.
 - **2** or **mail** for the mail system

- **3** or **daemon** for system daemons
 - **4** or **auth** for security/authentication messages
 - **5** or **syslog** for messages generated internally by syslogd
 - **6** or **lpr** for the line printer subsystem
 - **7** or **news** for the network news subsystem
 - **8** or **uucp** for the UUCP subsystem
 - **9** or **cron** for the clock daemon
 - **10** or **authpriv** for security authentication messages
 - **11** or **ftp** for the FTP daemon
 - **12** or **ntp** for the NTP subsystem
 - **13** or **security** for the syslog audit log
 - **14** or **console** for the syslog alert log
 - **15** or **solaris-cron** for the scheduling daemon
 - **16–23** or **local0** – **local7** for locally used facilities
- Optional: **payloadKey**: The record field to use as payload for the syslog message.



NOTE

Configuring the **payloadKey** parameter prevents other parameters from being forwarded to the syslog.

- **rfc**: The RFC to be used for sending logs using syslog. The default is RFC5424.
- **severity**: The [syslog severity](#) to set on outgoing syslog records. The value can be a decimal integer or a case-insensitive keyword:
 - **0** or **Emergency** for messages indicating the system is unusable
 - **1** or **Alert** for messages indicating action must be taken immediately
 - **2** or **Critical** for messages indicating critical conditions
 - **3** or **Error** for messages indicating error conditions
 - **4** or **Warning** for messages indicating warning conditions
 - **5** or **Notice** for messages indicating normal but significant conditions
 - **6** or **Informational** for messages indicating informational messages
 - **7** or **Debug** for messages indicating debug-level messages, the default
- **tag**: Tag specifies a record field to use as a tag on the syslog message.

- trimPrefix: Remove the specified prefix from the tag.

7.10.3. Additional RFC5424 syslog parameters

The following parameters apply to RFC5424:

- appName: The APP-NAME is a free-text string that identifies the application that sent the log. Must be specified for **RFC5424**.
- msgID: The MSGID is a free-text string that identifies the type of message. Must be specified for **RFC5424**.
- procID: The PROCID is a free-text string. A change in the value indicates a discontinuity in syslog reporting. Must be specified for **RFC5424**.

7.11. FORWARDING LOGS TO AMAZON CLOUDWATCH

You can forward logs to Amazon CloudWatch, a monitoring and log storage service hosted by Amazon Web Services (AWS). You can forward logs to CloudWatch in addition to, or instead of, the default logging subsystem managed Elasticsearch log store.

To configure log forwarding to CloudWatch, you must create a **ClusterLogForwarder** custom resource (CR) with an output for CloudWatch, and a pipeline that uses the output.

Procedure

1. Create a **Secret** YAML file that uses the **aws_access_key_id** and **aws_secret_access_key** fields to specify your base64-encoded AWS credentials. For example:

```
apiVersion: v1
kind: Secret
metadata:
  name: cw-secret
  namespace: openshift-logging
data:
  aws_access_key_id: QUtJQUIPU0ZPRE5ON0VYQU1QTEUK
  aws_secret_access_key:
d0phbHJYVXRuRkVNSS9LN01ERU5HL2JQeFJmaUNZRvBTvBMRUtFWQo=
```

2. Create the secret. For example:

```
$ oc apply -f cw-secret.yaml
```

3. Create or edit a YAML file that defines the **ClusterLogForwarder** CR object. In the file, specify the name of the secret. For example:

```
apiVersion: "logging.openshift.io/v1"
kind: ClusterLogForwarder
metadata:
  name: instance 1
  namespace: openshift-logging 2
spec:
  outputs:
    - name: cw 3
```

```

type: cloudwatch 4
cloudwatch:
  groupBy: logType 5
  groupPrefix: <group prefix> 6
  region: us-east-2 7
secret:
  name: cw-secret 8
pipelines:
- name: infra-logs 9
  inputRefs: 10
  - infrastructure
  - audit
  - application
  outputRefs:
  - cw 11

```

- 1 The name of the **ClusterLogForwarder** CR must be **instance**.
- 2 The namespace for the **ClusterLogForwarder** CR must be **openshift-logging**.
- 3 Specify a name for the output.
- 4 Specify the **cloudwatch** type.
- 5 Optional: Specify how to group the logs:
 - **logType** creates log groups for each log type
 - **namespaceName** creates a log group for each application name space. It also creates separate log groups for infrastructure and audit logs.
 - **namespaceUUID** creates a new log groups for each application namespace UUID. It also creates separate log groups for infrastructure and audit logs.
- 6 Optional: Specify a string to replace the default **infrastructureName** prefix in the names of the log groups.
- 7 Specify the AWS region.
- 8 Specify the name of the secret that contains your AWS credentials.
- 9 Optional: Specify a name for the pipeline.
- 10 Specify which log types to forward by using the pipeline: **application**, **infrastructure**, or **audit**.
- 11 Specify the name of the output to use when forwarding logs with this pipeline.

4. Create the CR object:

```
$ oc create -f <file-name>.yaml
```

Example: Using ClusterLogForwarder with Amazon CloudWatch

Here, you see an example **ClusterLogForwarder** custom resource (CR) and the log data that it outputs to Amazon CloudWatch.

Suppose that you are running an OpenShift Container Platform cluster named **mycluster**. The following command returns the cluster's **infrastructureName**, which you will use to compose **aws** commands later on:

```
$ oc get Infrastructure/cluster -ojson | jq .status.infrastructureName
"mycluster-7977k"
```

To generate log data for this example, you run a **busybox** pod in a namespace called **app**. The **busybox** pod writes a message to stdout every three seconds:

```
$ oc run busybox --image=busybox -- sh -c 'while true; do echo "My life is my message"; sleep 3; done'
$ oc logs -f busybox
My life is my message
My life is my message
My life is my message
...
```

You can look up the UUID of the **app** namespace where the **busybox** pod runs:

```
$ oc get ns/app -ojson | jq .metadata.uid
"794e1e1a-b9f5-4958-a190-e76a9b53d7bf"
```

In your **ClusterLogForwarder** custom resource (CR), you configure the **infrastructure**, **audit**, and **application** log types as inputs to the **all-logs** pipeline. You also connect this pipeline to **cw** output, which forwards the logs to a CloudWatch instance in the **us-east-2** region:

```
apiVersion: "logging.openshift.io/v1"
kind: ClusterLogForwarder
metadata:
  name: instance
  namespace: openshift-logging
spec:
  outputs:
    - name: cw
      type: cloudwatch
      cloudwatch:
        groupBy: logType
        region: us-east-2
      secret:
        name: cw-secret
  pipelines:
    - name: all-logs
      inputRefs:
        - infrastructure
        - audit
        - application
      outputRefs:
        - cw
```

Each region in CloudWatch contains three levels of objects:

- log group
 - log stream
 - log event

With **groupBy: logType** in the **ClusterLogForwarding** CR, the three log types in the **inputRefs** produce three log groups in Amazon Cloudwatch:

```
$ aws --output json logs describe-log-groups | jq .logGroups[].logGroupName
"mycluster-7977k.application"
"mycluster-7977k.audit"
"mycluster-7977k.infrastructure"
```

Each of the log groups contains log streams:

```
$ aws --output json logs describe-log-streams --log-group-name mycluster-7977k.application | jq
.logStreams[].logStreamName
"kubernetes.var.log.containers.busybox_app_busybox-
da085893053e20beddd6747acdbaf98e77c37718f85a7f6a4facf09ca195ad76.log"
```

```
$ aws --output json logs describe-log-streams --log-group-name mycluster-7977k.audit | jq
.logStreams[].logStreamName
"ip-10-0-131-228.us-east-2.compute.internal.k8s-audit.log"
"ip-10-0-131-228.us-east-2.compute.internal.linux-audit.log"
"ip-10-0-131-228.us-east-2.compute.internal.openshift-audit.log"
...
```

```
$ aws --output json logs describe-log-streams --log-group-name mycluster-7977k.infrastructure | jq
.logStreams[].logStreamName
"ip-10-0-131-228.us-east-2.compute.internal.kubernetes.var.log.containers.apiserver-69f9fd9b58-
zqzw5_openshift-oauth-apiserver_oauth-apiserver-
453c5c4ee026fe20a6139ba6b1cdd1bed25989c905bf5ac5ca211b7cbb5c3d7b.log"
"ip-10-0-131-228.us-east-2.compute.internal.kubernetes.var.log.containers.apiserver-797774f7c5-
lftrx_openshift-apiserver_openshift-apiserver-
ce51532df7d4e4d5f21c4f4be05f6575b93196336be0027067fd7d93d70f66a4.log"
"ip-10-0-131-228.us-east-2.compute.internal.kubernetes.var.log.containers.apiserver-797774f7c5-
lftrx_openshift-apiserver_openshift-apiserver-check-endpoints-
82a9096b5931b5c3b1d6dc4b66113252da4a6472c9fff48623baee761911a9ef.log"
...
```

Each log stream contains log events. To see a log event from the **busybox** Pod, you specify its log stream from the **application** log group:

```
$ aws logs get-log-events --log-group-name mycluster-7977k.application --log-stream-name
kubernetes.var.log.containers.busybox_app_busybox-
da085893053e20beddd6747acdbaf98e77c37718f85a7f6a4facf09ca195ad76.log
{
  "events": [
    {
      "timestamp": 1629422704178,
      "message": "{\"docker\":
{\"container_id\":\"da085893053e20beddd6747acdbaf98e77c37718f85a7f6a4facf09ca195ad76\"},\"kub
ernetes\":
```

```
{
  "container_name": "busybox",
  "namespace_name": "app",
  "pod_name": "busybox",
  "container_image": "docker.io/library/busybox:latest",
  "container_image_id": "docker.io/library/busybox@sha256:0f354ec1728d9ff32edcd7d1b8bbdfc798277ad36120dc3dc683be44524c8b60",
  "pod_id": "870be234-90a3-4258-b73f-4f4d6e2777c7",
  "host": "ip-10-0-216-3.us-east-2.compute.internal",
  "labels": {
    "run": "busybox",
    "master_url": "https://kubernetes.default.svc",
    "namespace_id": "794e1e1a-b9f5-4958-a190-e76a9b53d7bf",
    "namespace_labels": {
      "kubernetes_io/metadata_name": "app"
    },
    "message": "My life is my message",
    "level": "unknown",
    "hostname": "ip-10-0-216-3.us-east-2.compute.internal",
    "pipeline_metadata": {
      "collector": {
        "ipaddr4": "10.0.216.3",
        "inputname": "fluent-plugin-systemd",
        "name": "fluentd",
        "received_at": "2021-08-20T01:25:08.085760+00:00",
        "version": "1.7.4 1.6.0"
      },
      "@timestamp": "2021-08-20T01:25:04.178986+00:00",
      "viaq_index_name": "app-write",
      "viaq_msg_id": "NWRjZmUyMWQtZjgzNC00MjI4LTk3MjMtNTk3NmY3ZjU4NDk1",
      "log_type": "application",
      "time": "2021-08-20T01:25:04+00:00"
    },
    "ingestionTime": 1629422744016
  },
  ...
}
```

Example: Customizing the prefix in log group names

In the log group names, you can replace the default **infrastructureName** prefix, **mycluster-7977k**, with an arbitrary string like **demo-group-prefix**. To make this change, you update the **groupPrefix** field in the **ClusterLogForwarding** CR:

```
cloudwatch:
  groupBy: logType
  groupPrefix: demo-group-prefix
  region: us-east-2
```

The value of **groupPrefix** replaces the default **infrastructureName** prefix:

```
$ aws --output json logs describe-log-groups | jq .logGroups[].logGroupName
"demo-group-prefix.application"
"demo-group-prefix.audit"
"demo-group-prefix.infrastructure"
```

Example: Naming log groups after application namespace names

For each application namespace in your cluster, you can create a log group in CloudWatch whose name is based on the name of the application namespace.

If you delete an application namespace object and create a new one that has the same name, CloudWatch continues using the same log group as before.

If you consider successive application namespace objects that have the same name as equivalent to each other, use the approach described in this example. Otherwise, if you need to distinguish the resulting log groups from each other, see the following "Naming log groups for application namespace UUIDs" section instead.

To create application log groups whose names are based on the names of the application namespaces, you set the value of the **groupBy** field to **namespaceName** in the **ClusterLogForwarder** CR:

```
cloudwatch:
  groupBy: namespaceName
  region: us-east-2
```

Setting **groupBy** to **namespaceName** affects the application log group only. It does not affect the **audit** and **infrastructure** log groups.

In Amazon Cloudwatch, the namespace name appears at the end of each log group name. Because there is a single application namespace, "app", the following output shows a new **mycluster-7977k.app** log group instead of **mycluster-7977k.application**:

```
$ aws --output json logs describe-log-groups | jq .logGroups[].logGroupName
"mycluster-7977k.app"
"mycluster-7977k.audit"
"mycluster-7977k.infrastructure"
```

If the cluster in this example had contained multiple application namespaces, the output would show multiple log groups, one for each namespace.

The **groupBy** field affects the application log group only. It does not affect the **audit** and **infrastructure** log groups.

Example: Naming log groups after application namespace UUIDs

For each application namespace in your cluster, you can create a log group in CloudWatch whose name is based on the UUID of the application namespace.

If you delete an application namespace object and create a new one, CloudWatch creates a new log group.

If you consider successive application namespace objects with the same name as different from each other, use the approach described in this example. Otherwise, see the preceding "Example: Naming log groups for application namespace names" section instead.

To name log groups after application namespace UUIDs, you set the value of the **groupBy** field to **namespaceUUID** in the **ClusterLogForwarder** CR:

```
cloudwatch:
  groupBy: namespaceUUID
  region: us-east-2
```

In Amazon Cloudwatch, the namespace UUID appears at the end of each log group name. Because there is a single application namespace, "app", the following output shows a new **mycluster-7977k.794e1e1a-b9f5-4958-a190-e76a9b53d7bf** log group instead of **mycluster-7977k.application**:

```
$ aws --output json logs describe-log-groups | jq .logGroups[].logGroupName
"mycluster-7977k.794e1e1a-b9f5-4958-a190-e76a9b53d7bf" // uid of the "app" namespace
"mycluster-7977k.audit"
"mycluster-7977k.infrastructure"
```

The **groupBy** field affects the application log group only. It does not affect the **audit** and **infrastructure** log groups.

7.12. FORWARDING LOGS TO LOKI

You can forward logs to an external Loki logging system in addition to, or instead of, the internal default OpenShift Container Platform Elasticsearch instance.

To configure log forwarding to Loki, you must create a **ClusterLogForwarder** custom resource (CR) with an output to Loki, and a pipeline that uses the output. The output to Loki can use the HTTP (insecure) or HTTPS (secure HTTP) connection.

Prerequisites

- You must have a Loki logging system running at the URL you specify with the **url** field in the CR.

Procedure

- Create or edit a YAML file that defines the **ClusterLogForwarder** CR object:

```
apiVersion: "logging.openshift.io/v1"
kind: ClusterLogForwarder
metadata:
  name: instance 1
  namespace: openshift-logging 2
spec:
  outputs:
    - name: loki-insecure 3
      type: "loki" 4
      url: http://loki.insecure.com:3100 5
    - name: loki-secure
      type: "loki"
      url: https://loki.secure.com:3100 6
      secret:
        name: loki-secret 7
  pipelines:
    - name: application-logs 8
      inputRefs: 9
      - application
      - audit
      outputRefs:
      - loki-secure 10
  loki:
    tenantKey: kubernetes.namespace_name 11
    labelKeys: kubernetes.labels.foo 12
```

- The name of the **ClusterLogForwarder** CR must be **instance**.
- The namespace for the **ClusterLogForwarder** CR must be **openshift-logging**.
- Specify a name for the output.
- Specify the type as **"loki"**.
- Specify the URL and port of the Loki system as a valid absolute URL. You can use the **http** (insecure) or **https** (secure HTTP) protocol. If the cluster-wide proxy using the CIDR annotation is enabled, the output must be a server name or FQDN, not an IP Address.
- For a secure connection, you can specify an **https** or **http** URL that you authenticate by specifying a **secret**.
-

For an **https** prefix, specify the name of the secret required by the endpoint for TLS communication. The secret must exist in the **openshift-logging** project, and must have

- 8 Optional: Specify a name for the pipeline.
- 9 Specify which log types to forward by using the pipeline: **application**, **infrastructure**, or **audit**.
- 10 Specify the name of the output to use when forwarding logs with this pipeline.
- 11 Optional: Specify a meta-data key field to generate values for the **TenantID** field in Loki. For example, setting **tenantKey: kubernetes.namespace_name** uses the names of the Kubernetes namespaces as values for tenant IDs in Loki. To see which other log record fields you can specify, see the "Log Record Fields" link in the following "Additional resources" section.
- 12 Optional: Specify a list of meta-data field keys to replace the default Loki labels. Loki label names must match the regular expression **[a-zA-Z_:[a-zA-Z0-9_:]***. Illegal characters in meta-data keys are replaced with **_** to form the label name. For example, the **kubernetes.labels.foo** meta-data key becomes Loki label **kubernetes_labels_foo**. If you do not set **labelKeys**, the default value is: **[log_type, kubernetes.namespace_name, kubernetes.pod_name, kubernetes_host]**. Keep the set of labels small because Loki limits the size and number of labels allowed. See [Configuring Loki, limits_config](#). You can still query based on any log record field using query filters.



NOTE


Because Loki requires log streams to be correctly ordered by timestamp, **labelKeys** always includes the **kubernetes_host** label set, even if you do not specify it. This inclusion ensures that each stream originates from a single host, which prevents timestamps from becoming disordered due to clock differences on different hosts.

2. Create the CR object:

```
$ oc create -f <file-name>.yaml
```

7.12.1. Troubleshooting Loki "entry out of order" errors

If your Fluentd forwards a large block of messages to a Loki logging system that exceeds the rate limit, Loki generates "entry out of order" errors. To fix this issue, you update some values in the Loki server configuration file, **loki.yaml**.



NOTE

loki.yaml is not available on Grafana-hosted Loki. This topic does not apply to Grafana-hosted Loki servers.

Conditions

- The **ClusterLogForwarder** custom resource is configured to forward logs to Loki.
- Your system sends a block of messages that is larger than 2 MB to Loki, such as:
 -


```
"values":[[{"1630410392689800468",{"kind":"Event","apiVersion":\
.....
.....
.....
.....
\"received_at\":\"2021-08-31T11:46:32.800278+00:00\",\"version\":\"1.7.4
1.6.0\"}},{\"@timestamp\":\"2021-08-
31T11:46:32.799692+00:00\",\"viaq_index_name\":\"audit-
write\",\"viaq_msg_id\":\"MzFjYjkZjltNjY0MCM0YWU4LWlwMTetNGNmM2E5ZmViMGU4\",\"lo
g_type\":\"audit\"}]]}}
```

- When you enter **oc logs -c fluentd**, the Fluentd logs in your OpenShift Logging cluster show the following messages:

```
429 Too Many Requests Ingestion rate limit exceeded (limit: 8388608 bytes/sec) while
attempting to ingest '2140' lines totaling '3285284' bytes
```

```
429 Too Many Requests Ingestion rate limit exceeded' or '500 Internal Server Error rpc error:
code = ResourceExhausted desc = grpc: received message larger than max (5277702 vs.
4194304)'
```

- When you open the logs on the Loki server, they display **entry out of order** messages like these:

```
,\nentry with timestamp 2021-08-18 05:58:55.061936 +0000 UTC ignored, reason: 'entry out
of order' for stream:
```

```
{fluentd_thread=\"flush_thread_0\", log_type=\"audit\"},\nentry with timestamp 2021-08-18
06:01:18.290229 +0000 UTC ignored, reason: 'entry out of order' for stream:
{fluentd_thread=\"flush_thread_0\", log_type=\"audit\"}
```

Procedure

1. Update the following fields in the **loki.yaml** configuration file on the Loki server with the values shown here:
 - **grpc_server_max_recv_msg_size: 8388608**
 - **chunk_target_size: 8388608**
 - **ingestion_rate_mb: 8**
 - **ingestion_burst_size_mb: 16**
2. Apply the changes in **loki.yaml** to the Loki server.

Example loki.yaml file

```
auth_enabled: false

server:
  http_listen_port: 3100
  grpc_listen_port: 9096
  grpc_server_max_recv_msg_size: 8388608

ingester:
```

```
wal:
  enabled: true
  dir: /tmp/wal
lifecycle:
  address: 127.0.0.1
  ring:
    kvstore:
      store: inmemory
      replication_factor: 1
    final_sleep: 0s
  chunk_idle_period: 1h # Any chunk not receiving new logs in this time will be flushed
  chunk_target_size: 8388608
  max_chunk_age: 1h # All chunks will be flushed when they hit this age, default is 1h
  chunk_retain_period: 30s # Must be greater than index read cache TTL if using an index cache
  (Default index read cache TTL is 5m)
  max_transfer_retries: 0 # Chunk transfers disabled

schema_config:
  configs:
    - from: 2020-10-24
      store: boltdb-shipper
      object_store: filesystem
      schema: v11
      index:
        prefix: index_
        period: 24h

storage_config:
  boltdb_shipper:
    active_index_directory: /tmp/loki/boltdb-shipper-active
    cache_location: /tmp/loki/boltdb-shipper-cache
    cache_ttl: 24h # Can be increased for faster performance over longer query periods, uses
    more disk space
    shared_store: filesystem
  filesystem:
    directory: /tmp/loki/chunks

compactor:
  working_directory: /tmp/loki/boltdb-shipper-compactor
  shared_store: filesystem

limits_config:
  reject_old_samples: true
  reject_old_samples_max_age: 12h
  ingestion_rate_mb: 8
  ingestion_burst_size_mb: 16

chunk_store_config:
  max_look_back_period: 0s

table_manager:
  retention_deletes_enabled: false
  retention_period: 0s

ruler:
  storage:
```

```

type: local
local:
  directory: /tmp/loki/rules
rule_path: /tmp/loki/rules-temp
alertmanager_url: http://localhost:9093
ring:
  kvstore:
    store: inmemory
enable_api: true

```

Additional resources

- [Configuring Loki](#)

Additional resources

- [Log Record Fields.](#)
- [Configuring Loki server](#)

7.13. FORWARDING APPLICATION LOGS FROM SPECIFIC PROJECTS

You can use the Cluster Log Forwarder to send a copy of the application logs from specific projects to an external log aggregator. You can do this in addition to, or instead of, using the default Elasticsearch log store. You must also configure the external log aggregator to receive log data from OpenShift Container Platform.

To configure forwarding application logs from a project, you must create a **ClusterLogForwarder** custom resource (CR) with at least one input from a project, optional outputs for other log aggregators, and pipelines that use those inputs and outputs.

Prerequisites

- You must have a logging server that is configured to receive the logging data using the specified protocol or format.

Procedure

1. Create or edit a YAML file that defines the **ClusterLogForwarder** CR object:

```

apiVersion: logging.openshift.io/v1
kind: ClusterLogForwarder
metadata:
  name: instance 1
  namespace: openshift-logging 2
spec:
  outputs:
    - name: fluentd-server-secure 3
      type: fluentdForward 4
      url: 'tls://fluentdserver.security.example.com:24224' 5
      secret: 6
        name: fluentd-secret
    - name: fluentd-server-insecure
      type: fluentdForward

```

```

url: 'tcp://fluentdserver.home.example.com:24224'
inputs: 7
- name: my-app-logs
  application:
    namespaces:
      - my-project
pipelines:
- name: forward-to-fluentd-insecure 8
  inputRefs: 9
  - my-app-logs
  outputRefs: 10
  - fluentd-server-insecure
  parse: json 11
  labels:
    project: "my-project" 12
- name: forward-to-fluentd-secure 13
  inputRefs:
  - application
  - audit
  - infrastructure
  outputRefs:
  - fluentd-server-secure
  - default
  labels:
    clusterId: "C1234"

```

- 1 The name of the **ClusterLogForwarder** CR must be **instance**.
- 2 The namespace for the **ClusterLogForwarder** CR must be **openshift-logging**.
- 3 Specify a name for the output.
- 4 Specify the output type: **elasticsearch**, **fluentdForward**, **syslog**, or **kafka**.
- 5 Specify the URL and port of the external log aggregator as a valid absolute URL. If the cluster-wide proxy using the CIDR annotation is enabled, the output must be a server name or FQDN, not an IP address.
- 6 If using a **tls** prefix, you must specify the name of the secret required by the endpoint for TLS communication. The secret must exist in the **openshift-logging** project and have **tls.crt**, **tls.key**, and **ca-bundle.crt** keys that each point to the certificates they represent.
- 7 Configuration for an input to filter application logs from the specified projects.
- 8 Configuration for a pipeline to use the input to send project application logs to an external Fluentd instance.
- 9 The **my-app-logs** input.
- 10 The name of the output to use.
- 11 Optional: Specify whether to forward structured JSON log entries as JSON objects in the **structured** field. The log entry must contain valid structured JSON; otherwise, OpenShift Logging removes the **structured** field and instead sends the log entry to the default index, **app-00000x**.

- 12 Optional: String. One or more labels to add to the logs.
- 13 Configuration for a pipeline to send logs to other log aggregators.
 - Optional: Specify a name for the pipeline.
 - Specify which log types to forward by using the pipeline: **application**, **infrastructure**, or **audit**.
 - Specify the name of the output to use when forwarding logs with this pipeline.
 - Optional: Specify the **default** output to forward logs to the internal Elasticsearch instance.
 - Optional: String. One or more labels to add to the logs.

2. Create the CR object:

```
$ oc create -f <file-name>.yaml
```

7.14. FORWARDING APPLICATION LOGS FROM SPECIFIC PODS

As a cluster administrator, you can use Kubernetes pod labels to gather log data from specific pods and forward it to a log collector.

Suppose that you have an application composed of pods running alongside other pods in various namespaces. If those pods have labels that identify the application, you can gather and output their log data to a specific log collector.

To specify the pod labels, you use one or more **matchLabels** key-value pairs. If you specify multiple key-value pairs, the pods must match all of them to be selected.

Procedure

1. Create or edit a YAML file that defines the **ClusterLogForwarder** CR object. In the file, specify the pod labels using simple equality-based selectors under **inputs[].name.application.selector.matchLabels**, as shown in the following example.

Example ClusterLogForwarder CR YAML file

```
apiVersion: logging.openshift.io/v1
kind: ClusterLogForwarder
metadata:
  name: instance 1
  namespace: openshift-logging 2
spec:
  pipelines:
    - inputRefs: [ myAppLogData ] 3
      outputRefs: [ default ] 4
      parse: json 5
  inputs: 6
    - name: myAppLogData
      application:
```

```

selector:
  matchLabels: 7
    environment: production
    app: nginx
  namespaces: 8
    - app1
    - app2
outputs: 9
  - default
  ...

```

- 1 The name of the **ClusterLogForwarder** CR must be **instance**.
 - 2 The namespace for the **ClusterLogForwarder** CR must be **openshift-logging**.
 - 3 Specify one or more comma-separated values from **inputs[].name**.
 - 4 Specify one or more comma-separated values from **outputs[]**.
 - 5 Optional: Specify whether to forward structured JSON log entries as JSON objects in the **structured** field. The log entry must contain valid structured JSON; otherwise, OpenShift Logging removes the **structured** field and instead sends the log entry to the default index, **app-00000x**.
 - 6 Define a unique **inputs[].name** for each application that has a unique set of pod labels.
 - 7 Specify the key-value pairs of pod labels whose log data you want to gather. You must specify both a key and value, not just a key. To be selected, the pods must match all the key-value pairs.
 - 8 Optional: Specify one or more namespaces.
 - 9 Specify one or more outputs to forward your log data to. The optional **default** output shown here sends log data to the internal Elasticsearch instance.
2. Optional: To restrict the gathering of log data to specific namespaces, use **inputs[].name.application.namespaces**, as shown in the preceding example.
 3. Optional: You can send log data from additional applications that have different pod labels to the same pipeline.
 - a. For each unique combination of pod labels, create an additional **inputs[].name** section similar to the one shown.
 - b. Update the **selectors** to match the pod labels of this application.
 - c. Add the new **inputs[].name** value to **inputRefs**. For example:

```

- inputRefs: [ myAppLogData, myOtherAppLogData ]

```

4. Create the CR object:

```

$ oc create -f <file-name>.yaml

```

Additional resources

- For more information on **matchLabels** in Kubernetes, see [Resources that support set-based requirements](#).

Additional resources

- [Network policy audit logging](#)

7.15. TROUBLESHOOTING LOG FORWARDING

When you create a **ClusterLogForwarder** custom resource (CR), if the Red Hat OpenShift Logging Operator does not redeploy the Fluentd pods automatically, you can delete the Fluentd pods to force them to redeploy.

Prerequisites

- You have created a **ClusterLogForwarder** custom resource (CR) object.

Procedure

- Delete the Fluentd pods to force them to redeploy.

```
❯ $ oc delete pod --selector logging-infra=collector
```

CHAPTER 8. ENABLING JSON LOGGING

You can configure the Log Forwarding API to parse JSON strings into a structured object.

8.1. PARSING JSON LOGS

Logs including JSON logs are usually represented as a string inside the **message** field. That makes it hard for users to query specific fields inside a JSON document. OpenShift Logging's Log Forwarding API enables you to parse JSON logs into a structured object and forward them to either OpenShift Logging-managed Elasticsearch or any other third-party system supported by the Log Forwarding API.

To illustrate how this works, suppose that you have the following structured JSON log entry.

Example structured JSON log entry

```
{"level":"info","name":"fred","home":"bedrock"}
```

Normally, the **ClusterLogForwarder** custom resource (CR) forwards that log entry in the **message** field. The **message** field contains the JSON-quoted string equivalent of the JSON log entry, as shown in the following example.

Example message field

```
{"message": "{\"level\":\"info\",\"name\":\"fred\",\"home\":\"bedrock\"",  
  "more fields..."}
```

To enable parsing JSON log, you add **parse: json** to a pipeline in the **ClusterLogForwarder** CR, as shown in the following example.

Example snippet showing parse: json

```
pipelines:  
- inputRefs: [ application ]  
  outputRefs: myFluentd  
  parse: json
```

When you enable parsing JSON logs by using **parse: json**, the CR copies the JSON-structured log entry in a **structured** field, as shown in the following example. This does not modify the original **message** field.

Example structured output containing the structured JSON log entry

```
{"structured": { "level": "info", "name": "fred", "home": "bedrock" },  
  "more fields..."}
```



IMPORTANT

If the log entry does not contain valid structured JSON, the **structured** field will be absent.

To enable parsing JSON logs for specific logging platforms, see [Forwarding logs to third-party systems](#) .

8.2. CONFIGURING JSON LOG DATA FOR ELASTICSEARCH

If your JSON logs follow more than one schema, storing them in a single index might cause type conflicts and cardinality problems. To avoid that, you must configure the **ClusterLogForwarder** custom resource (CR) to group each schema into a single output definition. This way, each schema is forwarded to a separate index.



IMPORTANT

If you forward JSON logs to the default Elasticsearch instance managed by OpenShift Logging, it generates new indices based on your configuration. To avoid performance issues associated with having too many indices, consider keeping the number of possible schemas low by standardizing to common schemas.

Structure types

You can use the following structure types in the **ClusterLogForwarder** CR to construct index names for the Elasticsearch log store:

- **structuredTypeKey** (string, optional) is the name of a message field. The value of that field, if present, is used to construct the index name.
 - **kubernetes.labels.<key>** is the Kubernetes pod label whose value is used to construct the index name.
 - **openshift.labels.<key>** is the **pipeline.label.<key>** element in the **ClusterLogForwarder** CR whose value is used to construct the index name.
 - **kubernetes.container_name** uses the container name to construct the index name.
- **structuredTypeName**: (string, optional) If **structuredTypeKey** is not set or its key is not present, OpenShift Logging uses the value of **structuredTypeName** as the structured type. When you use both **structuredTypeKey** and **structuredTypeName** together, **structuredTypeName** provides a fallback index name if the key in **structuredTypeKey** is missing from the JSON log data.



NOTE

Although you can set the value of **structuredTypeKey** to any field shown in the "Log Record Fields" topic, the most useful fields are shown in the preceding list of structure types.

A structuredTypeKey: kubernetes.labels.<key> example

Suppose the following:

- Your cluster is running application pods that produce JSON logs in two different formats, "apache" and "google".
- The user labels these application pods with **logFormat=apache** and **logFormat=google**.
- You use the following snippet in your **ClusterLogForwarder** CR YAML file.

```
outputDefaults:
  elasticsearch:
```

```

    structuredTypeKey: kubernetes.labels.logFormat 1
    structuredTypeName: nologformat
  pipelines:
  - inputRefs: <application>
    outputRefs: default
    parse: json 2

```

- 1** Uses the value of the key-value pair that is formed by the Kubernetes **logFormat** label.
- 2** Enables parsing JSON logs.

In that case, the following structured log record goes to the **app-apache-write** index:

```

{
  "structured":{"name":"fred","home":"bedrock"},
  "kubernetes":{"labels":{"logFormat": "apache", ...}}
}

```

And the following structured log record goes to the **app-google-write** index:

```

{
  "structured":{"name":"wilma","home":"bedrock"},
  "kubernetes":{"labels":{"logFormat": "google", ...}}
}

```

A structuredTypeKey: openshift.labels.<key> example

Suppose that you use the following snippet in your **ClusterLogForwarder** CR YAML file.

```

outputDefaults:
  elasticsearch:
    structuredTypeKey: openshift.labels.myLabel 1
    structuredTypeName: nologformat
  pipelines:
  - name: application-logs
    inputRefs:
    - application
    - audit
    outputRefs:
    - elasticsearch-secure
    - default
    parse: json
    labels:
      myLabel: myValue 2

```

- 1** Uses the value of the key-value pair that is formed by the OpenShift **myLabel** label.
- 2** The **myLabel** element gives its string value, **myValue**, to the structured log record.

In that case, the following structured log record goes to the **app-myValue-write** index:

```

{
  "structured":{"name":"fred","home":"bedrock"},

```

```
"opensearch":{"labels":{"myLabel": "myValue", ...}}
}
```

Additional considerations

- The Elasticsearch *index* for structured records is formed by prepending "app-" to the structured type and appending "-write".
- Unstructured records are not sent to the structured index. They are indexed as usual in the application, infrastructure, or audit indices.
- If there is no non-empty structured type, forward an *unstructured* record with no **structured** field.

It is important not to overload Elasticsearch with too many indices. Only use distinct structured types for distinct log *formats*, **not** for each application or namespace. For example, most Apache applications use the same JSON log format and structured type, such as **LogApache**.

8.3. FORWARDING JSON LOGS TO THE ELASTICSEARCH LOG STORE

For an Elasticsearch log store, if your JSON log entries *follow different schemas*, configure the **ClusterLogForwarder** custom resource (CR) to group each JSON schema into a single output definition. This way, Elasticsearch uses a separate index for each schema.



IMPORTANT

Because forwarding different schemas to the same index can cause type conflicts and cardinality problems, you must perform this configuration before you forward data to the Elasticsearch store.


To avoid performance issues associated with having too many indices, consider keeping the number of possible schemas low by standardizing to common schemas.

Procedure

1. Add the following snippet to your **ClusterLogForwarder** CR YAML file.

```
outputDefaults:
  elasticsearch:
    structuredTypeKey: <log record field>
    structuredTypeName: <name>
  pipelines:
  - inputRefs:
  - application
  outputRefs: default
  parse: json
```

2. Optional: Use **structuredTypeKey** to specify one of the log record fields, as described in the preceding topic, [Configuring JSON log data for Elasticsearch](#). Otherwise, remove this line.
3. Optional: Use **structuredTypeName** to specify a **<name>**, as described in the preceding topic, [Configuring JSON log data for Elasticsearch](#). Otherwise, remove this line.

**IMPORTANT**

To parse JSON logs, you must set either **structuredTypeKey** or **structuredTypeName**, or both **structuredTypeKey** and **structuredTypeName**.

4. For **inputRefs**, specify which log types to forward by using that pipeline, such as **application**, **infrastructure**, or **audit**.
5. Add the **parse: json** element to pipelines.
6. Create the CR object:

```
$ oc create -f <file-name>.yaml
```

The Red Hat OpenShift Logging Operator redeploys the Fluentd pods. However, if they do not redeploy, delete the Fluentd pods to force them to redeploy.

```
$ oc delete pod --selector logging-infra=collector
```

Additional resources

- [Forwarding logs to third-party systems](#)

CHAPTER 9. COLLECTING AND STORING KUBERNETES EVENTS

The OpenShift Container Platform Event Router is a pod that watches Kubernetes events and logs them for collection by the logging subsystem. You must manually deploy the Event Router.

The Event Router collects events from all projects and writes them to **STDOUT**. The collector then forwards those events to the store defined in the **ClusterLogForwarder** custom resource (CR).



IMPORTANT

The Event Router adds additional load to Fluentd and can impact the number of other log messages that can be processed.

9.1. DEPLOYING AND CONFIGURING THE EVENT ROUTER

Use the following steps to deploy the Event Router into your cluster. You should always deploy the Event Router to the **openshift-logging** project to ensure it collects events from across the cluster.

The following Template object creates the service account, cluster role, and cluster role binding required for the Event Router. The template also configures and deploys the Event Router pod. You can use this template without making changes, or change the deployment object CPU and memory requests.

Prerequisites

- You need proper permissions to create service accounts and update cluster role bindings. For example, you can run the following template with a user that has the **cluster-admin** role.
- The logging subsystem for Red Hat OpenShift must be installed.

Procedure

1. Create a template for the Event Router:

```
kind: Template
apiVersion: template.openshift.io/v1
metadata:
  name: eventrouter-template
  annotations:
    description: "A pod forwarding kubernetes events to OpenShift Logging stack."
    tags: "events,EFK,logging,cluster-logging"
objects:
  - kind: ServiceAccount 1
    apiVersion: v1
    metadata:
      name: eventrouter
      namespace: ${NAMESPACE}
  - kind: ClusterRole 2
    apiVersion: rbac.authorization.k8s.io/v1
    metadata:
      name: event-reader
    rules:
      - apiGroups: [""]
```

```

resources: ["events"]
verbs: ["get", "watch", "list"]
- kind: ClusterRoleBinding 3
  apiVersion: rbac.authorization.k8s.io/v1
  metadata:
    name: event-reader-binding
  subjects:
  - kind: ServiceAccount
    name: eventrouter
    namespace: ${NAMESPACE}
  roleRef:
    kind: ClusterRole
    name: event-reader
- kind: ConfigMap 4
  apiVersion: v1
  metadata:
    name: eventrouter
    namespace: ${NAMESPACE}
  data:
    config.json: |-
      {
        "sink": "stdout"
      }
- kind: Deployment 5
  apiVersion: apps/v1
  metadata:
    name: eventrouter
    namespace: ${NAMESPACE}
  labels:
    component: "eventrouter"
    logging-infra: "eventrouter"
    provider: "openshift"
  spec:
    selector:
      matchLabels:
        component: "eventrouter"
        logging-infra: "eventrouter"
        provider: "openshift"
    replicas: 1
    template:
      metadata:
        labels:
          component: "eventrouter"
          logging-infra: "eventrouter"
          provider: "openshift"
        name: eventrouter
      spec:
        serviceAccount: eventrouter
        containers:
        - name: kube-eventrouter
          image: ${IMAGE}
          imagePullPolicy: IfNotPresent
          resources:
            requests:
              cpu: ${CPU}
              memory: ${MEMORY}

```

```

    volumeMounts:
      - name: config-volume
        mountPath: /etc/eventrouter
    volumes:
      - name: config-volume
        configMap:
          name: eventrouter
  parameters:
    - name: IMAGE 6
      displayName: Image
      value: "registry.redhat.io/openshift-logging/eventrouter-rhel8:v0.4"
    - name: CPU 7
      displayName: CPU
      value: "100m"
    - name: MEMORY 8
      displayName: Memory
      value: "128Mi"
    - name: NAMESPACE
      displayName: Namespace
      value: "openshift-logging" 9

```

- 1** Creates a Service Account in the **openshift-logging** project for the Event Router.
- 2** Creates a ClusterRole to monitor for events in the cluster.
- 3** Creates a ClusterRoleBinding to bind the ClusterRole to the service account.
- 4** Creates a config map in the **openshift-logging** project to generate the required **config.json** file.
- 5** Creates a deployment in the **openshift-logging** project to generate and configure the Event Router pod.
- 6** Specifies the image, identified by a tag such as **v0.4**.
- 7** Specifies the minimum amount of CPU to allocate to the Event Router pod. Defaults to **100m**.
- 8** Specifies the minimum amount of memory to allocate to the Event Router pod. Defaults to **128Mi**.
- 9** Specifies the **openshift-logging** project to install objects in.

2. Use the following command to process and apply the template:

```
$ oc process -f <templatefile> | oc apply -n openshift-logging -f -
```

For example:

```
$ oc process -f eventrouter.yaml | oc apply -n openshift-logging -f -
```

Example output

```
serviceaccount/eventrouter created
```

```
clusterrole.authorization.openshift.io/event-reader created
clusterrolebinding.authorization.openshift.io/event-reader-binding created
configmap/eventrouter created
deployment.apps/eventrouter created
```

3. Validate that the Event Router installed in the **openshift-logging** project:

- a. View the new Event Router pod:

```
$ oc get pods --selector component=eventrouter -o name -n openshift-logging
```

Example output

```
pod/cluster-logging-eventrouter-d649f97c8-qvv8r
```

- b. View the events collected by the Event Router:

```
$ oc logs <cluster_logging_eventrouter_pod> -n openshift-logging
```

For example:

```
$ oc logs cluster-logging-eventrouter-d649f97c8-qvv8r -n openshift-logging
```

Example output

```
{"verb":"ADDED","event":{"metadata":{"name":"openshift-service-catalog-controller-remover.1632d931e88fcd8f","namespace":"openshift-service-catalog-removed","selfLink":"/api/v1/namespaces/openshift-service-catalog-removed/events/openshift-service-catalog-controller-manager-remover.1632d931e88fcd8f","uid":"787d7b26-3d2f-4017-b0b0-420db4ae62c0","resourceVersion":"21399","creationTimestamp":"2020-09-08T15:40:26Z"},"involvedObject":{"kind":"Job","namespace":"openshift-service-catalog-removed","name":"openshift-service-catalog-controller-manager-remover","uid":"fac9f479-4ad5-4a57-8adc-cb25d3d9cf8f","apiVersion":"batch/v1","resourceVersion":"21280"},"reason":"Completed","message":"Job completed","source":{"component":"job-controller"},"firstTimestamp":"2020-09-08T15:40:26Z","lastTimestamp":"2020-09-08T15:40:26Z","count":1,"type":"Normal"}}
```

You can also use Kibana to view events by creating an index pattern using the Elasticsearch **infra** index.

CHAPTER 10. UPDATING OPENSIFT LOGGING

10.1. SUPPORTED VERSIONS

For version compatibility and support information, see [Red Hat OpenShift Container Platform Life Cycle Policy](#)


To upgrade from cluster logging in OpenShift Container Platform version 4.6 and earlier to OpenShift Logging 5.x, you update the OpenShift Container Platform cluster to version 4.7 or 4.8. Then, you update the following operators:

- From Elasticsearch Operator 4.x to OpenShift Elasticsearch Operator 5.x
- From Cluster Logging Operator 4.x to Red Hat OpenShift Logging Operator 5.x

To upgrade from a previous version of OpenShift Logging to the current version, you update OpenShift Elasticsearch Operator and Red Hat OpenShift Logging Operator to their current versions.

10.2. UPDATING LOGGING TO THE CURRENT VERSION

To update Logging to the current version, you change the subscriptions for the OpenShift Elasticsearch Operator and Red Hat OpenShift Logging Operator.



IMPORTANT

You must update the OpenShift Elasticsearch Operator *before* you update the Red Hat OpenShift Logging Operator. You must also update *both* Operators to the same version.

If you update the Operators in the wrong order, Kibana does not update and the Kibana custom resource (CR) is not created. To work around this problem, you delete the Red Hat OpenShift Logging Operator pod. When the Red Hat OpenShift Logging Operator pod redeploys, it creates the Kibana CR and Kibana becomes available again.

Prerequisites

- The OpenShift Container Platform version is 4.7 or later.
- The Logging status is healthy:
 - All pods are **ready**.
 - The Elasticsearch cluster is healthy.
- Your [Elasticsearch and Kibana data is backed up](#).

Procedure

1. Update the OpenShift Elasticsearch Operator:
 - a. From the web console, click **Operators** → **Installed Operators**.
 - b. Select the **openshift-operators-redhat** project.
 - c. Click the **OpenShift Elasticsearch Operator**.

- d. Click **Subscription** → **Channel**.
 - e. In the **Change Subscription Update Channel** window, select **stable-5.x** and click **Save**.
 - f. Wait for a few seconds, then click **Operators** → **Installed Operators**.
 - g. Verify that the OpenShift Elasticsearch Operator version is 5.x.x.
 - h. Wait for the **Status** field to report **Succeeded**.
2. Update the Red Hat OpenShift Logging Operator:
 - a. From the web console, click **Operators** → **Installed Operators**.
 - b. Select the **openshift-logging** project.
 - c. Click the **Red Hat OpenShift Logging Operator**.
 - d. Click **Subscription** → **Channel**.
 - e. In the **Change Subscription Update Channel** window, select **stable-5.x** and click **Save**.
 - f. Wait for a few seconds, then click **Operators** → **Installed Operators**.
 - g. Verify that the Red Hat OpenShift Logging Operator version is 5.y.z
 - h. Wait for the **Status** field to report **Succeeded**.
 3. Check the logging components:

- a. Ensure that all Elasticsearch pods are in the **Ready** status:

```
$ oc get pod -n openshift-logging --selector component=elasticsearch
```

Example output

```
NAME                                READY STATUS RESTARTS AGE
elasticsearch-cdm-1pbrl44l-1-55b7546f4c-mshhk 2/2 Running 0      31m
elasticsearch-cdm-1pbrl44l-2-5c6d87589f-gx5hk 2/2 Running 0      30m
elasticsearch-cdm-1pbrl44l-3-88df5d47-m45jc 2/2 Running 0      29m
```

- b. Ensure that the Elasticsearch cluster is healthy:

```
$ oc exec -n openshift-logging -c elasticsearch elasticsearch-cdm-1pbrl44l-1-55b7546f4c-mshhk -- health
```

```
{
  "cluster_name" : "elasticsearch",
  "status" : "green",
}
```

- c. Ensure that the Elasticsearch cron jobs are created:

```
$ oc project openshift-logging
```

```
$ oc get cronjob
```

NAME	SCHEDULE	SUSPEND	ACTIVE	LAST SCHEDULE	AGE
elasticsearch-im-app	*/15 * * * *	False	0	<none>	56s
elasticsearch-im-audit	*/15 * * * *	False	0	<none>	56s
elasticsearch-im-infra	*/15 * * * *	False	0	<none>	56s

- d. Verify that the log store is updated to 5.x and the indices are **green**:

```
$ oc exec -c elasticsearch <any_es_pod_in_the_cluster> -- indices
```

- e. Verify that the output includes the **app-00000x**, **infra-00000x**, **audit-00000x**, **.security** indices.

Example 10.1. Sample output with indices in a green status

```
Tue Jun 30 14:30:54 UTC 2020
health status index                                uuid                                pri rep
docs.count docs.deleted store.size pri.store.size
green open  infra-000008
bnBvUFEXTWi92z3zWAZieQ 3 1    222195    0    289    144
green open  infra-000004
rtDSzoqsSl6saisSK7Au1Q 3 1    226717    0    297    148
green open  infra-000012
RSf_kUwDSR2xEuKRZMPqZQ 3 1    227623    0    295    147
green open  .kibana_7
1SJdCqIZTPWIIAaOUd78yg 1 1     4        0    0      0
green open  infra-000010
iXwL3bnqTuGEABbUDa6OVw 3 1    248368    0    317    158
green open  infra-000009
YN9EsULWSNaxWeeNvOs0RA 3 1    258799    0    337    168
green open  infra-000014
YP0U6R7FQ_GVQVQZ6Yh9lg 3 1    223788    0    292    146
green open  infra-000015
JRBbAbEmSMqK5X40df9HbQ 3 1    224371    0    291    145
green open  .orphaned.2020.06.30
n_xQC2dWQzConkvQqei3YA 3 1     9        0    0      0
green open  infra-000007
llkkAVSzSOMosWTSAJM_hg 3 1    228584    0    296    148
green open  infra-000005
d9BoGQdiQASsS3BBFm2iIRA 3 1    227987    0    297    148
green open  infra-000003
goREK1QUKIQPAIVkWVaQ 3 1    226719    0    295    147
green open  .security
zeT65uOuRTKZMjg_bbUc1g 1 1     5        0    0      0
green open  .kibana-377444158_kubeadmin          wwMhDwJkR-
mRZQO84K0gUQ 3 1     1        0    0      0
green open  infra-000006
KBSXGQKiO7hdapDE23g 3 1    226676    0    295    147
green open  infra-000001
bSxSWR5xYZB6IVg 3 1    341800    0    443    220
green open  .kibana-6
RVp77TemSSemGJcsSUmuf3A 1 1     4        0    0      0
green open  infra-000011
J7XWBauWSTe0jnzX02fU6A 3 1    226100    0    293    146
```

```

green open app-000001
axSAFfONQDmKwatkjPXdtw 3 1 103186 0 126 57
green open infra-000016
m9c1iRLiStWSF1GopaRyCg 3 1 13685 0 19 9
green open infra-000002
ewmbYg 3 1 228994 0 296 148 Hz6WvINtTvKcQzw-
green open infra-000013
jraYtanyIGw 3 1 228166 0 298 148 KR9mMFUpQI-
green open audit-000001
eERqLdLmQOiQDFES1LBATQ 3 1 0 0 0 0

```

- f. Verify that the log collector is updated:

```
$ oc get ds collector -o json | grep collector
```

- g. Verify that the output includes a **collectort** container:

```
"containerName": "collector"
```

- h. Verify that the log visualizer is updated to 5.x using the Kibana CRD:

```
$ oc get kibana kibana -o json
```

- i. Verify that the output includes a Kibana pod with the **ready** status:

Example 10.2. Sample output with a ready Kibana pod

```

[
  {
    "clusterCondition": {
      "kibana-5fdd766ffd-nb2jj": [
        {
          "lastTransitionTime": "2020-06-30T14:11:07Z",
          "reason": "ContainerCreating",
          "status": "True",
          "type": ""
        },
        {
          "lastTransitionTime": "2020-06-30T14:11:07Z",
          "reason": "ContainerCreating",
          "status": "True",
          "type": ""
        }
      ]
    },
    "deployment": "kibana",
    "pods": {
      "failed": [],
      "notReady": [],
      "ready": []
    },
    "replicaSets": [
      "kibana-5fdd766ffd"
    ]
  }
]

```

```
    | | "replicas": 1  
    | | }  
    | | ]
```

CHAPTER 11. VIEWING CLUSTER DASHBOARDS

The **Logging/Elasticsearch Nodes** and **OpenShift Logging** dashboards in the OpenShift Container Platform web console show in-depth details about your Elasticsearch instance and the individual Elasticsearch nodes that you can use to prevent and diagnose problems.

The **OpenShift Logging** dashboard contains charts that show details about your Elasticsearch instance at a cluster level, including cluster resources, garbage collection, shards in the cluster, and Fluentd statistics.

The **Logging/Elasticsearch Nodes** dashboard contains charts that show details about your Elasticsearch instance, many at node level, including details on indexing, shards, resources, and so forth.



NOTE

For more detailed data, click the **Grafana UI** link in a dashboard to launch the Grafana dashboard. Grafana is shipped with [OpenShift cluster monitoring](#).

11.1. ACCESSING THE ELASTICSEARCH AND OPENSIFT LOGGING DASHBOARDS

You can view the **Logging/Elasticsearch Nodes** and **OpenShift Logging** dashboards in the OpenShift Container Platform web console.

Procedure

To launch the dashboards:

1. In the OpenShift Container Platform web console, click **Observe** → **Dashboards**.
2. On the **Dashboards** page, select **Logging/Elasticsearch Nodes** or **OpenShift Logging** from the **Dashboard** menu.
For the **Logging/Elasticsearch Nodes** dashboard, you can select the Elasticsearch node you want to view and set the data resolution.

The appropriate dashboard is displayed, showing multiple charts of data.

3. Optional: Select a different time range to display or refresh rate for the data from the **Time Range** and **Refresh Interval** menus.



NOTE

For more detailed data, click the **Grafana UI** link to launch the Grafana dashboard.

For information on the dashboard charts, see [About the OpenShift Logging dashboard](#) and [About the Logging/Elasticsearch Nodes dashboard](#).

11.2. ABOUT THE OPENSIFT LOGGING DASHBOARD

The **OpenShift Logging** dashboard contains charts that show details about your Elasticsearch instance at a cluster-level that you can use to diagnose and anticipate problems.

Table 11.1. OpenShift Logging charts

Metric	Description
Elastic Cluster Status	<p>The current Elasticsearch status:</p> <ul style="list-style-type: none"> ● ONLINE - Indicates that the Elasticsearch instance is online. ● OFFLINE - Indicates that the Elasticsearch instance is offline.
Elastic Nodes	The total number of Elasticsearch nodes in the Elasticsearch instance.
Elastic Shards	The total number of Elasticsearch shards in the Elasticsearch instance.
Elastic Documents	The total number of Elasticsearch documents in the Elasticsearch instance.
Total Index Size on Disk	The total disk space that is being used for the Elasticsearch indices.
Elastic Pending Tasks	The total number of Elasticsearch changes that have not been completed, such as index creation, index mapping, shard allocation, or shard failure.
Elastic JVM GC time	The amount of time that the JVM spent executing Elasticsearch garbage collection operations in the cluster.
Elastic JVM GC Rate	The total number of times that JVM executed garbage activities per second.
Elastic Query/Fetch Latency Sum	<ul style="list-style-type: none"> ● Query latency: The average time each Elasticsearch search query takes to execute. ● Fetch latency: The average time each Elasticsearch search query spends fetching data. <p>Fetch latency typically takes less time than query latency. If fetch latency is consistently increasing, it might indicate slow disks, data enrichment, or large requests with too many results.</p>
Elastic Query Rate	The total queries executed against the Elasticsearch instance per second for each Elasticsearch node.
CPU	The amount of CPU used by Elasticsearch, Fluentd, and Kibana, shown for each component.

Metric	Description
Elastic JVM Heap Used	The amount of JVM memory used. In a healthy cluster, the graph shows regular drops as memory is freed by JVM garbage collection.
Elasticsearch Disk Usage	The total disk space used by the Elasticsearch instance for each Elasticsearch node.
File Descriptors In Use	The total number of file descriptors used by Elasticsearch, Fluentd, and Kibana.
FluentD emit count	The total number of Fluentd messages per second for the Fluentd default output, and the retry count for the default output.
FluentD Buffer Availability	The percent of the Fluentd buffer that is available for chunks. A full buffer might indicate that Fluentd is not able to process the number of logs received.
Elastic rx bytes	The total number of bytes that Elasticsearch has received from FluentD, the Elasticsearch nodes, and other sources.
Elastic Index Failure Rate	The total number of times per second that an Elasticsearch index fails. A high rate might indicate an issue with indexing.
FluentD Output Error Rate	The total number of times per second that FluentD is not able to output logs.

11.3. CHARTS ON THE LOGGING/ELASTICSEARCH NODES DASHBOARD

The **Logging/Elasticsearch Nodes** dashboard contains charts that show details about your Elasticsearch instance, many at node-level, for further diagnostics.

Elasticsearch status

The **Logging/Elasticsearch Nodes** dashboard contains the following charts about the status of your Elasticsearch instance.

Table 11.2. Elasticsearch status fields

Metric	Description
--------	-------------

Metric	Description
Cluster status	<p>The cluster health status during the selected time period, using the Elasticsearch green, yellow, and red statuses:</p> <ul style="list-style-type: none"> ● 0 - Indicates that the Elasticsearch instance is in green status, which means that all shards are allocated. ● 1 - Indicates that the Elasticsearch instance is in yellow status, which means that replica shards for at least one shard are not allocated. ● 2 - Indicates that the Elasticsearch instance is in red status, which means that at least one primary shard and its replicas are not allocated.
Cluster nodes	The total number of Elasticsearch nodes in the cluster.
Cluster data nodes	The number of Elasticsearch data nodes in the cluster.
Cluster pending tasks	The number of cluster state changes that are not finished and are waiting in a cluster queue, for example, index creation, index deletion, or shard allocation. A growing trend indicates that the cluster is not able to keep up with changes.

Elasticsearch cluster index shard status

Each Elasticsearch index is a logical group of one or more shards, which are basic units of persisted data. There are two types of index shards: primary shards, and replica shards. When a document is indexed into an index, it is stored in one of its primary shards and copied into every replica of that shard. The number of primary shards is specified when the index is created, and the number cannot change during index lifetime. You can change the number of replica shards at any time.

The index shard can be in several states depending on its lifecycle phase or events occurring in the cluster. When the shard is able to perform search and indexing requests, the shard is active. If the shard cannot perform these requests, the shard is non-active. A shard might be non-active if the shard is initializing, reallocating, unassigned, and so forth.

Index shards consist of a number of smaller internal blocks, called index segments, which are physical representations of the data. An index segment is a relatively small, immutable Lucene index that is created when Lucene commits newly-indexed data. Lucene, a search library used by Elasticsearch, merges index segments into larger segments in the background to keep the total number of segments low. If the process of merging segments is slower than the speed at which new segments are created, it could indicate a problem.

When Lucene performs data operations, such as a search operation, Lucene performs the operation against the index segments in the relevant index. For that purpose, each segment contains specific data structures that are loaded in the memory and mapped. Index mapping can have a significant impact on the memory used by segment data structures.

The **Logging/Elasticsearch Nodes** dashboard contains the following charts about the Elasticsearch index shards.

Table 11.3. Elasticsearch cluster shard status charts

Metric	Description
Cluster active shards	The number of active primary shards and the total number of shards, including replicas, in the cluster. If the number of shards grows higher, the cluster performance can start degrading.
Cluster initializing shards	The number of non-active shards in the cluster. A non-active shard is one that is initializing, being reallocated to a different node, or is unassigned. A cluster typically has non-active shards for short periods. A growing number of non-active shards over longer periods could indicate a problem.
Cluster relocating shards	The number of shards that Elasticsearch is relocating to a new node. Elasticsearch relocates nodes for multiple reasons, such as high memory use on a node or after a new node is added to the cluster.
Cluster unassigned shards	The number of unassigned shards. Elasticsearch shards might be unassigned for reasons such as a new index being added or the failure of a node.

Elasticsearch node metrics

Each Elasticsearch node has a finite amount of resources that can be used to process tasks. When all the resources are being used and Elasticsearch attempts to perform a new task, Elasticsearch put the tasks into a queue until some resources become available.

The **Logging/Elasticsearch Nodes** dashboard contains the following charts about resource usage for a selected node and the number of tasks waiting in the Elasticsearch queue.

Table 11.4. Elasticsearch node metric charts

Metric	Description
ThreadPool tasks	The number of waiting tasks in individual queues, shown by task type. A long-term accumulation of tasks in any queue could indicate node resource shortages or some other problem.
CPU usage	The amount of CPU being used by the selected Elasticsearch node as a percentage of the total CPU allocated to the host container.
Memory usage	The amount of memory being used by the selected Elasticsearch node.

Metric	Description
Disk usage	The total disk space being used for index data and metadata on the selected Elasticsearch node.
Documents indexing rate	The rate that documents are indexed on the selected Elasticsearch node.
Indexing latency	The time taken to index the documents on the selected Elasticsearch node. Indexing latency can be affected by many factors, such as JVM Heap memory and overall load. A growing latency indicates a resource capacity shortage in the instance.
Search rate	The number of search requests run on the selected Elasticsearch node.
Search latency	The time taken to complete search requests on the selected Elasticsearch node. Search latency can be affected by many factors. A growing latency indicates a resource capacity shortage in the instance.
Documents count (with replicas)	The number of Elasticsearch documents stored on the selected Elasticsearch node, including documents stored in both the primary shards and replica shards that are allocated on the node.
Documents deleting rate	The number of Elasticsearch documents being deleted from any of the index shards that are allocated to the selected Elasticsearch node.
Documents merging rate	The number of Elasticsearch documents being merged in any of index shards that are allocated to the selected Elasticsearch node.

Elasticsearch node fielddata

Fielddata is an Elasticsearch data structure that holds lists of terms in an index and is kept in the JVM Heap. Because fielddata building is an expensive operation, Elasticsearch caches the fielddata structures. Elasticsearch can evict a fielddata cache when the underlying index segment is deleted or merged, or if there is not enough JVM HEAP memory for all the fielddata caches.

The **Logging/Elasticsearch Nodes** dashboard contains the following charts about Elasticsearch fielddata.

Table 11.5. Elasticsearch node fielddata charts

Metric	Description
--------	-------------

Metric	Description
Fielddata memory size	The amount of JVM Heap used for the fielddata cache on the selected Elasticsearch node.
Fielddata evictions	The number of fielddata structures that were deleted from the selected Elasticsearch node.

Elasticsearch node query cache

If the data stored in the index does not change, search query results are cached in a node-level query cache for reuse by Elasticsearch.

The **Logging/Elasticsearch Nodes** dashboard contains the following charts about the Elasticsearch node query cache.

Table 11.6. Elasticsearch node query charts

Metric	Description
Query cache size	The total amount of memory used for the query cache for all the shards allocated to the selected Elasticsearch node.
Query cache evictions	The number of query cache evictions on the selected Elasticsearch node.
Query cache hits	The number of query cache hits on the selected Elasticsearch node.
Query cache misses	The number of query cache misses on the selected Elasticsearch node.

Elasticsearch index throttling

When indexing documents, Elasticsearch stores the documents in index segments, which are physical representations of the data. At the same time, Elasticsearch periodically merges smaller segments into a larger segment as a way to optimize resource use. If the indexing is faster than the ability to merge segments, the merge process does not complete quickly enough, which can lead to issues with searches and performance. To prevent this situation, Elasticsearch throttles indexing, typically by reducing the number of threads allocated to indexing down to a single thread.

The **Logging/Elasticsearch Nodes** dashboard contains the following charts about Elasticsearch index throttling.

Table 11.7. Index throttling charts

Metric	Description
Indexing throttling	The amount of time that Elasticsearch has been throttling the indexing operations on the selected Elasticsearch node.

Metric	Description
Merging throttling	The amount of time that Elasticsearch has been throttling the segment merge operations on the selected Elasticsearch node.

Node JVM Heap statistics

The **Logging/Elasticsearch Nodes** dashboard contains the following charts about JVM Heap operations.

Table 11.8. JVM Heap statistic charts

Metric	Description
Heap used	The amount of the total allocated JVM Heap space that is used on the selected Elasticsearch node.
GC count	The number of garbage collection operations that have been run on the selected Elasticsearch node, by old and young garbage collection.
GC time	The amount of time that the JVM spent running garbage collection operations on the selected Elasticsearch node, by old and young garbage collection.

CHAPTER 12. TROUBLESHOOTING LOGGING

12.1. VIEWING OPENSIFT LOGGING STATUS

You can view the status of the Red Hat OpenShift Logging Operator and for a number of logging subsystem components.

12.1.1. Viewing the status of the Red Hat OpenShift Logging Operator

You can view the status of your Red Hat OpenShift Logging Operator.

Prerequisites

- The Red Hat OpenShift Logging and Elasticsearch Operators must be installed.

Procedure

1. Change to the **openshift-logging** project.

```
$ oc project openshift-logging
```

2. To view the OpenShift Logging status:
 - a. Get the OpenShift Logging status:

```
$ oc get clusterlogging instance -o yaml
```

Example output

```
apiVersion: logging.openshift.io/v1
kind: ClusterLogging

....

status: 1
collection:
logs:
  fluentdStatus:
    daemonSet: fluentd 2
    nodes:
      fluentd-2rhqp: ip-10-0-169-13.ec2.internal
      fluentd-6fgjh: ip-10-0-165-244.ec2.internal
      fluentd-6l2ff: ip-10-0-128-218.ec2.internal
      fluentd-54nx5: ip-10-0-139-30.ec2.internal
      fluentd-flpnn: ip-10-0-147-228.ec2.internal
      fluentd-n2frh: ip-10-0-157-45.ec2.internal
    pods:
      failed: []
      notReady: []
      ready:
        - fluentd-2rhqp
        - fluentd-54nx5
        - fluentd-6fgjh
```

```

- fluentd-6l2ff
- fluentd-flpnn
- fluentd-n2frh
logstore: 3
elasticsearchStatus:
- ShardAllocationEnabled: all
cluster:
  activePrimaryShards: 5
  activeShards: 5
  initializingShards: 0
  numDataNodes: 1
  numNodes: 1
  pendingTasks: 0
  relocatingShards: 0
  status: green
  unassignedShards: 0
clusterName: elasticsearch
nodeConditions:
  elasticsearch-cdm-mkkdys93-1:
nodeCount: 1
pods:
  client:
    failed:
    notReady:
    ready:
    - elasticsearch-cdm-mkkdys93-1-7f7c6-mjm7c
  data:
    failed:
    notReady:
    ready:
    - elasticsearch-cdm-mkkdys93-1-7f7c6-mjm7c
  master:
    failed:
    notReady:
    ready:
    - elasticsearch-cdm-mkkdys93-1-7f7c6-mjm7c
visualization: 4
kibanaStatus:
- deployment: kibana
pods:
  failed: []
  notReady: []
  ready:
  - kibana-7fb4fd4cc9-f2nls
replicaSets:
- kibana-7fb4fd4cc9
replicas: 1

```

- 1 In the output, the cluster status fields appear in the **status** stanza.
- 2 Information on the Fluentd pods.
- 3 Information on the Elasticsearch pods, including Elasticsearch cluster health, **green**, **yellow**, or **red**.
- 4 Information on the Kibana pods.

12.1.1.1. Example condition messages

The following are examples of some condition messages from the **Status.Nodes** section of the OpenShift Logging instance.

A status message similar to the following indicates a node has exceeded the configured low watermark and no shard will be allocated to this node:

Example output

```
nodes:
- conditions:
- lastTransitionTime: 2019-03-15T15:57:22Z
  message: Disk storage usage for node is 27.5gb (36.74%). Shards will be not
    be allocated on this node.
  reason: Disk Watermark Low
  status: "True"
  type: NodeStorage
  deploymentName: example-elasticsearch-clientdatamaster-0-1
  upgradeStatus: {}
```

A status message similar to the following indicates a node has exceeded the configured high watermark and shards will be relocated to other nodes:

Example output

```
nodes:
- conditions:
- lastTransitionTime: 2019-03-15T16:04:45Z
  message: Disk storage usage for node is 27.5gb (36.74%). Shards will be relocated
    from this node.
  reason: Disk Watermark High
  status: "True"
  type: NodeStorage
  deploymentName: cluster-logging-operator
  upgradeStatus: {}
```

A status message similar to the following indicates the Elasticsearch node selector in the CR does not match any nodes in the cluster:

Example output

```
Elasticsearch Status:
Shard Allocation Enabled: shard allocation unknown
Cluster:
  Active Primary Shards: 0
  Active Shards:        0
  Initializing Shards:  0
  Num Data Nodes:       0
  Num Nodes:            0
  Pending Tasks:        0
  Relocating Shards:    0
  Status:                cluster health unknown
  Unassigned Shards:    0
Cluster Name:           elasticsearch
```



```

Node Conditions:
  elasticsearch-cdm-mkkdys93-1:
    Last Transition Time: 2019-06-26T03:37:32Z
    Message:          0/5 nodes are available: 5 node(s) didn't match node selector.
    Reason:           Unschedulable
    Status:           True
    Type:             Unschedulable
  elasticsearch-cdm-mkkdys93-2:
Node Count: 2
Pods:
  Client:
    Failed:
    Not Ready:
      elasticsearch-cdm-mkkdys93-1-75dd69dccd-f7f49
      elasticsearch-cdm-mkkdys93-2-67c64f5f4c-n58vl
    Ready:
  Data:
    Failed:
    Not Ready:
      elasticsearch-cdm-mkkdys93-1-75dd69dccd-f7f49
      elasticsearch-cdm-mkkdys93-2-67c64f5f4c-n58vl
    Ready:
  Master:
    Failed:
    Not Ready:
      elasticsearch-cdm-mkkdys93-1-75dd69dccd-f7f49
      elasticsearch-cdm-mkkdys93-2-67c64f5f4c-n58vl
    Ready:

```

A status message similar to the following indicates that the requested PVC could not bind to PV:

Example output

```

Node Conditions:
  elasticsearch-cdm-mkkdys93-1:
    Last Transition Time: 2019-06-26T03:37:32Z
    Message:          pod has unbound immediate PersistentVolumeClaims (repeated 5 times)
    Reason:           Unschedulable
    Status:           True
    Type:             Unschedulable

```

A status message similar to the following indicates that the Fluentd pods cannot be scheduled because the node selector did not match any nodes:

Example output

```

Status:
Collection:
Logs:
Fluentd Status:
  Daemon Set: fluentd
Nodes:
Pods:

```

```
Failed:
Not Ready:
Ready:
```

12.1.2. Viewing the status of logging subsystem components

You can view the status for a number of logging subsystem components.

Prerequisites

- The Red Hat OpenShift Logging and Elasticsearch Operators must be installed.

Procedure

1. Change to the **openshift-logging** project.

```
$ oc project openshift-logging
```

2. View the status of the logging subsystem for Red Hat OpenShift environment:

```
$ oc describe deployment cluster-logging-operator
```

Example output

```
Name:          cluster-logging-operator
...

Conditions:
  Type          Status Reason
  ----          -
  Available     True   MinimumReplicasAvailable
  Progressing  True   NewReplicaSetAvailable
...

Events:
  Type    Reason          Age    From          Message
  ----    -
  Normal  ScalingReplicaSet 62m   deployment-controller  Scaled up replica set cluster-logging-operator-574b8987df to 1----
```

3. View the status of the logging subsystem replica set:

- a. Get the name of a replica set:

Example output

```
$ oc get replicaset
```

Example output

```
NAME                                DESIRED CURRENT READY AGE
```

```

cluster-logging-operator-574b8987df    1    1    1    159m
elasticsearch-cdm-uhr537yu-1-6869694fb 1    1    1    157m
elasticsearch-cdm-uhr537yu-2-857b6d676f 1    1    1    156m
elasticsearch-cdm-uhr537yu-3-5b6fdd8cfd 1    1    1    155m
kibana-5bd5544f87                      1    1    1    157m

```

- b. Get the status of the replica set:

```
$ oc describe replicaset cluster-logging-operator-574b8987df
```

Example output

```

Name:          cluster-logging-operator-574b8987df
...

Replicas:      1 current / 1 desired
Pods Status:   1 Running / 0 Waiting / 0 Succeeded / 0 Failed
...

Events:
  Type      Reason      Age From          Message
  ----      -
Normal SuccessfulCreate 66m replicaset-controller Created pod: cluster-logging-operator-574b8987df-qjhqv----

```

12.2. VIEWING THE STATUS OF THE ELASTICSEARCH LOG STORE

You can view the status of the OpenShift Elasticsearch Operator and for a number of Elasticsearch components.

12.2.1. Viewing the status of the log store

You can view the status of your log store.

Prerequisites

- The Red Hat OpenShift Logging and Elasticsearch Operators must be installed.

Procedure

1. Change to the **openshift-logging** project.

```
$ oc project openshift-logging
```

2. To view the status:

- a. Get the name of the log store instance:

```
$ oc get Elasticsearch
```

Example output

```
NAME      AGE
elasticsearch 5h9m
```

- b. Get the log store status:

```
$ oc get Elasticsearch <Elasticsearch-instance> -o yaml
```

For example:

```
$ oc get Elasticsearch elasticsearch -n openshift-logging -o yaml
```

The output includes information similar to the following:

Example output

```
status: 1
cluster: 2
  activePrimaryShards: 30
  activeShards: 60
  initializingShards: 0
  numDataNodes: 3
  numNodes: 3
  pendingTasks: 0
  relocatingShards: 0
  status: green
  unassignedShards: 0
clusterHealth: ""
conditions: [] 3
nodes: 4
- deploymentName: elasticsearch-cdm-zjf34ved-1
  upgradeStatus: {}
- deploymentName: elasticsearch-cdm-zjf34ved-2
  upgradeStatus: {}
- deploymentName: elasticsearch-cdm-zjf34ved-3
  upgradeStatus: {}
pods: 5
client:
  failed: []
  notReady: []
  ready:
  - elasticsearch-cdm-zjf34ved-1-6d7fbf844f-sn422
  - elasticsearch-cdm-zjf34ved-2-dfbd988bc-qkzjz
  - elasticsearch-cdm-zjf34ved-3-c8f566f7c-t7zkt
data:
  failed: []
  notReady: []
  ready:
  - elasticsearch-cdm-zjf34ved-1-6d7fbf844f-sn422
  - elasticsearch-cdm-zjf34ved-2-dfbd988bc-qkzjz
  - elasticsearch-cdm-zjf34ved-3-c8f566f7c-t7zkt
master:
  failed: []
  notReady: []
  ready:
```

```

- elasticsearch-cdm-zjf34ved-1-6d7fbf844f-sn422
- elasticsearch-cdm-zjf34ved-2-dfbd988bc-qkzjz
- elasticsearch-cdm-zjf34ved-3-c8f566f7c-t7zkt
shardAllocationEnabled: all

```

- 1 In the output, the cluster status fields appear in the **status** stanza.
- 2 The status of the log store:
 - The number of active primary shards.
 - The number of active shards.
 - The number of shards that are initializing.
 - The number of log store data nodes.
 - The total number of log store nodes.
 - The number of pending tasks.
 - The log store status: **green, red, yellow**.
 - The number of unassigned shards.
- 3 Any status conditions, if present. The log store status indicates the reasons from the scheduler if a pod could not be placed. Any events related to the following conditions are shown:
 - Container Waiting for both the log store and proxy containers.
 - Container Terminated for both the log store and proxy containers.
 - Pod unschedulable. Also, a condition is shown for a number of issues; see **Example condition messages**.
- 4 The log store nodes in the cluster, with **upgradeStatus**.
- 5 The log store client, data, and master pods in the cluster, listed under 'failed', **notReady**, or **ready** state.

12.2.1.1. Example condition messages

The following are examples of some condition messages from the **Status** section of the Elasticsearch instance.

The following status message indicates that a node has exceeded the configured low watermark, and no shard will be allocated to this node.

```

status:
  nodes:
    - conditions:
      - lastTransitionTime: 2019-03-15T15:57:22Z
        message: Disk storage usage for node is 27.5gb (36.74%). Shards will be not
          be allocated on this node.
        reason: Disk Watermark Low

```

```

status: "True"
type: NodeStorage
deploymentName: example-elasticsearch-cdm-0-1
upgradeStatus: {}

```

The following status message indicates that a node has exceeded the configured high watermark, and shards will be relocated to other nodes.

```

status:
  nodes:
  - conditions:
    - lastTransitionTime: 2019-03-15T16:04:45Z
      message: Disk storage usage for node is 27.5gb (36.74%). Shards will be relocated
        from this node.
      reason: Disk Watermark High
      status: "True"
      type: NodeStorage
    deploymentName: example-elasticsearch-cdm-0-1
    upgradeStatus: {}

```

The following status message indicates that the log store node selector in the CR does not match any nodes in the cluster:

```

status:
  nodes:
  - conditions:
    - lastTransitionTime: 2019-04-10T02:26:24Z
      message: '0/8 nodes are available: 8 node(s) didn't match node selector.'
      reason: Unschedulable
      status: "True"
      type: Unschedulable

```

The following status message indicates that the log store CR uses a non-existent persistent volume claim (PVC).

```

status:
  nodes:
  - conditions:
    - last Transition Time: 2019-04-10T05:55:51Z
      message: pod has unbound immediate PersistentVolumeClaims (repeated 5 times)
      reason: Unschedulable
      status: True
      type: Unschedulable

```

The following status message indicates that your log store cluster does not have enough nodes to support the redundancy policy.

```

status:
  clusterHealth: ""
  conditions:
  - lastTransitionTime: 2019-04-17T20:01:31Z
    message: Wrong RedundancyPolicy selected. Choose different RedundancyPolicy or
      add more nodes with data roles

```

```
reason: Invalid Settings
status: "True"
type: InvalidRedundancy
```

This status message indicates your cluster has too many control plane nodes:

```
status:
clusterHealth: green
conditions:
- lastTransitionTime: '2019-04-17T20:12:34Z'
  message: >-
    Invalid master nodes count. Please ensure there are no more than 3 total
    nodes with master roles
  reason: Invalid Settings
  status: 'True'
  type: InvalidMasters
```

The following status message indicates that Elasticsearch storage does not support the change you tried to make.

For example:

```
status:
clusterHealth: green
conditions:
- lastTransitionTime: "2021-05-07T01:05:13Z"
  message: Changing the storage structure for a custom resource is not supported
  reason: StorageStructureChangeIgnored
  status: 'True'
  type: StorageStructureChangeIgnored
```

The **reason** and **type** fields specify the type of unsupported change:

StorageClassNameChangeIgnored

Unsupported change to the storage class name.

StorageSizeChangeIgnored

Unsupported change the storage size.

StorageStructureChangeIgnored

Unsupported change between ephemeral and persistent storage structures.



IMPORTANT

If you try to configure the **ClusterLogging** custom resource (CR) to switch from ephemeral to persistent storage, the OpenShift Elasticsearch Operator creates a persistent volume claim (PVC) but does not create a persistent volume (PV). To clear the **StorageStructureChangeIgnored** status, you must revert the change to the **ClusterLogging** CR and delete the PVC.

12.2.2. Viewing the status of the log store components

You can view the status for a number of the log store components.

Elasticsearch indices

You can view the status of the Elasticsearch indices.

1. Get the name of an Elasticsearch pod:

```
$ oc get pods --selector component=elasticsearch -o name
```

Example output

```
pod/elasticsearch-cdm-1godmszn-1-6f8495-vp4lw
pod/elasticsearch-cdm-1godmszn-2-5769cf-9ms2n
pod/elasticsearch-cdm-1godmszn-3-f66f7d-zqkz7
```

2. Get the status of the indices:

```
$ oc exec elasticsearch-cdm-4vjor49p-2-6d4d7db474-q2w7z -- indices
```

Example output

```
Defaulting container name to elasticsearch.
Use 'oc describe pod/elasticsearch-cdm-4vjor49p-2-6d4d7db474-q2w7z -n openshift-logging' to see all of the containers in this pod.

green open  infra-000002                S4QANnf1QP6NgCegfnrnBQ
3 1 119926      0 157      78
green open  audit-000001                8_EQx77iQCSTzFOXtxRqFw
3 1 0          0 0          0
green open  .security                    iDjSCH7aSUGhldq0LheLBQ 1
1 5 0          0 0          0
green open  .kibana_-377444158_kubeadmin  yBywZ9GfSrKebz5gWBZbjw 3 1 1 0 0 0
green open  infra-000001                z6Dpe__ORgiopEpW6YI44A
3 1 871000     0 874      436
green open  app-000001                   hlrazQCeSISewG3c2VlvsQ
3 1 2453       0 3          1
green open  .kibana_1                    JCitcBMSQxKOvlq6iQW6wg
1 1 0          0 0          0
green open  .kibana_-1595131456_user1    glYFIEGRRRe-
ka0W3okS-mQ 3 1 1 0 0 0
```

Log store pods

You can view the status of the pods that host the log store.

1. Get the name of a pod:

```
$ oc get pods --selector component=elasticsearch -o name
```

Example output

```
pod/elasticsearch-cdm-1godmszn-1-6f8495-vp4lw
pod/elasticsearch-cdm-1godmszn-2-5769cf-9ms2n
pod/elasticsearch-cdm-1godmszn-3-f66f7d-zqkz7
```


2. Get the status of a pod:

```
$ oc describe pod elasticsearch-cdm-1godmszn-1-6f8495-vp4lw
```

The output includes the following status information:

Example output

```
....
Status:      Running

....

Containers:
  elasticsearch:
    Container ID:  cri-o://b7d44e0a9ea486e27f47763f5bb4c39dfd2
    State:        Running
      Started:    Mon, 08 Jun 2020 10:17:56 -0400
    Ready:       True
    Restart Count: 0
    Readiness:    exec [/usr/share/elasticsearch/probe/readiness.sh] delay=10s timeout=30s
                  period=5s #success=1 #failure=3

....

  proxy:
    Container ID:  cri-
o://3f77032abaddbb1652c116278652908dc01860320b8a4e741d06894b2f8f9aa1
    State:        Running
      Started:    Mon, 08 Jun 2020 10:18:38 -0400
    Ready:       True
    Restart Count: 0

....

Conditions:
  Type           Status
  Initialized    True
  Ready          True
  ContainersReady True
  PodScheduled   True

....

Events:         <none>
```

Log storage pod deployment configuration

You can view the status of the log store deployment configuration.

1. Get the name of a deployment configuration:

```
$ oc get deployment --selector component=elasticsearch -o name
```

Example output

```
deployment.extensions/elasticsearch-cdm-1gon-1
deployment.extensions/elasticsearch-cdm-1gon-2
deployment.extensions/elasticsearch-cdm-1gon-3
```

2. Get the deployment configuration status:

```
$ oc describe deployment elasticsearch-cdm-1gon-1
```

The output includes the following status information:

Example output

```
....
Containers:
  elasticsearch:
    Image: registry.redhat.io/openshift-logging/elasticsearch6-rhel8
    Readiness: exec [/usr/share/elasticsearch/probe/readiness.sh] delay=10s timeout=30s
              period=5s #success=1 #failure=3
....

Conditions:
  Type           Status  Reason
  ----           -
  Progressing    Unknown DeploymentPaused
  Available      True    MinimumReplicasAvailable
....

Events:          <none>
```

Log store replica set

You can view the status of the log store replica set.

1. Get the name of a replica set:

```
$ oc get replicaSet --selector component=elasticsearch -o name

replicaset.extensions/elasticsearch-cdm-1gon-1-6f8495
replicaset.extensions/elasticsearch-cdm-1gon-2-5769cf
replicaset.extensions/elasticsearch-cdm-1gon-3-f66f7d
```

2. Get the status of the replica set:

```
$ oc describe replicaSet elasticsearch-cdm-1gon-1-6f8495
```

The output includes the following status information:

Example output

```
....
Containers:
  elasticsearch:
```

```

Image: registry.redhat.io/openshift-logging/elasticsearch6-
rhel8@sha256:4265742c7cdd85359140e2d7d703e4311b6497eec7676957f455d6908e7b1
c25
Readiness: exec [/usr/share/elasticsearch/probe/readiness.sh] delay=10s timeout=30s
period=5s #success=1 #failure=3
....
Events: <none>

```

12.2.3. Elasticsearch cluster status

The Grafana dashboard in the **Observe** section of the OpenShift Container Platform web console displays the status of the Elasticsearch cluster.

To get the status of the OpenShift Elasticsearch cluster, visit the Grafana dashboard in the **Observe** section of the OpenShift Container Platform web console at **<cluster_url>/monitoring/dashboards/grafana-dashboard-cluster-logging**.

Elasticsearch status fields

eo_elasticsearch_cr_cluster_management_state

Shows whether the Elasticsearch cluster is in a managed or unmanaged state. For example:

```

eo_elasticsearch_cr_cluster_management_state{state="managed"} 1
eo_elasticsearch_cr_cluster_management_state{state="unmanaged"} 0

```

eo_elasticsearch_cr_restart_total

Shows the number of times the Elasticsearch nodes have restarted for certificate restarts, rolling restarts, or scheduled restarts. For example:

```

eo_elasticsearch_cr_restart_total{reason="cert_restart"} 1
eo_elasticsearch_cr_restart_total{reason="rolling_restart"} 1
eo_elasticsearch_cr_restart_total{reason="scheduled_restart"} 3

```

es_index_namespaces_total

Shows the total number of Elasticsearch index namespaces. For example:

```

Total number of Namespaces.
es_index_namespaces_total 5

```

es_index_document_count

Shows the number of records for each namespace. For example:

```

es_index_document_count{namespace="namespace_1"} 25
es_index_document_count{namespace="namespace_2"} 10
es_index_document_count{namespace="namespace_3"} 5

```

The "Secret Elasticsearch fields are either missing or empty" message

If Elasticsearch is missing the **admin-cert**, **admin-key**, **logging-es.crt**, or **logging-es.key** files, the dashboard shows a status message similar to the following example:

```
message": "Secret \"elasticsearch\" fields are either missing or empty: [admin-cert, admin-key,
logging-es.crt, logging-es.key]",
"reason": "Missing Required Secrets",
```

12.3. UNDERSTANDING LOGGING SUBSYSTEM ALERTS

All of the logging collector alerts are listed on the Alerting UI of the OpenShift Container Platform web console.

12.3.1. Viewing logging collector alerts

Alerts are shown in the OpenShift Container Platform web console, on the **Alerts** tab of the Alerting UI. Alerts are in one of the following states:

- **Firing.** The alert condition is true for the duration of the timeout. Click the **Options** menu at the end of the firing alert to view more information or silence the alert.
- **Pending** The alert condition is currently true, but the timeout has not been reached.
- **Not Firing.** The alert is not currently triggered.

Procedure

To view the logging subsystem and other OpenShift Container Platform alerts:

1. In the OpenShift Container Platform console, click **Observe** → **Alerting**.
2. Click the **Alerts** tab. The alerts are listed, based on the filters selected.

Additional resources

- For more information on the Alerting UI, see [Managing alerts](#).

12.3.2. About logging collector alerts

The following alerts are generated by the logging collector. You can view these alerts in the OpenShift Container Platform web console, on the **Alerts** page of the Alerting UI.

Table 12.1. Fluentd Prometheus alerts

Alert	Message	Description	Severity
FluentDHighErrorRate	<value> of records have resulted in an error by fluentd <instance>.	The number of FluentD output errors is high, by default more than 10 in the previous 15 minutes.	Warning

Alert	Message	Description	Severity
FluentdNodeDown	Prometheus could not scrape fluentd <instance> for more than 10m.	Fluentd is reporting that Prometheus could not scrape a specific Fluentd instance.	Critical
FluentdQueueLengthIncreasing	In the last 12h, fluentd <instance> buffer queue length constantly increased more than 1. Current value is <value>.	Fluentd is reporting that the queue size is increasing.	Critical
FluentDVeryHighErrorRate	<value> of records have resulted in an error by fluentd <instance>.	The number of FluentD output errors is very high, by default more than 25 in the previous 15 minutes.	Critical

12.3.3. About Elasticsearch alerting rules

You can view these alerting rules in Prometheus.

Table 12.2. Alerting rules

Alert	Description	Severity
ElasticsearchClusterNotHealthy	The cluster health status has been RED for at least 2 minutes. The cluster does not accept writes, shards may be missing, or the master node hasn't been elected yet.	Critical
ElasticsearchClusterNotHealthy	The cluster health status has been YELLOW for at least 20 minutes. Some shard replicas are not allocated.	Warning
ElasticsearchDiskSpaceRunningLow	The cluster is expected to be out of disk space within the next 6 hours.	Critical
ElasticsearchHighFileDescriptorUsage	The cluster is predicted to be out of file descriptors within the next hour.	Warning
ElasticsearchJVMHeapUseHigh	The JVM Heap usage on the specified node is high.	Alert
ElasticsearchNodeDiskWatermarkReached	The specified node has hit the low watermark due to low free disk space. Shards can not be allocated to this node anymore. You should consider adding more disk space to the node.	Info

Alert	Description	Severity
ElasticsearchNodeDiskWatermarkReached	The specified node has hit the high watermark due to low free disk space. Some shards will be re-allocated to different nodes if possible. Make sure more disk space is added to the node or drop old indices allocated to this node.	Warning
ElasticsearchNodeDiskWatermarkReached	The specified node has hit the flood watermark due to low free disk space. Every index that has a shard allocated on this node is enforced a read-only block. The index block must be manually released when the disk use falls below the high watermark.	Critical
ElasticsearchJVMHeapUseHigh	The JVM Heap usage on the specified node is too high.	Alert
ElasticsearchWriteRequestsRejectionJumps	Elasticsearch is experiencing an increase in write rejections on the specified node. This node might not be keeping up with the indexing speed.	Warning
AggregatedLoggingSystemCPUHigh	The CPU used by the system on the specified node is too high.	Alert
ElasticsearchProcessCPUHigh	The CPU used by Elasticsearch on the specified node is too high.	Alert

12.4. COLLECTING LOGGING DATA FOR RED HAT SUPPORT

When opening a support case, it is helpful to provide debugging information about your cluster to Red Hat Support.

The [must-gather tool](#) enables you to collect diagnostic information for project-level resources, cluster-level resources, and each of the logging subsystem components.

For prompt support, supply diagnostic information for both OpenShift Container Platform and OpenShift Logging.



NOTE

Do not use the **hack/logging-dump.sh** script. The script is no longer supported and does not collect data.

12.4.1. About the must-gather tool

The **oc adm must-gather** CLI command collects the information from your cluster that is most likely needed for debugging issues.

For your logging subsystem, **must-gather** collects the following information:

- Project-level resources, including pods, configuration maps, service accounts, roles, role bindings, and events at the project level

- Cluster-level resources, including nodes, roles, and role bindings at the cluster level
- OpenShift Logging resources in the **openshift-logging** and **openshift-operators-redhat** namespaces, including health status for the log collector, the log store, and the log visualizer

When you run **oc adm must-gather**, a new pod is created on the cluster. The data is collected on that pod and saved in a new directory that starts with **must-gather.local**. This directory is created in the current working directory.

12.4.2. Prerequisites

- The logging subsystem and Elasticsearch must be installed.

12.4.3. Collecting OpenShift Logging data

You can use the **oc adm must-gather** CLI command to collect information about your logging subsystem.

Procedure

To collect logging subsystem information with **must-gather**:

1. Navigate to the directory where you want to store the **must-gather** information.
2. Run the **oc adm must-gather** command against the OpenShift Logging image:

```
$ oc adm must-gather --image=$(oc -n openshift-logging get deployment.apps/cluster-logging-operator -o jsonpath='{.spec.template.spec.containers[?(@.name == "cluster-logging-operator")].image}')
```

The **must-gather** tool creates a new directory that starts with **must-gather.local** within the current directory. For example: **must-gather.local.4157245944708210408**.

3. Create a compressed file from the **must-gather** directory that was just created. For example, on a computer that uses a Linux operating system, run the following command:

```
$ tar -cvaf must-gather.tar.gz must-gather.local.4157245944708210408
```

4. Attach the compressed file to your support case on the [Red Hat Customer Portal](#).

12.5. TROUBLESHOOTING FOR CRITICAL ALERTS

12.5.1. Elasticsearch Cluster Health is Red

At least one primary shard and its replicas are not allocated to a node.

Troubleshooting

1. Check the Elasticsearch cluster health and verify that the cluster **status** is red.

```
oc exec -n openshift-logging -c elasticsearch <elasticsearch_pod_name> -- health
```

2. List the nodes that have joined the cluster.

```
oc exec -n openshift-logging -c elasticsearch <elasticsearch_pod_name> -- es_util --
query=_cat/nodes?v
```

3. List the Elasticsearch pods and compare them with the nodes in the command output from the previous step.

```
oc -n openshift-logging get pods -l component=elasticsearch
```

4. If some of the Elasticsearch nodes have not joined the cluster, perform the following steps.
 - a. Confirm that Elasticsearch has an elected control plane node.

```
oc exec -n openshift-logging -c elasticsearch <elasticsearch_pod_name> -- es_util --
query=_cat/master?v
```

- b. Review the pod logs of the elected control plane node for issues.

```
oc logs <elasticsearch_master_pod_name> -c elasticsearch -n openshift-logging
```

- c. Review the logs of nodes that have not joined the cluster for issues.

```
oc logs <elasticsearch_node_name> -c elasticsearch -n openshift-logging
```

5. If all the nodes have joined the cluster, perform the following steps, check if the cluster is in the process of recovering.

```
oc exec -n openshift-logging -c elasticsearch <elasticsearch_pod_name> -- es_util --
query=_cat/recovery?active_only=true
```

If there is no command output, the recovery process might be delayed or stalled by pending tasks.

6. Check if there are pending tasks.

```
oc exec -n openshift-logging -c elasticsearch <elasticsearch_pod_name> -- health |grep
number_of_pending_tasks
```

7. If there are pending tasks, monitor their status.

If their status changes and indicates that the cluster is recovering, continue waiting. The recovery time varies according to the size of the cluster and other factors.

Otherwise, if the status of the pending tasks does not change, this indicates that the recovery has stalled.

8. If it seems like the recovery has stalled, check if **cluster.routing.allocation.enable** is set to **none**.

```
oc exec -n openshift-logging -c elasticsearch <elasticsearch_pod_name> -- es_util --
query=_cluster/settings?pretty
```

9. If **cluster.routing.allocation.enable** is set to **none**, set it to **all**.


```
oc exec -n openshift-logging -c elasticsearch <elasticsearch_pod_name> -- es_util --
query=_cluster/settings?pretty -X PUT -d '{"persistent":
{"cluster.routing.allocation.enable":"all"}}'
```

10. Check which indices are still red.

```
oc exec -n openshift-logging -c elasticsearch <elasticsearch_pod_name> -- es_util --
query=_cat/indices?v
```

11. If any indices are still red, try to clear them by performing the following steps.

- a. Clear the cache.

```
oc exec -n openshift-logging -c elasticsearch <elasticsearch_pod_name> -- es_util --
query=<elasticsearch_index_name>/_cache/clear?pretty
```

- b. Increase the max allocation retries.

```
oc exec -n openshift-logging -c elasticsearch <elasticsearch_pod_name> -- es_util --
query=<elasticsearch_index_name>/_settings?pretty -X PUT -d
'{"index.allocation.max_retries":10}'
```

- c. Delete all the scroll items.

```
oc exec -n openshift-logging -c elasticsearch <elasticsearch_pod_name> -- es_util --
query=_search/scroll/_all -X DELETE
```

- d. Increase the timeout.

```
oc exec -n openshift-logging -c elasticsearch <elasticsearch_pod_name> -- es_util --
query=<elasticsearch_index_name>/_settings?pretty -X PUT -d
'{"index.unassigned.node_left.delayed_timeout":"10m"}'
```

12. If the preceding steps do not clear the red indices, delete the indices individually.

- a. Identify the red index name.

```
oc exec -n openshift-logging -c elasticsearch <elasticsearch_pod_name> -- es_util --
query=_cat/indices?v
```

- b. Delete the red index.

```
oc exec -n openshift-logging -c elasticsearch <elasticsearch_pod_name> -- es_util --
query=<elasticsearch_red_index_name> -X DELETE
```

13. If there are no red indices and the cluster status is red, check for a continuous heavy processing load on a data node.

- a. Check if the Elasticsearch JVM Heap usage is high.

```
oc exec -n openshift-logging -c elasticsearch <elasticsearch_pod_name> -- es_util --
query=_nodes/stats?pretty
```

In the command output, review the **node_name.jvm.mem.heap_used_percent** field to determine the JVM Heap usage.

- b. Check for high CPU utilization.

Additional resources

- Search for "Free up or increase disk space" in the Elasticsearch topic, [Fix a red or yellow cluster status](#).

12.5.2. Elasticsearch Cluster Health is Yellow

Replica shards for at least one primary shard are not allocated to nodes.

Troubleshooting

1. Increase the node count by adjusting **nodeCount** in the **ClusterLogging** CR.

Additional resources

- [About the Cluster Logging custom resource](#)
- [Configuring persistent storage for the log store](#)
- Search for "Free up or increase disk space" in the Elasticsearch topic, [Fix a red or yellow cluster status](#).

12.5.3. Elasticsearch Node Disk Low Watermark Reached

Elasticsearch does not allocate shards to nodes that [reach the low watermark](#).

Troubleshooting

1. Identify the node on which Elasticsearch is deployed.

```
oc -n openshift-logging get po -o wide
```

2. Check if there are **unassigned shards**.

```
oc exec -n openshift-logging -c elasticsearch <elasticsearch_pod_name> -- es_util --  
query=_cluster/health?pretty | grep unassigned_shards
```

3. If there are unassigned shards, check the disk space on each node.

```
for pod in `oc -n openshift-logging get po -l component=elasticsearch -o  
jsonpath='{.items[*].metadata.name}'`; do echo $pod; oc -n openshift-logging exec -c  
elasticsearch $pod -- df -h /elasticsearch/persistent; done
```

4. Check the **nodes.node_name.fs** field to determine the free disk space on that node. If the used disk percentage is above 85%, the node has exceeded the low watermark, and shards can no longer be allocated to this node.
5. Try to increase the disk space on all nodes.

6. If increasing the disk space is not possible, try adding a new data node to the cluster.
7. If adding a new data node is problematic, decrease the total cluster redundancy policy.
 - a. Check the current **redundancyPolicy**.

```
oc -n openshift-logging get es elasticsearch -o jsonpath='{.spec.redundancyPolicy}'
```



NOTE

If you are using a **ClusterLogging** CR, enter:

```
oc -n openshift-logging get cl -o jsonpath='{.items[*].spec.logStore.elasticsearch.redundancyPolicy}'
```

- b. If the cluster **redundancyPolicy** is higher than **SingleRedundancy**, set it to **SingleRedundancy** and save this change.
8. If the preceding steps do not fix the issue, delete the old indices.
 - a. Check the status of all indices on Elasticsearch.

```
oc exec -n openshift-logging -c elasticsearch <elasticsearch_pod_name> -- indices
```

- b. Identify an old index that can be deleted.
 - c. Delete the index.

```
oc exec -n openshift-logging -c elasticsearch <elasticsearch_pod_name> -- es_util -- query=<elasticsearch_index_name> -X DELETE
```

Additional resources

- Search for "redundancyPolicy" in the "Sample **ClusterLogging** custom resource (CR)" in [About the Cluster Logging custom resource](#)

12.5.4. Elasticsearch Node Disk High Watermark Reached

Elasticsearch attempts to relocate shards away from a node [that has reached the high watermark](#) .

Troubleshooting

1. Identify the node on which Elasticsearch is deployed.

```
oc -n openshift-logging get po -o wide
```

2. Check the disk space on each node.

```
for pod in `oc -n openshift-logging get po -l component=elasticsearch -o jsonpath='{.items[*].metadata.name}'`; do echo $pod; oc -n openshift-logging exec -c elasticsearch $pod -- df -h /elasticsearch/persistent; done
```

3. Check if the cluster is rebalancing.

```
oc exec -n openshift-logging -c elasticsearch <elasticsearch_pod_name> -- es_util --
query=_cluster/health?pretty | grep relocating_shards
```

If the command output shows relocating shards, the High Watermark has been exceeded. The default value of the High Watermark is 90%.

The shards relocate to a node with low disk usage that has not crossed any watermark threshold limits.

4. To allocate shards to a particular node, free up some space.
5. Try to increase the disk space on all nodes.
6. If increasing the disk space is not possible, try adding a new data node to the cluster.
7. If adding a new data node is problematic, decrease the total cluster redundancy policy.
 - a. Check the current **redundancyPolicy**.

```
oc -n openshift-logging get es elasticsearch -o jsonpath='{.spec.redundancyPolicy}'
```



NOTE

If you are using a **ClusterLogging** CR, enter:

```
oc -n openshift-logging get cl -o
jsonpath='{.items[*].spec.logStore.elasticsearch.redundancyPolicy}'
```

- b. If the cluster **redundancyPolicy** is higher than **SingleRedundancy**, set it to **SingleRedundancy** and save this change.
8. If the preceding steps do not fix the issue, delete the old indices.
 - a. Check the status of all indices on Elasticsearch.

```
oc exec -n openshift-logging -c elasticsearch <elasticsearch_pod_name> -- indices
```

- b. Identify an old index that can be deleted.
 - c. Delete the index.

```
oc exec -n openshift-logging -c elasticsearch <elasticsearch_pod_name> -- es_util --
query=<elasticsearch_index_name> -X DELETE
```

Additional resources

- Search for "redundancyPolicy" in the "Sample **ClusterLogging** custom resource (CR)" in [About the Cluster Logging custom resource](#)

12.5.5. Elasticsearch Node Disk Flood Watermark Reached

Elasticsearch enforces a read-only index block on every index that has both of these conditions:

- One or more shards are allocated to the node.
- One or more disks exceed the [flood stage](#).

Troubleshooting

1. Check the disk space of the Elasticsearch node.

```
for pod in `oc -n openshift-logging get po -l component=elasticsearch -o
jsonpath='{.items[*].metadata.name}'`; do echo $pod; oc -n openshift-logging exec -c
elasticsearch $pod -- df -h /elasticsearch/persistent; done
```

Check the **nodes.node_name.fs** field to determine the free disk space on that node.

2. If the used disk percentage is above 95%, it signifies that the node has crossed the flood watermark. Writing is blocked for shards allocated on this particular node.
3. Try to increase the disk space on all nodes.
4. If increasing the disk space is not possible, try adding a new data node to the cluster.
5. If adding a new data node is problematic, decrease the total cluster redundancy policy.
 - a. Check the current **redundancyPolicy**.

```
oc -n openshift-logging get es elasticsearch -o jsonpath='{.spec.redundancyPolicy}'
```



NOTE

If you are using a **ClusterLogging** CR, enter:

```
oc -n openshift-logging get cl -o
jsonpath='{.items[*].spec.logStore.elasticsearch.redundancyPolicy}'
```

- b. If the cluster **redundancyPolicy** is higher than **SingleRedundancy**, set it to **SingleRedundancy** and save this change.
6. If the preceding steps do not fix the issue, delete the old indices.
 - a. Check the status of all indices on Elasticsearch.


```
oc exec -n openshift-logging -c elasticsearch <elasticsearch_pod_name> -- indices
```
 - b. Identify an old index that can be deleted.
 - c. Delete the index.


```
oc exec -n openshift-logging -c elasticsearch <elasticsearch_pod_name> -- es_util --
query=<elasticsearch_index_name> -X DELETE
```
 7. Continue freeing up and monitoring the disk space until the used disk space drops below 90%. Then, unblock write to this particular node.

```
oc exec -n openshift-logging -c elasticsearch <elasticsearch_pod_name> -- es_util --
query=_all/_settings?pretty -X PUT -d '{"index.blocks.read_only_allow_delete": null}'
```

Additional resources

- Search for "redundancyPolicy" in the "Sample **ClusterLogging** custom resource (CR)" in [About the Cluster Logging custom resource](#)

12.5.6. Elasticsearch JVM Heap Use is High

The Elasticsearch node JVM Heap memory used is above 75%.

Troubleshooting

Consider [increasing the heap size](#).

12.5.7. Aggregated Logging System CPU is High

System CPU usage on the node is high.

Troubleshooting

Check the CPU of the cluster node. Consider allocating more CPU resources to the node.

12.5.8. Elasticsearch Process CPU is High

Elasticsearch process CPU usage on the node is high.

Troubleshooting

Check the CPU of the cluster node. Consider allocating more CPU resources to the node.

12.5.9. Elasticsearch Disk Space is Running Low

The Elasticsearch Cluster is predicted to be out of disk space within the next 6 hours based on current disk usage.

Troubleshooting

1. Get the disk space of the Elasticsearch node.

```
for pod in `oc -n openshift-logging get po -l component=elasticsearch -o
jsonpath='{.items[*].metadata.name}'`; do echo $pod; oc -n openshift-logging exec -c
elasticsearch $pod -- df -h /elasticsearch/persistent; done
```

2. In the command output, check the **nodes.node_name.fs** field to determine the free disk space on that node.
3. Try to increase the disk space on all nodes.
4. If increasing the disk space is not possible, try adding a new data node to the cluster.
5. If adding a new data node is problematic, decrease the total cluster redundancy policy.
 - a. Check the current **redundancyPolicy**.

```
oc -n openshift-logging get es elasticsearch -o jsonpath='{.spec.redundancyPolicy}'
```



NOTE

If you are using a **ClusterLogging** CR, enter:

```
oc -n openshift-logging get cl -o
jsonpath='{.items[*].spec.logStore.elasticsearch.redundancyPolicy}'
```

- b. If the cluster **redundancyPolicy** is higher than **SingleRedundancy**, set it to **SingleRedundancy** and save this change.
6. If the preceding steps do not fix the issue, delete the old indices.
 - a. Check the status of all indices on Elasticsearch.

```
oc exec -n openshift-logging -c elasticsearch <elasticsearch_pod_name> -- indices
```

- b. Identify an old index that can be deleted.
- c. Delete the index.

```
oc exec -n openshift-logging -c elasticsearch <elasticsearch_pod_name> -- es_util --
query=<elasticsearch_index_name> -X DELETE
```

Additional resources

- Search for "redundancyPolicy" in the "Sample **ClusterLogging** custom resource (CR)" in [About the Cluster Logging custom resource](#)
- Search for "ElasticsearchDiskSpaceRunningLow" in [About Elasticsearch alerting rules](#).
- Search for "Free up or increase disk space" in the Elasticsearch topic, [Fix a red or yellow cluster status](#).

12.5.10. Elasticsearch FileDescriptor Usage is high

Based on current usage trends, the predicted number of file descriptors on the node is insufficient.

Troubleshooting

Check and, if needed, configure the value of **max_file_descriptors** for each node, as described in the Elasticsearch [File descriptors](#) topic.

Additional resources

- Search for "ElasticsearchHighFileDescriptorUsage" in [About Elasticsearch alerting rules](#).
- Search for "File Descriptors In Use" in [OpenShift Logging dashboards](#).

CHAPTER 13. UNINSTALLING OPENSIFT LOGGING

You can remove the logging subsystem from your OpenShift Container Platform cluster.

13.1. UNINSTALLING THE LOGGING SUBSYSTEM FOR RED HAT OPENSIFT

You can stop log aggregation by deleting the **ClusterLogging** custom resource (CR). After deleting the CR, there are other logging subsystem components that remain, which you can optionally remove.

Deleting the **ClusterLogging** CR does not remove the persistent volume claims (PVCs). To preserve or delete the remaining PVCs, persistent volumes (PVs), and associated data, you must take further action.

Prerequisites


- The Red Hat OpenShift Logging and Elasticsearch Operators must be installed.

Procedure

To remove OpenShift Logging:


1. Use the OpenShift Container Platform web console to remove the **ClusterLogging** CR:


- a. Switch to the **Administration** → **Custom Resource Definitions** page.
- b. On the **Custom Resource Definitions** page, click **ClusterLogging**.
- c. On the **Custom Resource Definition Details** page, click **Instances**.


- d. Click the Options menu  next to the instance and select **Delete ClusterLogging**.

2. Optional: Delete the custom resource definitions (CRD):

- a. Switch to the **Administration** → **Custom Resource Definitions** page.


- b. Click the Options menu  next to **ClusterLogForwarder** and select **Delete Custom Resource Definition**.

- c. Click the Options menu  next to **ClusterLogging** and select **Delete Custom Resource Definition**.

- d. Click the Options menu  next to **Elasticsearch** and select **Delete Custom Resource Definition**.

3. Optional: Remove the Red Hat OpenShift Logging Operator and OpenShift Elasticsearch Operator:

- a. Switch to the **Operators** → **Installed Operators** page.

- b. Click the Options menu  next to the Red Hat OpenShift Logging Operator and select **Uninstall Operator**.
 - c. Click the Options menu  next to the OpenShift Elasticsearch Operator and select **Uninstall Operator**.
4. Optional: Remove the OpenShift Logging and Elasticsearch projects.
- a. Switch to the **Home → Projects** page.
 - b. Click the Options menu  next to the **openshift-logging** project and select **Delete Project**.
 - c. Confirm the deletion by typing **openshift-logging** in the dialog box and click **Delete**.
 - d. Click the Options menu  next to the **openshift-operators-redhat** project and select **Delete Project**.


**IMPORTANT**

Do not delete the **openshift-operators-redhat** project if other global operators are installed in this namespace.

- e. Confirm the deletion by typing **openshift-operators-redhat** in the dialog box and click **Delete**.
5. To keep the PVCs for reuse with other pods, keep the labels or PVC names that you need to reclaim the PVCs.
6. Optional: If you do not want to keep the PVCs, you can delete them.

**WARNING**

Releasing or deleting PVCs can delete PVs and cause data loss.

- a. Switch to the **Storage → Persistent Volume Claims** page.
- b. Click the Options menu  next to each PVC and select **Delete Persistent Volume Claim**.
- c. If you want to recover storage space, you can delete the PVs.

Additional resources

- [Reclaiming a persistent volume manually](#)

CHAPTER 14. LOG RECORD FIELDS

The following fields can be present in log records exported by the logging subsystem. Although log records are typically formatted as JSON objects, the same data model can be applied to other encodings.

To search these fields from Elasticsearch and Kibana, use the full dotted field name when searching. For example, with an Elasticsearch `/_search URL`, to look for a Kubernetes pod name, use `/_search/q=kubernetes.pod_name:name-of-my-pod`.

The top level fields may be present in every record.

CHAPTER 15. MESSAGE

The original log entry text, UTF-8 encoded. This field may be absent or empty if a non-empty **structured** field is present. See the description of **structured** for more.

Data type	text
Example value	HAPPY

CHAPTER 16. STRUCTURED

Original log entry as a structured object. This field may be present if the forwarder was configured to parse structured JSON logs. If the original log entry was a valid structured log, this field will contain an equivalent JSON structure. Otherwise this field will be empty or absent, and the **message** field will contain the original log message. The **structured** field can have any subfields that are included in the log message, there are no restrictions defined here.

Data type	group
Example value	map[message:starting fluentd worker pid=21631 ppid=21618 worker=0 pid:21631 ppid:21618 worker:0]

CHAPTER 17. @TIMESTAMP

A UTC value that marks when the log payload was created or, if the creation time is not known, when the log payload was first collected. The "@" prefix denotes a field that is reserved for a particular use. By default, most tools look for "@timestamp" with Elasticsearch.

Data type	date
Example value	2015-01-24 14:06:05.071000000 Z

CHAPTER 18. HOSTNAME

The name of the host where this log message originated. In a Kubernetes cluster, this is the same as **kubernetes.host**.

Data type	keyword
-----------	---------

CHAPTER 19. IPADDR4

The IPv4 address of the source server. Can be an array.

Data type	ip
-----------	----

CHAPTER 20. IPADDR6

The IPv6 address of the source server, if available. Can be an array.

Data type	ip
-----------	----

CHAPTER 21. LEVEL

The logging level from various sources, including **rsyslog(severitytext property)**, a Python logging module, and others.

The following values come from **syslog.h**, and are preceded by their [numeric equivalents](#):

- **0 = emerg**, system is unusable.
- **1 = alert**, action must be taken immediately.
- **2 = crit**, critical conditions.
- **3 = err**, error conditions.
- **4 = warn**, warning conditions.
- **5 = notice**, normal but significant condition.
- **6 = info**, informational.
- **7 = debug**, debug-level messages.

The two following values are not part of **syslog.h** but are widely used:

- **8 = trace**, trace-level messages, which are more verbose than **debug** messages.
- **9 = unknown**, when the logging system gets a value it doesn't recognize.

Map the log levels or priorities of other logging systems to their nearest match in the preceding list. For example, from [python logging](#), you can match **CRITICAL** with **crit**, **ERROR** with **err**, and so on.

Data type	keyword
Example value	info

CHAPTER 22. PID

The process ID of the logging entity, if available.

Data type	keyword
-----------	---------

CHAPTER 23. SERVICE

The name of the service associated with the logging entity, if available. For example, syslog's **APP-NAME** and rsyslog's **programname** properties are mapped to the service field.

Data type	keyword
-----------	---------

CHAPTER 24. TAGS

Optional. An operator-defined list of tags placed on each log by the collector or normalizer. The payload can be a string with whitespace-delimited string tokens or a JSON list of string tokens.

Data type	text
-----------	------

CHAPTER 25. FILE

The path to the log file from which the collector reads this log entry. Normally, this is a path in the **/var/log** file system of a cluster node.

Data type	text
-----------	------

CHAPTER 26. OFFSET

The offset value. Can represent bytes to the start of the log line in the file (zero- or one-based), or log line numbers (zero- or one-based), so long as the values are strictly monotonically increasing in the context of a single log file. The values are allowed to wrap, representing a new version of the log file (rotation).

Data type	long
-----------	------

CHAPTER 27. KUBERNETES

The namespace for Kubernetes-specific metadata

Data type	group
-----------	-------

27.1. KUBERNETES.POD_NAME

The name of the pod

Data type	keyword
-----------	---------

27.2. KUBERNETES.POD_ID

The Kubernetes ID of the pod

Data type	keyword
-----------	---------

27.3. KUBERNETES.NAMESPACE_NAME

The name of the namespace in Kubernetes

Data type	keyword
-----------	---------

27.4. KUBERNETES.NAMESPACE_ID

The ID of the namespace in Kubernetes

Data type	keyword
-----------	---------

27.5. KUBERNETES.HOST

The Kubernetes node name

Data type	keyword
-----------	---------

27.6. KUBERNETES.CONTAINER_NAME

The name of the container in Kubernetes

Data type	keyword
-----------	---------

27.7. KUBERNETES.ANNOTATIONS

Annotations associated with the Kubernetes object

Data type	group
-----------	-------

27.8. KUBERNETES.LABELS

Labels present on the original Kubernetes Pod

Data type	group
-----------	-------

27.9. KUBERNETES.EVENT

The Kubernetes event obtained from the Kubernetes master API. This event description loosely follows **type Event** in [Event v1 core](#).

Data type	group
-----------	-------

27.9.1. kubernetes.event.verb

The type of event, **ADDED**, **MODIFIED**, or **DELETED**

Data type	keyword
Example value	ADDED

27.9.2. kubernetes.event.metadata

Information related to the location and time of the event creation

Data type	group
-----------	-------

27.9.2.1. kubernetes.event.metadata.name

The name of the object that triggered the event creation

Data type	keyword
Example value	java-mainclass-1.14d888a4cfc24890

27.9.2.2. kubernetes.event.metadata.namespace

The name of the namespace where the event originally occurred. Note that it differs from **kubernetes.namespace_name**, which is the namespace where the **eventrouter** application is deployed.

Data type	keyword
Example value	default

27.9.2.3. kubernetes.event.metadata.selfLink

A link to the event

Data type	keyword
Example value	/api/v1/namespaces/javaj/events/java-mainclass-1.14d888a4cfc24890

27.9.2.4. kubernetes.event.metadata.uid

The unique ID of the event

Data type	keyword
Example value	d828ac69-7b58-11e7-9cf5-5254002f560c

27.9.2.5. kubernetes.event.metadata.resourceVersion

A string that identifies the server's internal version of the event. Clients can use this string to determine when objects have changed.

Data type	integer
Example value	311987

27.9.3. kubernetes.event.involvedObject

The object that the event is about.

Data type	group
-----------	-------

27.9.3.1. kubernetes.event.involvedObject.kind

The type of object

Data type	keyword
Example value	ReplicationController

27.9.3.2. `kubernetes.event.involvedObject.namespace`

The namespace name of the involved object. Note that it may differ from **`kubernetes.namespace_name`**, which is the namespace where the **`eventrouter`** application is deployed.

Data type	keyword
Example value	default

27.9.3.3. `kubernetes.event.involvedObject.name`

The name of the object that triggered the event

Data type	keyword
Example value	java-mainclass-1

27.9.3.4. `kubernetes.event.involvedObject.uid`

The unique ID of the object

Data type	keyword
Example value	e6bff941-76a8-11e7-8193-5254002f560c

27.9.3.5. `kubernetes.event.involvedObject.apiVersion`

The version of kubernetes master API

Data type	keyword
Example value	v1

27.9.3.6. `kubernetes.event.involvedObject.resourceVersion`

A string that identifies the server's internal version of the pod that triggered the event. Clients can use this string to determine when objects have changed.

Data type	keyword
Example value	308882

27.9.4. kubernetes.event.reason

A short machine-understandable string that gives the reason for generating this event

Data type	keyword
Example value	SuccessfulCreate

27.9.5. kubernetes.event.source_component

The component that reported this event

Data type	keyword
Example value	replication-controller

27.9.6. kubernetes.event.firstTimestamp

The time at which the event was first recorded

Data type	date
Example value	2017-08-07 10:11:57.000000000 Z

27.9.7. kubernetes.event.count

The number of times this event has occurred

Data type	integer
Example value	1

27.9.8. kubernetes.event.type

The type of event, **Normal** or **Warning**. New types could be added in the future.

Data type	keyword
Example value	Normal

CHAPTER 28. OPENSIFT

The namespace for openshift-logging specific metadata

Data type	group
-----------	-------

28.1. OPENSIFT.LABELS

Labels added by the Cluster Log Forwarder configuration

Data type	group
-----------	-------