



# Red Hat OpenStack Platform 13

## Configuration Reference

Configuring Red Hat OpenStack Platform environments



# Red Hat OpenStack Platform 13 Configuration Reference

---

Configuring Red Hat OpenStack Platform environments

OpenStack Documentation Team

[rhos-docs@redhat.com](mailto:rhos-docs@redhat.com)

OpenStack Team

[rhos-docs@redhat.com](mailto:rhos-docs@redhat.com)

## Legal Notice

Copyright © 2021 Red Hat, Inc.

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

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

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

## Abstract

This document is for system administrators who want to look up configuration options. It contains lists of configuration options available with OpenStack and uses auto-generation to generate options and the descriptions from the code for each project.

---

# Table of Contents

<b>PREFACE</b> .....	<b>12</b>
<b>CHAPTER 1. BARBICAN</b> .....	<b>13</b>
1.1. BARBICAN.CONF	13
1.1.1. DEFAULT	13
1.1.2. certificate	22
1.1.3. certificate_event	22
1.1.4. cors	23
1.1.5. crypto	24
1.1.6. dogtag_plugin	24
1.1.7. keystone_authtoken	25
1.1.8. keystone_notifications	29
1.1.9. kmip_plugin	30
1.1.10. matchmaker_redis	31
1.1.11. oslo_messaging_amqp	31
1.1.12. oslo_messaging_kafka	34
1.1.13. oslo_messaging_notifications	35
1.1.14. oslo_messaging_rabbit	36
1.1.15. oslo_messaging_zmq	40
1.1.16. oslo_middleware	43
1.1.17. oslo_policy	43
1.1.18. p11_crypto_plugin	44
1.1.19. queue	45
1.1.20. quotas	46
1.1.21. retry_scheduler	46
1.1.22. secretstore	47
1.1.23. simple_crypto_plugin	47
1.1.24. snakeoil_ca_plugin	48
1.1.25. ssl	48
<b>CHAPTER 2. CEILOMETER</b> .....	<b>49</b>
2.1. CEILOMETER.CONF	49
2.1.1. DEFAULT	49
2.1.2. compute	57
2.1.3. coordination	58
2.1.4. dispatcher_gnocchi	58
2.1.5. event	59
2.1.6. hardware	59
2.1.7. ipmi	60
2.1.8. matchmaker_redis	60
2.1.9. meter	61
2.1.10. notification	61
2.1.11. oslo_concurrency	62
2.1.12. oslo_messaging_amqp	63
2.1.13. oslo_messaging_kafka	66
2.1.14. oslo_messaging_notifications	67
2.1.15. oslo_messaging_rabbit	68
2.1.16. oslo_messaging_zmq	72
2.1.17. polling	75
2.1.18. publisher	75
2.1.19. publisher_notifier	76

2.1.20. rgw_admin_credentials	76
2.1.21. service_credentials	76
2.1.22. service_types	78
2.1.23. vmware	78
2.1.24. xenapi	79
<b>CHAPTER 3. CINDER</b> .....	<b>80</b>
3.1. CINDER.CONF	80
3.1.1. DEFAULT	80
3.1.2. backend	106
3.1.3. backend_defaults	106
3.1.4. barbican	142
3.1.5. brcd_fabric_example	143
3.1.6. cisco_fabric_example	144
3.1.7. coordination	144
3.1.8. cors	145
3.1.9. database	146
3.1.10. fc-zone-manager	147
3.1.11. healthcheck	148
3.1.12. key_manager	149
3.1.13. keystone_authtoken	150
3.1.14. matchmaker_redis	154
3.1.15. nova	154
3.1.16. oslo_concurrency	155
3.1.17. oslo_messaging_amqp	156
3.1.18. oslo_messaging_kafka	159
3.1.19. oslo_messaging_notifications	160
3.1.20. oslo_messaging_rabbit	160
3.1.21. oslo_messaging_zmq	165
3.1.22. oslo_middleware	168
3.1.23. oslo_policy	168
3.1.24. oslo_reports	169
3.1.25. oslo_versionedobjects	169
3.1.26. profiler	170
3.1.27. service_user	171
3.1.28. ssl	172
3.1.29. vault	172
<b>CHAPTER 4. DESIGNATE</b> .....	<b>174</b>
4.1. DESIGNATE.CONF	174
4.1.1. DEFAULT	174
4.1.2. backend:agent:bind9	182
4.1.3. backend:agent:denominator	183
4.1.4. backend:agent:djbdns	183
4.1.5. backend:agent:gdnsd	184
4.1.6. backend:agent:knot2	184
4.1.7. coordination	184
4.1.8. cors	185
4.1.9. database	185
4.1.10. handler:neutron_floatingip	187
4.1.11. handler:nova_fixed	188
4.1.12. healthcheck	188
4.1.13. heartbeat_emitter	189

4.1.14. keystone_authtoken	189
4.1.15. matchmaker_redis	192
4.1.16. monasca:statsd	193
4.1.17. network_api:neutron	193
4.1.18. oslo_concurrency	194
4.1.19. oslo_messaging_amqp	195
4.1.20. oslo_messaging_kafka	198
4.1.21. oslo_messaging_notifications	199
4.1.22. oslo_messaging_rabbit	199
4.1.23. oslo_messaging_zmq	204
4.1.24. oslo_middleware	207
4.1.25. oslo_policy	207
4.1.26. pool_manager_cache:memcache	208
4.1.27. pool_manager_cache:sqlalchemy	208
4.1.28. producer_task:delayed_notify	210
4.1.29. producer_task:periodic_exists	210
4.1.30. producer_task:periodic_secondary_refresh	210
4.1.31. producer_task:worker_periodic_recovery	211
4.1.32. producer_task:zone_purge	211
4.1.33. proxy	211
4.1.34. service:agent	212
4.1.35. service:api	213
4.1.36. service:central	214
4.1.37. service:mdns	215
4.1.38. service:pool_manager	216
4.1.39. service:producer	217
4.1.40. service:sink	218
4.1.41. service:worker	218
4.1.42. service:zone_manager	219
4.1.43. ssl	220
4.1.44. storage:sqlalchemy	220
<b>CHAPTER 5. GLANCE</b> .....	<b>223</b>
5.1. GLANCE-API.CONF	223
5.1.1. DEFAULT	223
5.1.2. cors	268
5.1.3. database	269
5.1.4. glance_store	271
5.1.5. image_format	310
5.1.6. keystone_authtoken	311
5.1.7. matchmaker_redis	314
5.1.8. oslo_concurrency	315
5.1.9. oslo_messaging_amqp	315
5.1.10. oslo_messaging_kafka	319
5.1.11. oslo_messaging_notifications	320
5.1.12. oslo_messaging_rabbit	320
5.1.13. oslo_messaging_zmq	324
5.1.14. oslo_middleware	327
5.1.15. oslo_policy	328
5.1.16. paste_deploy	329
5.1.17. profiler	330
5.1.18. store_type_location_strategy	332
5.1.19. task	333

5.1.20. taskflow_executor	335
5.2. GLANCE-REGISTRY.CONF	337
5.2.1. DEFAULT	337
5.2.2. database	366
5.2.3. keystone_authtoken	368
5.2.4. matchmaker_redis	372
5.2.5. oslo_messaging_amqp	372
5.2.6. oslo_messaging_kafka	376
5.2.7. oslo_messaging_notifications	377
5.2.8. oslo_messaging_rabbit	377
5.2.9. oslo_messaging_zmq	382
5.2.10. oslo_policy	384
5.2.11. paste_deploy	385
5.2.12. profiler	387
5.3. GLANCE-SCRUBBER.CONF	389
5.3.1. DEFAULT	389
5.3.2. database	411
5.3.3. glance_store	413
5.3.4. oslo_concurrency	451
5.3.5. oslo_policy	452
5.4. GLANCE-CACHE.CONF	453
5.4.1. DEFAULT	453
5.4.2. glance_store	482
5.4.3. oslo_policy	520
<b>CHAPTER 6. HEAT</b> .....	<b>522</b>
6.1. HEAT.CONF	522
6.1.1. DEFAULT	522
6.1.2. auth_password	534
6.1.3. clients	534
6.1.4. clients_aodh	535
6.1.5. clients_barbican	535
6.1.6. clients_ceilometer	536
6.1.7. clients_cinder	536
6.1.8. clients_designate	537
6.1.9. clients_glance	537
6.1.10. clients_heat	538
6.1.11. clients_keystone	538
6.1.12. clients_magnum	539
6.1.13. clients_manila	539
6.1.14. clients_mistral	540
6.1.15. clients_monasca	540
6.1.16. clients_neutron	541
6.1.17. clients_nova	541
6.1.18. clients_octavia	542
6.1.19. clients_sahara	542
6.1.20. clients_senlin	543
6.1.21. clients_swift	543
6.1.22. clients_trove	544
6.1.23. clients_zaqar	544
6.1.24. cors	545
6.1.25. database	545
6.1.26. ec2authtoken	547



---

6.1.27. eventlet_opts	548
6.1.28. healthcheck	548
6.1.29. heat_api	549
6.1.30. heat_api_cfn	549
6.1.31. heat_api_cloudwatch	550
6.1.32. keystone_auth token	551
6.1.33. matchmaker_redis	554
6.1.34. noauth	555
6.1.35. oslo_messaging_amqp	555
6.1.36. oslo_messaging_kafka	559
6.1.37. oslo_messaging_notifications	560
6.1.38. oslo_messaging_rabbit	560
6.1.39. oslo_messaging_zmq	564
6.1.40. oslo_middleware	567
6.1.41. oslo_policy	568
6.1.42. paste_deploy	569
6.1.43. profiler	569
6.1.44. revision	571
6.1.45. ssl	572
6.1.46. trustee	572
6.1.47. volumes	573
<b>CHAPTER 7. IRONIC</b> .....	<b>575</b>
7.1. IRONIC.CONF	575
7.1.1. DEFAULT	575
7.1.2. agent	591
7.1.3. ansible	592
7.1.4. api	594
7.1.5. audit	595
7.1.6. cimc	596
7.1.7. cinder	596
7.1.8. cisco_ucs	598
7.1.9. conductor	599
7.1.10. console	601
7.1.11. cors	602
7.1.12. database	603
7.1.13. deploy	604
7.1.14. dhcp	606
7.1.15. disk_partitioner	606
7.1.16. disk_utils	606
7.1.17. drac	607
7.1.18. glance	607
7.1.19. ilo	611
7.1.20. inspector	612
7.1.21. ipmi	615
7.1.22. irmc	615
7.1.23. ironic_lib	617
7.1.24. iscsi	617
7.1.25. keystone	617
7.1.26. keystone_auth token	618
7.1.27. matchmaker_redis	621
7.1.28. metrics	622
7.1.29. metrics_statsd	623

---

7.1.30. neutron	623
7.1.31. oneview	627
7.1.32. oslo_concurrency	627
7.1.33. oslo_messaging_amqp	628
7.1.34. oslo_messaging_kafka	631
7.1.35. oslo_messaging_notifications	632
7.1.36. oslo_messaging_rabbit	632
7.1.37. oslo_messaging_zmq	637
7.1.38. oslo_policy	640
7.1.39. profiler	640
7.1.40. pxe	642
7.1.41. service_catalog	644
7.1.42. snmp	646
7.1.43. ssl	647
7.1.44. swift	647
7.1.45. xclarity	649
<b>CHAPTER 8. IRONIC-INSPECTOR</b> .....	<b>651</b>
8.1. INSPECTOR.CONF	651
8.1.1. DEFAULT	651
8.1.2. capabilities	654
8.1.3. cors	655
8.1.4. database	656
8.1.5. discovery	657
8.1.6. dnsmasq_pxe_filter	657
8.1.7. iptables	658
8.1.8. ironic	659
8.1.9. keystone_authtoken	661
8.1.10. oslo_policy	665
8.1.11. pci_devices	666
8.1.12. processing	666
8.1.13. pxe_filter	668
8.1.14. swift	668
<b>CHAPTER 9. KEYSTONE</b> .....	<b>672</b>
9.1. KEYSTONE.CONF	672
9.1.1. DEFAULT	672
9.1.2. application_credential	681
9.1.3. assignment	682
9.1.4. auth	682
9.1.5. cache	684
9.1.6. catalog	685
9.1.7. cors	686
9.1.8. credential	687
9.1.9. database	688
9.1.10. domain_config	689
9.1.11. endpoint_filter	690
9.1.12. endpoint_policy	690
9.1.13. eventlet_server	690
9.1.14. federation	691
9.1.15. fernet_tokens	692
9.1.16. healthcheck	693
9.1.17. identity	694

---

9.1.18. identity_mapping	696
9.1.19. ldap	697
9.1.20. matchmaker_redis	703
9.1.21. memcache	704
9.1.22. oauth1	704
9.1.23. oslo_messaging_amqp	705
9.1.24. oslo_messaging_kafka	708
9.1.25. oslo_messaging_notifications	709
9.1.26. oslo_messaging_rabbit	710
9.1.27. oslo_messaging_zmq	714
9.1.28. oslo_middleware	717
9.1.29. oslo_policy	717
9.1.30. paste_deploy	718
9.1.31. policy	719
9.1.32. profiler	719
9.1.33. resource	721
9.1.34. revoke	723
9.1.35. role	723
9.1.36. saml	724
9.1.37. security_compliance	726
9.1.38. shadow_users	728
9.1.39. signing	728
9.1.40. token	730
9.1.41. tokenless_auth	732
9.1.42. trust	732
9.1.43. unified_limit	733
<b>CHAPTER 10. NEUTRON</b> .....	<b>734</b>
10.1. DHCP_AGENT.INI	734
10.1.1. DEFAULT	734
10.1.2. agent	739
10.1.3. ovs	739
10.2. L3_AGENT.INI	740
10.2.1. DEFAULT	740
10.2.2. agent	747
10.2.3. ovs	747
10.3. LINUXBRIDGE_AGENT.INI	748
10.3.1. DEFAULT	748
10.3.2. agent	751
10.3.3. linux_bridge	752
10.3.4. network_log	752
10.3.5. securitygroup	753
10.3.6. vxlan	753
10.4. METADATA_AGENT.INI	754
10.4.1. DEFAULT	755
10.4.2. agent	759
10.4.3. cache	759
10.5. METERING_AGENT.INI	760
10.5.1. DEFAULT	760
10.5.2. agent	764
10.5.3. ovs	764
10.6. ML2_CONF.INI	765
10.6.1. DEFAULT	765

---

10.6.2. l2pop	768
10.6.3. ml2	769
10.6.4. ml2_type_flat	770
10.6.5. ml2_type_geneve	770
10.6.6. ml2_type_gre	771
10.6.7. ml2_type_vlan	771
10.6.8. ml2_type_vxlan	771
10.6.9. securitygroup	771
10.7. NEUTRON.CONF	772
10.7.1. DEFAULT	772
10.7.2. agent	785
10.7.3. cors	786
10.7.4. database	787
10.7.5. keystone_auth token	788
10.7.6. matchmaker_redis	792
10.7.7. nova	792
10.7.8. oslo_concurrency	794
10.7.9. oslo_messaging_amqp	794
10.7.10. oslo_messaging_kafka	798
10.7.11. oslo_messaging_notifications	799
10.7.12. oslo_messaging_rabbit	799
10.7.13. oslo_messaging_zmq	803
10.7.14. oslo_middleware	806
10.7.15. oslo_policy	806
10.7.16. quotas	807
10.7.17. ssl	808
10.8. OPENVSWITCH_AGENT.INI	809
10.8.1. DEFAULT	809
10.8.2. agent	812
10.8.3. network_log	813
10.8.4. ovs	814
10.8.5. securitygroup	816
10.8.6. xenapi	816
10.9. SRIOV_AGENT.INI	817
10.9.1. DEFAULT	817
10.9.2. agent	820
10.9.3. sriov_nic	820
<b>CHAPTER 11. NOVA</b> .....	<b>821</b>
11.1. NOVA.CONF	821
11.1.1. DEFAULT	821
11.1.2. api	886
11.1.3. api_database	892
11.1.4. barbican	893
11.1.5. cache	894
11.1.6. cells	895
11.1.7. cinder	904
11.1.8. compute	907
11.1.9. conductor	909
11.1.10. console	910
11.1.11. consoleauth	910
11.1.12. cors	910
11.1.13. crypto	911

11.1.14. database	913
11.1.15. devices	914
11.1.16. ephemeral_storage_encryption	915
11.1.17. filter_scheduler	915
11.1.18. glance	928
11.1.19. guestfs	931
11.1.20. healthcheck	932
11.1.21. hyperv	932
11.1.22. ironic	940
11.1.23. key_manager	942
11.1.24. keystone	944
11.1.25. keystone_authtoken	945
11.1.26. libvirt	948
11.1.27. matchmaker_redis	973
11.1.28. metrics	973
11.1.29. mks	977
11.1.30. neutron	978
11.1.31. notifications	982
11.1.32. osapi_v21	984
11.1.33. oslo_concurrency	984
11.1.34. oslo_messaging_amqp	985
11.1.35. oslo_messaging_kafka	988
11.1.36. oslo_messaging_notifications	989
11.1.37. oslo_messaging_rabbit	989
11.1.38. oslo_messaging_zmq	994
11.1.39. oslo_middleware	997
11.1.40. oslo_policy	997
11.1.41. pci	998
11.1.42. placement	1001
11.1.43. quota	1004
11.1.44. rdp	1008
11.1.45. remote_debug	1010
11.1.46. scheduler	1011
11.1.47. serial_console	1014
11.1.48. service_user	1016
11.1.49. spice	1018
11.1.50. upgrade_levels	1021
11.1.51. vault	1022
11.1.52. vendordata_dynamic_auth	1022
11.1.53. vmware	1024
11.1.54. vnc	1028
11.1.55. workarounds	1033
11.1.56. wsgi	1037
11.1.57. xenserver	1041
11.1.58. xvp	1051
<b>CHAPTER 12. SAHARA</b> .....	<b>1053</b>
12.1. SAHARA.CONF	1053
12.1.1. DEFAULT	1053
12.1.2. castellan	1064
12.1.3. cinder	1064
12.1.4. conductor	1064
12.1.5. cors	1065

12.1.6. database	1065
12.1.7. glance	1067
12.1.8. heat	1067
12.1.9. keystone	1068
12.1.10. keystone_auth token	1068
12.1.11. manila	1072
12.1.12. matchmaker_redis	1072
12.1.13. neutron	1073
12.1.14. nova	1073
12.1.15. object_store_access	1074
12.1.16. oslo_concurrency	1074
12.1.17. oslo_messaging_amqp	1074
12.1.18. oslo_messaging_kafka	1078
12.1.19. oslo_messaging_notifications	1079
12.1.20. oslo_messaging_rabbit	1079
12.1.21. oslo_messaging_zmq	1083
12.1.22. oslo_middleware	1086
12.1.23. oslo_policy	1087
12.1.24. retries	1087
12.1.25. ssl	1088
12.1.26. swift	1088
12.1.27. timeouts	1089
12.1.28. trustee	1089



## PREFACE

This document describes the options available in the configuration files for each of the major services in Red Hat OpenStack Platform. The content is automatically generated based on the values in the configuration files themselves, and is provided for reference purposes only.



### **WARNING**

Manually editing configuration files is not supported. All configuration changes must be made through the Director. Red Hat provides this guide as a technical reference only.



# CHAPTER 1. BARBICAN

The following chapter contains information about the configuration options in the **barbican** service.

## 1.1. BARBICAN.CONF

This section contains options for the `/etc/barbican/barbican.conf` file.

### 1.1.1. DEFAULT

The following table outlines the options available under the **[DEFAULT]** group in the `/etc/barbican/barbican.conf` file.

Table 1.1. DEFAULT

Configuration option = Default value	Type	Description
<b>admin_role = admin</b>	string value	Role used to identify an authenticated user as administrator.
<b>allow_anonymous_access = False</b>	boolean value	Allow unauthenticated users to access the API with read-only privileges. This only applies when using ContextMiddleware.
<b>api_paste_config = api-paste.ini</b>	string value	File name for the paste.deploy config for api service
<b>backdoor_port = None</b>	string value	Enable eventlet backdoor. Acceptable values are 0, <port>, and <start>:<end>, where 0 results in listening on a random tcp port number; <port> results in listening on the specified port number (and not enabling backdoor if that port is in use); and <start>:<end> results in listening on the smallest unused port number within the specified range of port numbers. The chosen port is displayed in the service's log file.
<b>backdoor_socket = None</b>	string value	Enable eventlet backdoor, using the provided path as a unix socket that can receive connections. This option is mutually exclusive with <code>backdoor_port</code> in that only one should be provided. If both are provided then the existence of this option overrides the usage of that option.
<b>client_socket_timeout = 900</b>	integer value	Timeout for client connections' socket operations. If an incoming connection is idle for this number of seconds it will be closed. A value of 0 means wait forever.
<b>conn_pool_min_size = 2</b>	integer value	The pool size limit for connections expiration policy

Configuration option = Default value	Type	Description
<b>conn_pool_ttl = 1200</b>	integer value	The time-to-live in sec of idle connections in the pool
<b>control_exchange = openstack</b>	string value	The default exchange under which topics are scoped. May be overridden by an exchange name specified in the transport_url option.
<b>db_auto_create = True</b>	boolean value	Create the Barbican database on service startup.
<b>debug = False</b>	boolean value	If set to true, the logging level will be set to DEBUG instead of the default INFO level.
<b>default_limit_paging = 10</b>	integer value	Default page size for the <i>limit</i> paging URL parameter.
<b>default_log_levels = [u'amqp=WARN', u'amqplib=WARN', u'boto=WARN', u'qpid=WARN', u'sqlalchemy=WARN', u'suds=INFO', u'oslo.messaging=INFO', u'oslo_messaging=INFO', u'iso8601=WARN', u'requests.packages.urlli b3.connectionpool=WAR N', u'urllib3.connectionpool= WARN', u'websocket=WARN', u'requests.packages.urlli b3.util.retry=WARN', u'urllib3.util.retry=WARN', u'keystonemiddleware=W ARN', u'routes.middleware=WA RN', u'stevedore=WARN', u'taskflow=WARN', u'keystoneauth=WARN', u'oslo.cache=INFO', u'dogpile.core.dogpile=IN FO']</b>	list value	List of package logging levels in logger=LEVEL pairs. This option is ignored if log_config_append is set.
<b>executor_thread_pool_siz e = 64</b>	integer value	Size of executor thread pool when executor is threading or eventlet.
<b>fatal_deprecations = False</b>	boolean value	Enables or disables fatal status of deprecations.

Configuration option = Default value	Type	Description
<b>host_href =</b> <b>http://localhost:9311</b>	string value	Host name, for use in HATEOAS-style references Note: Typically this would be the load balanced endpoint that clients would use to communicate back with this service. If a deployment wants to derive host from wsgi request instead then make this blank. Blank is needed to override default config value which is <a href="http://localhost:9311">http://localhost:9311</a>
<code>`instance_format = [instance: %(uuid)s] `</code>	string value	The format for an instance that is passed with the log message.
<code>`instance_uuid_format = [instance: %(uuid)s] `</code>	string value	The format for an instance UUID that is passed with the log message.
<b>log-config-append = None</b>	string value	The name of a logging configuration file. This file is appended to any existing logging configuration files. For details about logging configuration files, see the Python logging module documentation. Note that when logging configuration files are used then all logging configuration is set in the configuration file and other logging configuration options are ignored (for example, <code>logging_context_format_string</code> ).
<b>log-date-format = %Y-%m-%d %H:%M:%S</b>	string value	Defines the format string for <code>%(asctime)s</code> in log records. Default: <code>%(default)s</code> . This option is ignored if <code>log_config_append</code> is set.
<b>log-dir = None</b>	string value	(Optional) The base directory used for relative <code>log_file</code> paths. This option is ignored if <code>log_config_append</code> is set.
<b>log-file = None</b>	string value	(Optional) Name of log file to send logging output to. If no default is set, logging will go to <code>stderr</code> as defined by <code>use_stderr</code> . This option is ignored if <code>log_config_append</code> is set.
<b>logging_context_format_string = %(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [%%(request_id)s %(user_identity)s] %(instance)s%(message)s</b>	string value	Format string to use for log messages with context.
<b>logging_debug_format_s uffix = %(funcName)s % (pathname)s:%(lineno)d</b>	string value	Additional data to append to log message when logging level for the message is <code>DEBUG</code> .

Configuration option = Default value	Type	Description
<b>logging_default_format_string =</b> <code>%(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [-] %(instance)s%(message)s</code>	string value	Format string to use for log messages when context is undefined.
<b>logging_exception_prefix =</b> <code>%(asctime)s.%(msecs)03d %(process)d ERROR %(name)s %(instance)s</code>	string value	Prefix each line of exception output with this format.
<b>logging_user_identity_format =</b> <code>%(user)s %(tenant)s %(domain)s %(user_domain)s %(project_domain)s</code>	string value	Defines the format string for <code>%(user_identity)s</code> that is used in <code>logging_context_format_string</code> .
<b>max_allowed_request_size_in_bytes =</b> <code>15000</code>	integer value	Maximum allowed http request size against the barbican-api.
<b>max_allowed_secret_in_bytes =</b> <code>10000</code>	integer value	Maximum allowed secret size in bytes.
<b>max_header_line =</b> <code>16384</code>	integer value	Maximum line size of message headers to be accepted. <code>max_header_line</code> may need to be increased when using large tokens (typically those generated when keystone is configured to use PKI tokens with big service catalogs).
<b>max_limit_paging =</b> <code>100</code>	integer value	Maximum page size for the <i>limit</i> paging URL parameter.
<b>publish_errors =</b> <code>False</code>	boolean value	Enables or disables publication of error events.
<b>rate_limit_burst =</b> <code>0</code>	integer value	Maximum number of logged messages per <code>rate_limit_interval</code> .
<b>rate_limit_except_level =</b> <code>CRITICAL</code>	string value	Log level name used by rate limiting: CRITICAL, ERROR, INFO, WARNING, DEBUG or empty string. Logs with level greater or equal to <code>rate_limit_except_level</code> are not filtered. An empty string means that all levels are filtered.
<b>rate_limit_interval =</b> <code>0</code>	integer value	Interval, number of seconds, of log rate limiting.

Configuration option = Default value	Type	Description
<b>rpc_ack_timeout_base = 15</b>	integer value	Number of seconds to wait for an ack from a cast/call. After each retry attempt this timeout is multiplied by some specified multiplier.
<b>rpc_ack_timeout_multiplier = 2</b>	integer value	Number to multiply base ack timeout by after each retry attempt.
<b>rpc_backend = rabbit</b>	string value	The messaging driver to use, defaults to rabbit. Other drivers include amqp and zmq.
<b>rpc_conn_pool_size = 30</b>	integer value	Size of RPC connection pool.
<b>rpc_message_ttl = 300</b>	integer value	Expiration timeout in seconds of a sent/received message after which it is not tracked anymore by a client/server.
<b>rpc_poll_timeout = 1</b>	integer value	The default number of seconds that poll should wait. Poll raises timeout exception when timeout expired.
<b>rpc_response_timeout = 60</b>	integer value	Seconds to wait for a response from a call.
<b>rpc_retry_attempts = 3</b>	integer value	Default number of message sending attempts in case of any problems occurred: positive value N means at most N retries, 0 means no retries, None or -1 (or any other negative values) mean to retry forever. This option is used only if acknowledgments are enabled.
<b>rpc_thread_pool_size = 100</b>	integer value	Maximum number of (green) threads to work concurrently.
<b>rpc_use_acks = False</b>	boolean value	Wait for message acknowledgements from receivers. This mechanism works only via proxy without PUB/SUB.
<b>rpc_zmq_bind_address = *</b>	string value	ZeroMQ bind address. Should be a wildcard (*), an ethernet interface, or IP. The "host" option should point or resolve to this address.
<b>rpc_zmq_bind_port_retries = 100</b>	integer value	Number of retries to find free port number before fail with ZMQBindError.
<b>rpc_zmq_contexts = 1</b>	integer value	Number of ZeroMQ contexts, defaults to 1.

Configuration option = Default value	Type	Description
<b>rpc_zmq_host = ccsosp-undercloud.usersys.redhat.com</b>	string value	Name of this node. Must be a valid hostname, FQDN, or IP address. Must match "host" option, if running Nova.
<b>rpc_zmq_ipc_dir = /var/run/openstack</b>	string value	Directory for holding IPC sockets.
<b>rpc_zmq_matchmaker = redis</b>	string value	MatchMaker driver.
<b>rpc_zmq_max_port = 65536</b>	integer value	Maximal port number for random ports range.
<b>rpc_zmq_min_port = 49153</b>	port value	Minimal port number for random ports range.
<b>rpc_zmq_serialization = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>rpc_zmq_topic_backlog = None</b>	integer value	Maximum number of ingress messages to locally buffer per topic. Default is unlimited.
<b>run_external_periodic_tasks = True</b>	boolean value	Some periodic tasks can be run in a separate process. Should we run them here?
<b>sql_connection = sqlite:///barbican.sqlite</b>	string value	SQLAlchemy connection string for the reference implementation registry server. Any valid SQLAlchemy connection string is fine. See: <a href="http://www.sqlalchemy.org/docs/05/reference/sqlalchemy/connections.html#sqlalchemy.create_engine">http://www.sqlalchemy.org/docs/05/reference/sqlalchemy/connections.html#sqlalchemy.create_engine</a> . Note: For absolute addresses, use <code>////</code> slashes after <code>sqlite:</code> .
<b>sql_idle_timeout = 3600</b>	integer value	Period in seconds after which SQLAlchemy should reestablish its connection to the database. MySQL uses a default <b>wait_timeout</b> of 8 hours, after which it will drop idle connections. This can result in <i>MySQL Gone Away</i> exceptions. If you notice this, you can lower this value to ensure that SQLAlchemy reconnects before MySQL can drop the connection.
<b>sql_max_retries = 60</b>	integer value	Maximum number of database connection retries during startup. Set to -1 to specify an infinite retry count.

Configuration option = Default value	Type	Description
<b>sql_pool_class = QueuePool</b>	string value	Accepts a class imported from the sqlalchemy.pool module, and handles the details of building the pool for you. If commented out, SQLAlchemy will select based on the database dialect. Other options are QueuePool (for SQLAlchemy-managed connections) and NullPool (to disabled SQLAlchemy management of connections). See <a href="http://docs.sqlalchemy.org/en/latest/core/pooling.html">http://docs.sqlalchemy.org/en/latest/core/pooling.html</a> for more details
<b>sql_pool_logging = False</b>	boolean value	Show SQLAlchemy pool-related debugging output in logs (sets DEBUG log level output) if specified.
<b>sql_pool_max_overflow = 10</b>	integer value	The maximum overflow size of the pool used by SQLAlchemy. When the number of checked-out connections reaches the size set in sql_pool_size, additional connections will be returned up to this limit. It follows then that the total number of simultaneous connections the pool will allow is sql_pool_size + sql_pool_max_overflow. Can be set to -1 to indicate no overflow limit, so no limit will be placed on the total number of concurrent connections. Comment out to allow SQLAlchemy to select the default.
<b>sql_pool_size = 5</b>	integer value	Size of pool used by SQLAlchemy. This is the largest number of connections that will be kept persistently in the pool. Can be set to 0 to indicate no size limit. To disable pooling, use a NullPool with sql_pool_class instead. Comment out to allow SQLAlchemy to select the default.
<b>sql_retry_interval = 1</b>	integer value	Interval between retries of opening a SQL connection.
<b>subscribe_on = []</b>	list value	List of publisher hosts SubConsumer can subscribe on. This option has higher priority then the default publishers list taken from the matchmaker.
<b>syslog-log-facility = LOG_USER</b>	string value	Syslog facility to receive log lines. This option is ignored if log_config_append is set.
<b>tcp_keepidle = 600</b>	integer value	Sets the value of TCP_KEEPIDLE in seconds for each server socket. Not supported on OS X.

Configuration option = Default value	Type	Description
<b>transport_url = None</b>	string value	<p>The network address and optional user credentials for connecting to the messaging backend, in URL format. The expected format is:</p> <pre>driver://[user:pass@]host:port[, [userN:passN@]hostN:portN]/virtual_host?query</pre> <p>Example: rabbit://rabbitmq:password@127.0.0.1:5672//</p> <p>For full details on the fields in the URL see the documentation of oslo_messaging.TransportURL at <a href="https://docs.openstack.org/oslo.messaging/latest/reference/transport.html">https://docs.openstack.org/oslo.messaging/latest/reference/transport.html</a></p>
<b>use-journal = False</b>	boolean value	Enable journald for logging. If running in a systemd environment you may wish to enable journal support. Doing so will use the journal native protocol which includes structured metadata in addition to log messages. This option is ignored if log_config_append is set.
<b>use-json = False</b>	boolean value	Use JSON formatting for logging. This option is ignored if log_config_append is set.
<b>use-syslog = False</b>	boolean value	Use syslog for logging. Existing syslog format is DEPRECATED and will be changed later to honor RFC5424. This option is ignored if log_config_append is set.
<b>use_dynamic_connections = False</b>	boolean value	This option makes direct connections dynamic or static. It makes sense only with use_router_proxy=False which means to use direct connections for direct message types (ignored otherwise).
<b>use_pub_sub = False</b>	boolean value	Use PUB/SUB pattern for fanout methods. PUB/SUB always uses proxy.
<b>use_router_proxy = False</b>	boolean value	Use ROUTER remote proxy.
<b>use_stderr = False</b>	boolean value	Log output to standard error. This option is ignored if log_config_append is set.



Configuration option = Default value	Type	Description
<b>watch-log-file = False</b>	boolean value	Uses logging handler designed to watch file system. When log file is moved or removed this handler will open a new log file with specified path instantaneously. It makes sense only if log_file option is specified and Linux platform is used. This option is ignored if log_config_append is set.
<b>wsgi_default_pool_size = 100</b>	integer value	Size of the pool of greenthreads used by wsgi
<b>wsgi_keep_alive = True</b>	boolean value	If False, closes the client socket connection explicitly.
<b>wsgi_log_format = % (client_ip)s "% (request_line)s" status: % (status_code)s len: % (body_length)s time: % (wall_seconds).7f</b>	string value	A python format string that is used as the template to generate log lines. The following values can be formatted into it: client_ip, date_time, request_line, status_code, body_length, wall_seconds.
<b>zmq_failover_connection s = 2</b>	integer value	How many additional connections to a host will be made for failover reasons. This option is actual only in dynamic connections mode.
<b>zmq_immediate = True</b>	boolean value	This option configures round-robin mode in zmq socket. True means not keeping a queue when server side disconnects. False means to keep queue and messages even if server is disconnected, when the server appears we send all accumulated messages to it.
<b>zmq_linger = -1</b>	integer value	Number of seconds to wait before all pending messages will be sent after closing a socket. The default value of -1 specifies an infinite linger period. The value of 0 specifies no linger period. Pending messages shall be discarded immediately when the socket is closed. Positive values specify an upper bound for the linger period.
<b>zmq_target_expire = 300</b>	integer value	Expiration timeout in seconds of a name service record about existing target (< 0 means no timeout).
<b>zmq_target_update = 180</b>	integer value	Update period in seconds of a name service record about existing target.

Configuration option = Default value	Type	Description
<b>zmq_tcp_keepalive = -1</b>	integer value	Enable/disable TCP keepalive (KA) mechanism. The default value of -1 (or any other negative value) means to skip any overrides and leave it to OS default; 0 and 1 (or any other positive value) mean to disable and enable the option respectively.
<b>zmq_tcp_keepalive_cnt = -1</b>	integer value	The number of retransmissions to be carried out before declaring that remote end is not available. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_idle = -1</b>	integer value	The duration between two keepalive transmissions in idle condition. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_intvl = -1</b>	integer value	The duration between two successive keepalive retransmissions, if acknowledgement to the previous keepalive transmission is not received. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.

### 1.1.2. certificate

The following table outlines the options available under the **[certificate]** group in the **/etc/barbican/barbican.conf** file.

Table 1.2. certificate

Configuration option = Default value	Type	Description
<b>enabled_certificate_plugins = [u'simple_certificate']</b>	multi valued	List of certificate plugins to load.
<b>namespace = barbican.certificate.plugin</b>	string value	Extension namespace to search for plugins.

### 1.1.3. certificate\_event

The following table outlines the options available under the **[certificate\_event]** group in the **/etc/barbican/barbican.conf** file.

Table 1.3. certificate\_event

Configuration option = Default value	Type	Description
<b>enabled_certificate_event_plugins = [u'simple_certificate_event']</b>	multi valued	List of certificate plugins to load.
<b>namespace = barbican.certificate.event.plugin</b>	string value	Extension namespace to search for eventing plugins.

#### 1.1.4. cors

The following table outlines the options available under the **[cors]** group in the **/etc/barbican/barbican.conf** file.

Table 1.4. cors

Configuration option = Default value	Type	Description
<b>allow_credentials = True</b>	boolean value	Indicate that the actual request can include user credentials
<b>allow_headers = [u'X-Auth-Token', u'X-Openstack-Request-Id', u'X-Project-Id', u'X-Identity-Status', u'X-User-Id', u'X-Storage-Token', u'X-Domain-Id', u'X-User-Domain-Id', u'X-Project-Domain-Id', u'X-Roles']</b>	list value	Indicate which header field names may be used during the actual request.
<b>allow_methods = [u'GET', u'PUT', u'POST', u'DELETE', u'PATCH']</b>	list value	Indicate which methods can be used during the actual request.
<b>allowed_origin = None</b>	list value	Indicate whether this resource may be shared with the domain received in the requests "origin" header. Format: "<protocol>://<host>[:<port>]", no trailing slash. Example: <a href="https://horizon.example.com">https://horizon.example.com</a>

Configuration option = Default value	Type	Description
<b>expose_headers = [u'X-Auth-Token', u'X-Openstack-Request-Id', u'X-Project-Id', u'X-Identity-Status', u'X-UserId', u'X-Storage-Token', u'X-Domain-Id', u'X-User-Domain-Id', u'X-Project-Domain-Id', u'X-Roles']</b>	list value	Indicate which headers are safe to expose to the API. Defaults to HTTP Simple Headers.
<b>max_age = 3600</b>	integer value	Maximum cache age of CORS preflight requests.

### 1.1.5. crypto

The following table outlines the options available under the **[crypto]** group in the `/etc/barbican/barbican.conf` file.

Table 1.5. crypto

Configuration option = Default value	Type	Description
<b>enabled_crypto_plugins = [u'simple_crypto']</b>	multi valued	List of crypto plugins to load.
<b>namespace = barbican.crypto.plugin</b>	string value	Extension namespace to search for plugins.

### 1.1.6. dogtag\_plugin

The following table outlines the options available under the **[dogtag\_plugin]** group in the `/etc/barbican/barbican.conf` file.

Table 1.6. dogtag\_plugin

Configuration option = Default value	Type	Description
<b>auto_approved_profiles = caServerCert</b>	string value	List of automatically approved enrollment profiles
<b>ca_expiration_time = 1</b>	string value	Time in days for CA entries to expire
<b>dogtag_host = localhost</b>	string value	Hostname for the Dogtag instance
<b>dogtag_port = 8443</b>	port value	Port for the Dogtag instance

Configuration option = Default value	Type	Description
<b>nss_db_path = /etc/barbican/alias</b>	string value	Path to the NSS certificate database
<b>nss_password = None</b>	string value	Password for the NSS certificate databases
<b>pem_path = /etc/barbican/kra_admin_ cert.pem</b>	string value	Path to PEM file for authentication
<b>plugin_name = Dogtag KRA</b>	string value	User friendly plugin name
<b>plugin_working_dir = /etc/barbican/dogtag</b>	string value	Working directory for Dogtag plugin
<b>retries = 3</b>	integer value	Retries when storing or generating secrets
<b>simple_cmc_profile = caOtherCert</b>	string value	Profile for simple CMC requests

### 1.1.7. keystone\_authtoken

The following table outlines the options available under the **[keystone\_authtoken]** group in the **/etc/barbican/barbican.conf** file.

Table 1.7. keystone\_authtoken

Configuration option = Default value	Type	Description
<b>auth_section = None</b>	string value	Config Section from which to load plugin specific options
<b>auth_type = None</b>	string value	Authentication type to load

Configuration option = Default value	Type	Description
<b>auth_uri = None</b>	string value	Complete "public" Identity API endpoint. This endpoint should not be an "admin" endpoint, as it should be accessible by all end users. Unauthenticated clients are redirected to this endpoint to authenticate. Although this endpoint should ideally be unversioned, client support in the wild varies. If you're using a versioned v2 endpoint here, then this should <b>not</b> be the same endpoint the service user utilizes for validating tokens, because normal end users may not be able to reach that endpoint. This option is deprecated in favor of <code>www_authenticate_uri</code> and will be removed in the S release.
<b>auth_version = None</b>	string value	API version of the admin Identity API endpoint.
<b>cache = None</b>	string value	Request environment key where the Swift cache object is stored. When <code>auth_token</code> middleware is deployed with a Swift cache, use this option to have the middleware share a caching backend with swift. Otherwise, use the <b><code>memcached_servers</code></b> option instead.
<b>cafile = None</b>	string value	A PEM encoded Certificate Authority to use when verifying HTTPs connections. Defaults to system CAs.
<b>certfile = None</b>	string value	Required if identity server requires client certificate
<b>check_revocations_for_cached = False</b>	boolean value	If true, the revocation list will be checked for cached tokens. This requires that PKI tokens are configured on the identity server.
<b>delay_auth_decision = False</b>	boolean value	Do not handle authorization requests within the middleware, but delegate the authorization decision to downstream WSGI components.
<b>enforce_token_bind = permissive</b>	string value	Used to control the use and type of token binding. Can be set to: "disabled" to not check token binding. "permissive" (default) to validate binding information if the bind type is of a form known to the server and ignore it if not. "strict" like "permissive" but if the bind type is unknown the token will be rejected. "required" any form of token binding is needed to be allowed. Finally the name of a binding method that must be present in tokens.

Configuration option = Default value	Type	Description
<b>hash_algorithms = [u'md5']</b>	list value	Hash algorithms to use for hashing PKI tokens. This may be a single algorithm or multiple. The algorithms are those supported by Python standard <code>hashlib.new()</code> . The hashes will be tried in the order given, so put the preferred one first for performance. The result of the first hash will be stored in the cache. This will typically be set to multiple values only while migrating from a less secure algorithm to a more secure one. Once all the old tokens are expired this option should be set to a single value for better performance.
<b>http_connect_timeout = None</b>	integer value	Request timeout value for communicating with Identity API server.
<b>http_request_max_retries = 3</b>	integer value	How many times are we trying to reconnect when communicating with Identity API Server.
<b>include_service_catalog = True</b>	boolean value	(Optional) Indicate whether to set the X-Service-Catalog header. If False, middleware will not ask for service catalog on token validation and will not set the X-Service-Catalog header.
<b>insecure = False</b>	boolean value	Verify HTTPS connections.
<b>keyfile = None</b>	string value	Required if identity server requires client certificate
<b>memcache_pool_conn_g et_timeout = 10</b>	integer value	(Optional) Number of seconds that an operation will wait to get a memcached client connection from the pool.
<b>memcache_pool_dead_re try = 300</b>	integer value	(Optional) Number of seconds memcached server is considered dead before it is tried again.
<b>memcache_pool_maxsize = 10</b>	integer value	(Optional) Maximum total number of open connections to every memcached server.
<b>memcache_pool_socket_ timeout = 3</b>	integer value	(Optional) Socket timeout in seconds for communicating with a memcached server.
<b>memcache_pool_unused_ timeout = 60</b>	integer value	(Optional) Number of seconds a connection to memcached is held unused in the pool before it is closed.
<b>memcache_secret_key = None</b>	string value	(Optional, mandatory if <code>memcache_security_strategy</code> is defined) This string is used for key derivation.

Configuration option = Default value	Type	Description
<b>memcache_security_strategy = None</b>	string value	(Optional) If defined, indicate whether token data should be authenticated or authenticated and encrypted. If MAC, token data is authenticated (with HMAC) in the cache. If ENCRYPT, token data is encrypted and authenticated in the cache. If the value is not one of these options or empty, <code>auth_token</code> will raise an exception on initialization.
<b>memcache_use_advanced_pool = False</b>	boolean value	(Optional) Use the advanced (eventlet safe) memcached client pool. The advanced pool will only work under python 2.x.
<b>memcached_servers = None</b>	list value	Optionally specify a list of memcached server(s) to use for caching. If left undefined, tokens will instead be cached in-process.
<b>region_name = None</b>	string value	The region in which the identity server can be found.
<b>revocation_cache_time = 10</b>	integer value	Determines the frequency at which the list of revoked tokens is retrieved from the Identity service (in seconds). A high number of revocation events combined with a low cache duration may significantly reduce performance. Only valid for PKI tokens. This option has been deprecated in the Ocata release and will be removed in the P release.
<b>service_token_roles = [u'service']</b>	list value	A choice of roles that must be present in a service token. Service tokens are allowed to request that an expired token can be used and so this check should tightly control that only actual services should be sending this token. Roles here are applied as an ANY check so any role in this list must be present. For backwards compatibility reasons this currently only affects the <code>allow_expired</code> check.
<b>service_token_roles_required = False</b>	boolean value	For backwards compatibility reasons we must let valid service tokens pass that don't pass the <code>service_token_roles</code> check as valid. Setting this true will become the default in a future release and should be enabled if possible.
<b>signing_dir = None</b>	string value	Directory used to cache files related to PKI tokens. This option has been deprecated in the Ocata release and will be removed in the P release.



Configuration option = Default value	Type	Description
<b>token_cache_time = 300</b>	integer value	In order to prevent excessive effort spent validating tokens, the middleware caches previously-seen tokens for a configurable duration (in seconds). Set to -1 to disable caching completely.
<b>www_authenticate_uri = None</b>	string value	Complete "public" Identity API endpoint. This endpoint should not be an "admin" endpoint, as it should be accessible by all end users. Unauthenticated clients are redirected to this endpoint to authenticate. Although this endpoint should ideally be unversioned, client support in the wild varies. If you're using a versioned v2 endpoint here, then this should <b>not</b> be the same endpoint the service user utilizes for validating tokens, because normal end users may not be able to reach that endpoint.

### 1.1.8. keystone\_notifications

The following table outlines the options available under the **[keystone\_notifications]** group in the **/etc/barbican/barbican.conf** file.

Table 1.8. keystone\_notifications

Configuration option = Default value	Type	Description
<b>allow_requeue = False</b>	boolean value	True enables requeue feature in case of notification processing error. Enable this only when underlying transport supports this feature.
<b>control_exchange = openstack</b>	string value	The default exchange under which topics are scoped. May be overridden by an exchange name specified in the transport_url option.
<b>enable = False</b>	boolean value	True enables keystone notification listener functionality.
<b>thread_pool_size = 10</b>	integer value	Define the number of max threads to be used for notification server processing functionality.

Configuration option = Default value	Type	Description
<b>topic = notifications</b>	string value	Keystone notification queue topic name. This name needs to match one of values mentioned in Keystone deployment's <i>notification_topics</i> configuration e.g. <i>notification_topics=notifications</i> , <i>barbican_notifications</i> Multiple servers may listen on a topic and messages will be dispatched to one of the servers in a round-robin fashion. That's why Barbican service should have its own dedicated notification queue so that it receives all of Keystone notifications.
<b>version = 1.0</b>	string value	Version of tasks invoked via notifications

### 1.1.9. kmip\_plugin

The following table outlines the options available under the **[kmip\_plugin]** group in the `/etc/barbican/barbican.conf` file.

Table 1.9. kmip\_plugin

Configuration option = Default value	Type	Description
<b>ca_certs = None</b>	string value	File path to concatenated "certification authority" certificates
<b>certfile = None</b>	string value	File path to local client certificate
<b>host = localhost</b>	string value	Address of the KMIP server
<b>keyfile = None</b>	string value	File path to local client certificate keyfile
<b>password = None</b>	string value	Password for authenticating with KMIP server
<b>pkcs1_only = False</b>	boolean value	Only support PKCS#1 encoding of asymmetric keys
<b>plugin_name = KMIP HSM</b>	string value	User friendly plugin name
<b>port = 5696</b>	port value	Port for the KMIP server
<b>ssl_version = PROTOCOL_TLSv1_2</b>	string value	SSL version, maps to the module ssl's constants
<b>username = None</b>	string value	Username for authenticating with KMIP server

### 1.1.10. matchmaker\_redis

The following table outlines the options available under the **[matchmaker\_redis]** group in the **/etc/barbican/barbican.conf** file.

Table 1.10. matchmaker\_redis

Configuration option = Default value	Type	Description
<b>check_timeout = 20000</b>	integer value	Time in ms to wait before the transaction is killed.
<b>host = 127.0.0.1</b>	string value	Host to locate redis.
<code>`password = `</code>	string value	Password for Redis server (optional).
<b>port = 6379</b>	port value	Use this port to connect to redis host.
<b>sentinel_group_name = oslo-messaging-zeromq</b>	string value	Redis replica set name.
<b>sentinel_hosts = []</b>	list value	List of Redis Sentinel hosts (fault tolerance mode), e.g., [host:port, host1:port ... ]
<b>socket_timeout = 10000</b>	integer value	Timeout in ms on blocking socket operations.
<b>wait_timeout = 2000</b>	integer value	Time in ms to wait between connection attempts.

### 1.1.11. oslo\_messaging\_amqp

The following table outlines the options available under the **[oslo\_messaging\_amqp]** group in the **/etc/barbican/barbican.conf** file.

Table 1.11. oslo\_messaging\_amqp

Configuration option = Default value	Type	Description
<b>addressing_mode = dynamic</b>	string value	Indicates the addressing mode used by the driver. Permitted values: <i>legacy</i> - use legacy non-routable addressing <i>routable</i> - use routable addresses <i>dynamic</i> - use legacy addresses if the message bus does not support routing otherwise use routable addressing
<b>allow_insecure_clients = False</b>	boolean value	Accept clients using either SSL or plain TCP

Configuration option = Default value	Type	Description
<b>anycast_address = anycast</b>	string value	Appended to the address prefix when sending to a group of consumers. Used by the message bus to identify messages that should be delivered in a round-robin fashion across consumers.
<b>broadcast_prefix = broadcast</b>	string value	address prefix used when broadcasting to all servers
<b>connection_retry_backoff = 2</b>	integer value	Increase the connection_retry_interval by this many seconds after each unsuccessful failover attempt.
<b>connection_retry_interval = 1</b>	integer value	Seconds to pause before attempting to re-connect.
<b>connection_retry_interval_max = 30</b>	integer value	Maximum limit for connection_retry_interval + connection_retry_backoff
<b>container_name = None</b>	string value	Name for the AMQP container. must be globally unique. Defaults to a generated UUID
<b>default_notification_exchange = None</b>	string value	Exchange name used in notification addresses. Exchange name resolution precedence: Target.exchange if set else default_notification_exchange if set else control_exchange if set else <i>notify</i>
<b>default_notify_timeout = 30</b>	integer value	The deadline for a sent notification message delivery. Only used when caller does not provide a timeout expiry.
<b>default_reply_retry = 0</b>	integer value	The maximum number of attempts to re-send a reply message which failed due to a recoverable error.
<b>default_reply_timeout = 30</b>	integer value	The deadline for an rpc reply message delivery.
<b>default_rpc_exchange = None</b>	string value	Exchange name used in RPC addresses. Exchange name resolution precedence: Target.exchange if set else default_rpc_exchange if set else control_exchange if set else <i>rpc</i>
<b>default_send_timeout = 30</b>	integer value	The deadline for an rpc cast or call message delivery. Only used when caller does not provide a timeout expiry.
<b>default_sender_link_timeout = 600</b>	integer value	The duration to schedule a purge of idle sender links. Detach link after expiry.

Configuration option = Default value	Type	Description
<b>group_request_prefix = unicast</b>	string value	address prefix when sending to any server in group
<b>idle_timeout = 0</b>	integer value	Timeout for inactive connections (in seconds)
<b>link_retry_delay = 10</b>	integer value	Time to pause between re-connecting an AMQP 1.0 link that failed due to a recoverable error.
<b>multicast_address = multicast</b>	string value	Appended to the address prefix when sending a fanout message. Used by the message bus to identify fanout messages.
<b>notify_address_prefix = openstack.org/om/notify</b>	string value	Address prefix for all generated Notification addresses
<b>notify_server_credit = 100</b>	integer value	Window size for incoming Notification messages
<code>`password = `</code>	string value	Password for message broker authentication
<b>pre_settled = [u'rpc-cast', u'rpc-reply']</b>	multi valued	Send messages of this type pre-settled. Pre-settled messages will not receive acknowledgement from the peer. Note well: pre-settled messages may be silently discarded if the delivery fails. Permitted values: <i>rpc-call</i> - send RPC Calls pre-settled <i>rpc-reply</i> - send RPC Replies pre-settled <i>rpc-cast</i> - Send RPC Casts pre-settled <i>notify</i> - Send Notifications pre-settled
<b>pseudo_vhost = True</b>	boolean value	Enable virtual host support for those message buses that do not natively support virtual hosting (such as qpidd). When set to true the virtual host name will be added to all message bus addresses, effectively creating a private <i>subnet</i> per virtual host. Set to False if the message bus supports virtual hosting using the <i>hostname</i> field in the AMQP 1.0 Open performative as the name of the virtual host.
<b>reply_link_credit = 200</b>	integer value	Window size for incoming RPC Reply messages.
<b>rpc_address_prefix = openstack.org/om/rpc</b>	string value	Address prefix for all generated RPC addresses
<b>rpc_server_credit = 100</b>	integer value	Window size for incoming RPC Request messages
<code>`sas_config_dir = `</code>	string value	Path to directory that contains the SASL configuration
<code>`sas_config_name = `</code>	string value	Name of configuration file (without .conf suffix)

Configuration option = Default value	Type	Description
<code>`sasl_default_realm = `</code>	string value	SASL realm to use if no realm present in username
<code>`sasl_mechanisms = `</code>	string value	Space separated list of acceptable SASL mechanisms
<b>server_request_prefix = exclusive</b>	string value	address prefix used when sending to a specific server
<b>ssl = False</b>	boolean value	Attempt to connect via SSL. If no other ssl-related parameters are given, it will use the system's CA-bundle to verify the server's certificate.
<code>`ssl_ca_file = `</code>	string value	CA certificate PEM file used to verify the server's certificate
<code>`ssl_cert_file = `</code>	string value	Self-identifying certificate PEM file for client authentication
<code>`ssl_key_file = `</code>	string value	Private key PEM file used to sign <code>ssl_cert_file</code> certificate (optional)
<b>ssl_key_password = None</b>	string value	Password for decrypting <code>ssl_key_file</code> (if encrypted)
<b>ssl_verify_vhost = False</b>	boolean value	By default SSL checks that the name in the server's certificate matches the hostname in the <code>transport_url</code> . In some configurations it may be preferable to use the virtual hostname instead, for example if the server uses the Server Name Indication TLS extension (rfc6066) to provide a certificate per virtual host. Set <code>ssl_verify_vhost</code> to True if the server's SSL certificate uses the virtual host name instead of the DNS name.
<b>trace = False</b>	boolean value	Debug: dump AMQP frames to stdout
<b>unicast_address = unicast</b>	string value	Appended to the address prefix when sending to a particular RPC/Notification server. Used by the message bus to identify messages sent to a single destination.
<code>`username = `</code>	string value	User name for message broker authentication

### 1.1.12. oslo\_messaging\_kafka

The following table outlines the options available under the **[oslo\_messaging\_kafka]** group in the **/etc/barbican/barbican.conf** file.

Table 1.12. oslo\_messaging\_kafka

Configuration option = Default value	Type	Description
<b>conn_pool_min_size = 2</b>	integer value	The pool size limit for connections expiration policy
<b>conn_pool_ttl = 1200</b>	integer value	The time-to-live in sec of idle connections in the pool
<b>consumer_group = oslo_messaging_consumer</b>	string value	Group id for Kafka consumer. Consumers in one group will coordinate message consumption
<b>kafka_consumer_timeout = 1.0</b>	floating point value	Default timeout(s) for Kafka consumers
<b>kafka_default_host = localhost</b>	string value	Default Kafka broker Host
<b>kafka_default_port = 9092</b>	port value	Default Kafka broker Port
<b>kafka_max_fetch_bytes = 1048576</b>	integer value	Max fetch bytes of Kafka consumer
<b>pool_size = 10</b>	integer value	Pool Size for Kafka Consumers
<b>producer_batch_size = 16384</b>	integer value	Size of batch for the producer async send
<b>producer_batch_timeout = 0.0</b>	floating point value	Upper bound on the delay for KafkaProducer batching in seconds

### 1.1.13. oslo\_messaging\_notifications

The following table outlines the options available under the **[oslo\_messaging\_notifications]** group in the **/etc/barbican/barbican.conf** file.

Table 1.13. oslo\_messaging\_notifications

Configuration option = Default value	Type	Description
<b>driver = []</b>	multi valued	The Drivers(s) to handle sending notifications. Possible values are messaging, messagingv2, routing, log, test, noop

Configuration option = Default value	Type	Description
<b>retry = -1</b>	integer value	The maximum number of attempts to re-send a notification message which failed to be delivered due to a recoverable error. 0 - No retry, -1 - indefinite
<b>topics = [u'notifications']</b>	list value	AMQP topic used for OpenStack notifications.
<b>transport_url = None</b>	string value	A URL representing the messaging driver to use for notifications. If not set, we fall back to the same configuration used for RPC.

### 1.1.14. oslo\_messaging\_rabbit

The following table outlines the options available under the **[oslo\_messaging\_rabbit]** group in the **/etc/barbican/barbican.conf** file.

Table 1.14. oslo\_messaging\_rabbit

Configuration option = Default value	Type	Description
<b>amqp_auto_delete = False</b>	boolean value	Auto-delete queues in AMQP.
<b>amqp_durable_queues = False</b>	boolean value	Use durable queues in AMQP.
<b>channel_max = None</b>	integer value	Maximum number of channels to allow
<b>connection_factory = single</b>	string value	Connection factory implementation
<b>default_notification_exchange = \${control_exchange}_notification</b>	string value	Exchange name for sending notifications
<b>default_notification_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending notification, -1 means infinite retry.
<b>default_rpc_exchange = \${control_exchange}_rpc</b>	string value	Exchange name for sending RPC messages



Configuration option = Default value	Type	Description
<b>default_rpc_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending RPC message, -1 means infinite retry. If actual retry attempts in not 0 the rpc request could be processed more than one time
<b>default_serializer_type = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>fake_rabbit = False</b>	boolean value	Deprecated, use <code>rpc_backend=kombu+memory</code> or <code>rpc_backend=fake</code>
<b>frame_max = None</b>	integer value	The maximum byte size for an AMQP frame
<b>heartbeat_interval = 3</b>	integer value	How often to send heartbeats for consumer's connections
<b>heartbeat_rate = 2</b>	integer value	How often times during the <code>heartbeat_timeout_threshold</code> we check the heartbeat.
<b>heartbeat_timeout_threshold = 60</b>	integer value	Number of seconds after which the Rabbit broker is considered down if heartbeat's keep-alive fails (0 disable the heartbeat). EXPERIMENTAL
<b>host_connection_reconnect_delay = 0.25</b>	floating point value	Set delay for reconnection to some host which has connection error
<b>kombu_compression = None</b>	string value	EXPERIMENTAL: Possible values are: <code>gzip</code> , <code>bz2</code> . If not set compression will not be used. This option may not be available in future versions.
<b>kombu_failover_strategy = round-robin</b>	string value	Determines how the next RabbitMQ node is chosen in case the one we are currently connected to becomes unavailable. Takes effect only if more than one RabbitMQ node is provided in config.
<b>kombu_missing_consumer_retry_timeout = 60</b>	integer value	How long to wait a missing client before abandoning to send it its replies. This value should not be longer than <code>rpc_response_timeout</code> .
<b>kombu_reconnect_delay = 1.0</b>	floating point value	How long to wait before reconnecting in response to an AMQP consumer cancel notification.
<b>notification_listener_prefetch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to notification listener.

Configuration option = Default value	Type	Description
<b>notification_persistence = False</b>	boolean value	Persist notification messages.
<b>notification_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending notification message
<b>pool_max_overflow = 0</b>	integer value	Maximum number of connections to create above <b>pool_max_size</b> .
<b>pool_max_size = 30</b>	integer value	Maximum number of connections to keep queued.
<b>pool_recycle = 600</b>	integer value	Lifetime of a connection (since creation) in seconds or None for no recycling. Expired connections are closed on acquire.
<b>pool_stale = 60</b>	integer value	Threshold at which inactive (since release) connections are considered stale in seconds or None for no staleness. Stale connections are closed on acquire.
<b>pool_timeout = 30</b>	integer value	Default number of seconds to wait for a connections to available
<b>rabbit_ha_queues = False</b>	boolean value	Try to use HA queues in RabbitMQ (x-ha-policy: all). If you change this option, you must wipe the RabbitMQ database. In RabbitMQ 3.0, queue mirroring is no longer controlled by the x-ha-policy argument when declaring a queue. If you just want to make sure that all queues (except those with auto-generated names) are mirrored across all nodes, run: <code>"rabbitmqctl set_policy HA ^(?!amq\).* {ha-mode}:all"</code>
<b>rabbit_host = localhost</b>	string value	The RabbitMQ broker address where a single node is used.
<b>rabbit_hosts = [u'\$rabbit_host:\$rabbit_port']</b>	list value	RabbitMQ HA cluster host:port pairs.
<b>rabbit_interval_max = 30</b>	integer value	Maximum interval of RabbitMQ connection retries. Default is 30 seconds.
<b>rabbit_login_method = AMQPLAIN</b>	string value	The RabbitMQ login method.

Configuration option = Default value	Type	Description
<b>rabbit_max_retries = 0</b>	integer value	Maximum number of RabbitMQ connection retries. Default is 0 (infinite retry count).
<b>rabbit_password = guest</b>	string value	The RabbitMQ password.
<b>rabbit_port = 5672</b>	port value	The RabbitMQ broker port where a single node is used.
<b>rabbit_qos_prefetch_count = 0</b>	integer value	Specifies the number of messages to prefetch. Setting to zero allows unlimited messages.
<b>rabbit_retry_backoff = 2</b>	integer value	How long to backoff for between retries when connecting to RabbitMQ.
<b>rabbit_retry_interval = 1</b>	integer value	How frequently to retry connecting with RabbitMQ.
<b>rabbit_transient_queues_ttl = 1800</b>	integer value	Positive integer representing duration in seconds for queue TTL (x-expires). Queues which are unused for the duration of the TTL are automatically deleted. The parameter affects only reply and fanout queues.
<b>rabbit_userid = guest</b>	string value	The RabbitMQ userid.
<b>rabbit_virtual_host = /</b>	string value	The RabbitMQ virtual host.
<b>rpc_listener_prefetch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to rpc listener.
<b>rpc_queue_expiration = 60</b>	integer value	Time to live for rpc queues without consumers in seconds.
<b>rpc_reply_exchange = \${control_exchange}_rpc_reply</b>	string value	Exchange name for receiving RPC replies
<b>rpc_reply_listener_prefetch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to rpc reply listener.
<b>rpc_reply_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending reply. -1 means infinite retry during rpc_timeout
<b>rpc_reply_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending reply.

Configuration option = Default value	Type	Description
<b>rpc_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending RPC message
<b>socket_timeout = 0.25</b>	floating point value	Set socket timeout in seconds for connection's socket
<b>ssl = None</b>	boolean value	Enable SSL
<code>`ssl_ca_file = `</code>	string value	SSL certification authority file (valid only if SSL enabled).
<code>`ssl_cert_file = `</code>	string value	SSL cert file (valid only if SSL enabled).
<code>`ssl_key_file = `</code>	string value	SSL key file (valid only if SSL enabled).
<b>ssl_options = None</b>	dict value	Arguments passed to <code>ssl.wrap_socket</code>
<code>`ssl_version = `</code>	string value	SSL version to use (valid only if SSL enabled). Valid values are TLSv1 and SSLv23. SSLv2, SSLv3, TLSv1_1, and TLSv1_2 may be available on some distributions.
<b>tcp_user_timeout = 0.25</b>	floating point value	Set TCP_USER_TIMEOUT in seconds for connection's socket

### 1.1.15. oslo\_messaging\_zmq

The following table outlines the options available under the **[oslo\_messaging\_zmq]** group in the `/etc/barbican/barbican.conf` file.

Table 1.15. oslo\_messaging\_zmq

Configuration option = Default value	Type	Description
<b>rpc_ack_timeout_base = 15</b>	integer value	Number of seconds to wait for an ack from a cast/call. After each retry attempt this timeout is multiplied by some specified multiplier.
<b>rpc_ack_timeout_multiplier = 2</b>	integer value	Number to multiply base ack timeout by after each retry attempt.

Configuration option = Default value	Type	Description
<b>rpc_message_ttl = 300</b>	integer value	Expiration timeout in seconds of a sent/received message after which it is not tracked anymore by a client/server.
<b>rpc_poll_timeout = 1</b>	integer value	The default number of seconds that poll should wait. Poll raises timeout exception when timeout expired.
<b>rpc_retry_attempts = 3</b>	integer value	Default number of message sending attempts in case of any problems occurred: positive value N means at most N retries, 0 means no retries, None or -1 (or any other negative values) mean to retry forever. This option is used only if acknowledgments are enabled.
<b>rpc_thread_pool_size = 100</b>	integer value	Maximum number of (green) threads to work concurrently.
<b>rpc_use_acks = False</b>	boolean value	Wait for message acknowledgements from receivers. This mechanism works only via proxy without PUB/SUB.
<b>rpc_zmq_bind_address = *</b>	string value	ZeroMQ bind address. Should be a wildcard (*), an ethernet interface, or IP. The "host" option should point or resolve to this address.
<b>rpc_zmq_bind_port_retries = 100</b>	integer value	Number of retries to find free port number before fail with ZMQBindError.
<b>rpc_zmq_contexts = 1</b>	integer value	Number of ZeroMQ contexts, defaults to 1.
<b>rpc_zmq_host = ccso-sp-undercloud.usersys.redhat.com</b>	string value	Name of this node. Must be a valid hostname, FQDN, or IP address. Must match "host" option, if running Nova.
<b>rpc_zmq_ipc_dir = /var/run/openstack</b>	string value	Directory for holding IPC sockets.
<b>rpc_zmq_matchmaker = redis</b>	string value	MatchMaker driver.
<b>rpc_zmq_max_port = 65536</b>	integer value	Maximal port number for random ports range.
<b>rpc_zmq_min_port = 49153</b>	port value	Minimal port number for random ports range.

Configuration option = Default value	Type	Description
<b>rpc_zmq_serialization = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>rpc_zmq_topic_backlog = None</b>	integer value	Maximum number of ingress messages to locally buffer per topic. Default is unlimited.
<b>subscribe_on = []</b>	list value	List of publisher hosts SubConsumer can subscribe on. This option has higher priority then the default publishers list taken from the matchmaker.
<b>use_dynamic_connections = False</b>	boolean value	This option makes direct connections dynamic or static. It makes sense only with <code>use_router_proxy=False</code> which means to use direct connections for direct message types (ignored otherwise).
<b>use_pub_sub = False</b>	boolean value	Use PUB/SUB pattern for fanout methods. PUB/SUB always uses proxy.
<b>use_router_proxy = False</b>	boolean value	Use ROUTER remote proxy.
<b>zmq_failover_connections = 2</b>	integer value	How many additional connections to a host will be made for failover reasons. This option is actual only in dynamic connections mode.
<b>zmq_immediate = True</b>	boolean value	This option configures round-robin mode in zmq socket. True means not keeping a queue when server side disconnects. False means to keep queue and messages even if server is disconnected, when the server appears we send all accumulated messages to it.
<b>zmq_linger = -1</b>	integer value	Number of seconds to wait before all pending messages will be sent after closing a socket. The default value of -1 specifies an infinite linger period. The value of 0 specifies no linger period. Pending messages shall be discarded immediately when the socket is closed. Positive values specify an upper bound for the linger period.
<b>zmq_target_expire = 300</b>	integer value	Expiration timeout in seconds of a name service record about existing target (< 0 means no timeout).

Configuration option = Default value	Type	Description
<b>zmq_target_update = 180</b>	integer value	Update period in seconds of a name service record about existing target.
<b>zmq_tcp_keepalive = -1</b>	integer value	Enable/disable TCP keepalive (KA) mechanism. The default value of -1 (or any other negative value) means to skip any overrides and leave it to OS default; 0 and 1 (or any other positive value) mean to disable and enable the option respectively.
<b>zmq_tcp_keepalive_cnt = -1</b>	integer value	The number of retransmissions to be carried out before declaring that remote end is not available. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_idle = -1</b>	integer value	The duration between two keepalive transmissions in idle condition. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_intvl = -1</b>	integer value	The duration between two successive keepalive retransmissions, if acknowledgement to the previous keepalive transmission is not received. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.

### 1.1.16. oslo\_middlewares

The following table outlines the options available under the **[oslo\_middlewares]** group in the **/etc/barbican/barbican.conf** file.

Table 1.16. oslo\_middlewares

Configuration option = Default value	Type	Description
<b>enable_proxy_headers_parsing = False</b>	boolean value	Whether the application is behind a proxy or not. This determines if the middleware should parse the headers or not.

### 1.1.17. oslo\_policy

The following table outlines the options available under the **[oslo\_policy]** group in the **/etc/barbican/barbican.conf** file.

Table 1.17. oslo\_policy

Configuration option = Default value	Type	Description
<b>enforce_scope = False</b>	boolean value	This option controls whether or not to enforce scope when evaluating policies. If <b>True</b> , the scope of the token used in the request is compared to the <b>scope_types</b> of the policy being enforced. If the scopes do not match, an <b>InvalidScope</b> exception will be raised. If <b>False</b> , a message will be logged informing operators that policies are being invoked with mismatching scope.
<b>policy_default_rule = default</b>	string value	Default rule. Enforced when a requested rule is not found.
<b>policy_dirs = [u'policy.d']</b>	multi valued	Directories where policy configuration files are stored. They can be relative to any directory in the search path defined by the config_dir option, or absolute paths. The file defined by policy_file must exist for these directories to be searched. Missing or empty directories are ignored.
<b>policy_file = policy.json</b>	string value	The file that defines policies.
<b>remote_content_type = application/x-www-form-urlencoded</b>	string value	Content Type to send and receive data for REST based policy check
<b>remote_ssl_ca_cert_file = None</b>	string value	Absolute path to ca cert file for REST based policy check
<b>remote_ssl_client_cert_file = None</b>	string value	Absolute path to client cert for REST based policy check
<b>remote_ssl_client_key_file = None</b>	string value	Absolute path client key file REST based policy check
<b>remote_ssl_verify_server_cert = False</b>	boolean value	server identity verification for REST based policy check

### 1.1.18. p11\_crypto\_plugin

The following table outlines the options available under the **[p11\_crypto\_plugin]** group in the **/etc/barbican/barbican.conf** file.

Table 1.18. p11\_crypto\_plugin



Configuration option = Default value	Type	Description
<b>aes_gcm_generate_iv = True</b>	boolean value	Generate IVs for CKM_AES_GCM mechanism.
<b>always_set_cka_sensitive = True</b>	boolean value	Always set CKA_SENSITIVE=CK_TRUE including CKA_EXTRACTABLE=CK_TRUE keys.
<b>encryption_mechanism = CKM_AES_CBC</b>	string value	Secret encryption mechanism
<b>hmac_key_type = CKK_AES</b>	string value	HMAC Key Type
<b>hmac_keygen_mechanism = CKM_AES_KEY_GEN</b>	string value	HMAC Key Generation Algorithm
<b>hmac_label = None</b>	string value	Master HMAC Key label (as stored in the HSM)
<b>library_path = None</b>	string value	Path to vendor PKCS11 library
<b>login = None</b>	string value	Password to login to PKCS11 session
<b>mkek_label = None</b>	string value	Master KEK label (as stored in the HSM)
<b>mkek_length = None</b>	integer value	Master KEK length in bytes.
<b>pkek_cache_limit = 100</b>	integer value	Project KEK Cache Item Limit
<b>pkek_cache_ttl = 900</b>	integer value	Project KEK Cache Time To Live, in seconds
<b>pkek_length = 32</b>	integer value	Project KEK length in bytes.
<b>plugin_name = PKCS11 HSM</b>	string value	User friendly plugin name
<b>rw_session = True</b>	boolean value	Flag for Read/Write Sessions
<b>seed_file = `</b>	string value	File to pull entropy for seeding RNG
<b>seed_length = 32</b>	integer value	Amount of data to read from file for seed
<b>slot_id = 1</b>	integer value	HSM Slot ID

### 1.1.19. queue

The following table outlines the options available under the **[queue]** group in the **/etc/barbican/barbican.conf** file.

Table 1.19. queue

Configuration option = Default value	Type	Description
<b>asynchronous_workers = 1</b>	integer value	Number of asynchronous worker processes
<b>enable = False</b>	boolean value	True enables queuing, False invokes workers synchronously
<b>namespace = barbican</b>	string value	Queue namespace
<b>server_name = barbican.queue</b>	string value	Server name for RPC task processing server
<b>topic = barbican.workers</b>	string value	Queue topic name
<b>version = 1.1</b>	string value	Version of tasks invoked via queue

### 1.1.20. quotas

The following table outlines the options available under the **[quotas]** group in the **/etc/barbican/barbican.conf** file.

Table 1.20. quotas

Configuration option = Default value	Type	Description
<b>quota_cas = -1</b>	integer value	Number of CAs allowed per project
<b>quota_consumers = -1</b>	integer value	Number of consumers allowed per project
<b>quota_containers = -1</b>	integer value	Number of containers allowed per project
<b>quota_orders = -1</b>	integer value	Number of orders allowed per project
<b>quota_secrets = -1</b>	integer value	Number of secrets allowed per project

### 1.1.21. retry\_scheduler

The following table outlines the options available under the **[retry\_scheduler]** group in the **/etc/barbican/barbican.conf** file.

Table 1.21. retry\_scheduler

Configuration option = Default value	Type	Description
<b>initial_delay_seconds = 10.0</b>	floating point value	Seconds (float) to wait before starting retry scheduler
<b>periodic_interval_max_seconds = 10.0</b>	floating point value	Seconds (float) to wait between periodic schedule events

### 1.1.22. secretstore

The following table outlines the options available under the **[secretstore]** group in the `/etc/barbican/barbican.conf` file.

Table 1.22. secretstore

Configuration option = Default value	Type	Description
<b>enable_multiple_secret_stores = False</b>	boolean value	Flag to enable multiple secret store plugin backend support. Default is False
<b>enabled_secretstore_plugins = [u'store_crypto']</b>	multi valued	List of secret store plugins to load.
<b>namespace = barbican.secretstore.plugin</b>	string value	Extension namespace to search for plugins.
<b>stores_lookup_suffix = None</b>	list value	List of suffix to use for looking up plugins which are supported with multiple backend support.

### 1.1.23. simple\_crypto\_plugin

The following table outlines the options available under the **[simple\_crypto\_plugin]** group in the `/etc/barbican/barbican.conf` file.

Table 1.23. simple\_crypto\_plugin

Configuration option = Default value	Type	Description
<b>kek = dGhpcnR5X3R3b19ieXRI X2tleWJsYWhibGFoYmxh aGg=</b>	string value	Key encryption key to be used by Simple Crypto Plugin
<b>plugin_name = Software Only Crypto</b>	string value	User friendly plugin name

### 1.1.24. snakeoil\_ca\_plugin

The following table outlines the options available under the **[snakeoil\_ca\_plugin]** group in the **/etc/barbican/barbican.conf** file.

Table 1.24. snakeoil\_ca\_plugin

Configuration option = Default value	Type	Description
<b>ca_cert_chain_path = None</b>	string value	Path to CA certificate chain file
<b>ca_cert_key_path = None</b>	string value	Path to CA certificate key file
<b>ca_cert_path = None</b>	string value	Path to CA certificate file
<b>ca_cert_pkcs7_path = None</b>	string value	Path to CA chain pkcs7 file
<b>subca_cert_key_directory = /etc/barbican/snakeoil-cas</b>	string value	Directory in which to store certs/keys for subcas

### 1.1.25. ssl

The following table outlines the options available under the **[ssl]** group in the **/etc/barbican/barbican.conf** file.

Table 1.25. ssl

Configuration option = Default value	Type	Description
<b>ca_file = None</b>	string value	CA certificate file to use to verify connecting clients.
<b>cert_file = None</b>	string value	Certificate file to use when starting the server securely.
<b>ciphers = None</b>	string value	Sets the list of available ciphers. value should be a string in the OpenSSL cipher list format.
<b>key_file = None</b>	string value	Private key file to use when starting the server securely.
<b>version = None</b>	string value	SSL version to use (valid only if SSL enabled). Valid values are TLSv1 and SSLv23. SSLv2, SSLv3, TLSv1_1, and TLSv1_2 may be available on some distributions.

## CHAPTER 2. CEILOMETER

The following chapter contains information about the configuration options in the **ceilometer** service.

### 2.1. CEILOMETER.CONF

This section contains options for the `/etc/ceilometer/ceilometer.conf` file.

#### 2.1.1. DEFAULT

The following table outlines the options available under the **[DEFAULT]** group in the `/etc/ceilometer/ceilometer.conf` file.

Table 2.1. DEFAULT

Configuration option = Default value	Type	Description
<b>backdoor_port = None</b>	string value	Enable eventlet backdoor. Acceptable values are 0, <port>, and <start>:<end>, where 0 results in listening on a random tcp port number; <port> results in listening on the specified port number (and not enabling backdoor if that port is in use); and <start>:<end> results in listening on the smallest unused port number within the specified range of port numbers. The chosen port is displayed in the service's log file.
<b>backdoor_socket = None</b>	string value	Enable eventlet backdoor, using the provided path as a unix socket that can receive connections. This option is mutually exclusive with <i>backdoor_port</i> in that only one should be provided. If both are provided then the existence of this option overrides the usage of that option.
<b>batch_polled_samples = True</b>	boolean value	To reduce polling agent load, samples are sent to the notification agent in a batch. To gain higher throughput at the cost of load set this to False.
<b>conn_pool_min_size = 2</b>	integer value	The pool size limit for connections expiration policy
<b>conn_pool_ttl = 1200</b>	integer value	The time-to-live in sec of idle connections in the pool
<b>control_exchange = openstack</b>	string value	The default exchange under which topics are scoped. May be overridden by an exchange name specified in the <i>transport_url</i> option.
<b>debug = False</b>	boolean value	If set to true, the logging level will be set to DEBUG instead of the default INFO level.

Configuration option = Default value	Type	Description
<b>default_log_levels =</b> [u'amqp=WARN', u'amqplib=WARN', u'boto=WARN', u'qpid=WARN', u'sqlalchemy=WARN', u'suds=INFO', u'oslo.messaging=INFO', u'oslo_messaging=INFO', u'iso8601=WARN', u'requests.packages.urlli b3.connectionpool=WAR N', u'urllib3.connectionpool= WARN', u'websocket=WARN', u'requests.packages.urlli b3.util.retry=WARN', u'urllib3.util.retry=WARN', u'keystonemiddleware=W ARN', u'routes.middleware=WA RN', u'stevedore=WARN', u'taskflow=WARN', u'keystoneauth=WARN', u'oslo.cache=INFO', u'dogpile.core.dogpile=IN FO']	list value	List of package logging levels in logger=LEVEL pairs. This option is ignored if log_config_append is set.
<b>event_pipeline_cfg_file =</b> event_pipeline.yaml	string value	Configuration file for event pipeline definition.
<b>executor_thread_pool_siz e = 64</b>	integer value	Size of executor thread pool when executor is threading or eventlet.
<b>fatal_deprecations =</b> False	boolean value	Enables or disables fatal status of deprecations.
<b>graceful_shutdown_timeo ut = 60</b>	integer value	Specify a timeout after which a gracefully shutdown server will exit. Zero value means endless wait.
<b>host = ccsosp- undercloud.usersys.redh at.com</b>	unknown value	Name of this node, which must be valid in an AMQP key. Can be an opaque identifier. For ZeroMQ only, must be a valid host name, FQDN, or IP address.
<b>http_timeout = 600</b>	integer value	Timeout seconds for HTTP requests. Set it to None to disable timeout.

Configuration option = Default value	Type	Description
<b>hypervisor_inspector = libvirt</b>	string value	Inspector to use for inspecting the hypervisor layer. Known inspectors are libvirt, hyperv, vsphere and xenapi.
<code>`instance_format = [instance: %(uuid)s] `</code>	string value	The format for an instance that is passed with the log message.
<code>`instance_uuid_format = [instance: %(uuid)s] `</code>	string value	The format for an instance UUID that is passed with the log message.
<b>libvirt_type = kvm</b>	string value	Libvirt domain type.
<code>`libvirt_uri = `</code>	string value	Override the default libvirt URI (which is dependent on libvirt_type).
<b>log-config-append = None</b>	string value	The name of a logging configuration file. This file is appended to any existing logging configuration files. For details about logging configuration files, see the Python logging module documentation. Note that when logging configuration files are used then all logging configuration is set in the configuration file and other logging configuration options are ignored (for example, logging_context_format_string).
<b>log-date-format = %Y-%m-%d %H:%M:%S</b>	string value	Defines the format string for <code>%(asctime)s</code> in log records. Default: <code>%(default)s</code> . This option is ignored if <code>log_config_append</code> is set.
<b>log-dir = None</b>	string value	(Optional) The base directory used for relative <code>log_file</code> paths. This option is ignored if <code>log_config_append</code> is set.
<b>log-file = None</b>	string value	(Optional) Name of log file to send logging output to. If no default is set, logging will go to <code>stderr</code> as defined by <code>use_stderr</code> . This option is ignored if <code>log_config_append</code> is set.
<b>log_options = True</b>	boolean value	Enables or disables logging values of all registered options when starting a service (at DEBUG level).
<b>logging_context_format_string = %(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [%request_id)s %(user_identity)s %(instance)s%(message)s</b>	string value	Format string to use for log messages with context.

Configuration option = Default value	Type	Description
<b>logging_debug_format_s uffix = %(funcName)s % (pathname)s:%(lineno)d</b>	string value	Additional data to append to log message when logging level for the message is DEBUG.
<b>logging_default_format_s tring = %(asctime)s.% (msecs)03d %(process)d %(levelname)s %(name)s [-] %(instance)s% (message)s</b>	string value	Format string to use for log messages when context is undefined.
<b>logging_exception_prefix = %(asctime)s.% (msecs)03d %(process)d ERROR %(name)s % (instance)s</b>	string value	Prefix each line of exception output with this format.
<b>logging_user_identity_for mat = %(user)s % (tenant)s %(domain)s % (user_domain)s % (project_domain)s</b>	string value	Defines the format string for %(user_identity)s that is used in logging_context_format_string.
<b>max_parallel_requests = 64</b>	integer value	Maximum number of parallel requests for services to handle at the same time.
<b>pipeline_cfg_file = pipeline.yaml</b>	string value	Configuration file for pipeline definition.
<b>publish_errors = False</b>	boolean value	Enables or disables publication of error events.
<b>rate_limit_burst = 0</b>	integer value	Maximum number of logged messages per rate_limit_interval.
<b>rate_limit_except_level = CRITICAL</b>	string value	Log level name used by rate limiting: CRITICAL, ERROR, INFO, WARNING, DEBUG or empty string. Logs with level greater or equal to rate_limit_except_level are not filtered. An empty string means that all levels are filtered.
<b>rate_limit_interval = 0</b>	integer value	Interval, number of seconds, of log rate limiting.
<b>reseller_prefix = AUTH_</b>	string value	Swift reseller prefix. Must be on par with reseller_prefix in proxy-server.conf.



Configuration option = Default value	Type	Description
<b>reserved_metadata_keys</b> = []	list value	List of metadata keys reserved for metering use. And these keys are additional to the ones included in the namespace.
<b>reserved_metadata_length</b> = 256	integer value	Limit on length of reserved metadata values.
<b>reserved_metadata_namespace</b> = [u'metering.']	list value	List of metadata prefixes reserved for metering use.
<b>rootwrap_config</b> = /etc/ceilometer/rootwrap.conf	string value	Path to the rootwrap configuration file to use for running commands as root
<b>rpc_ack_timeout_base</b> = 15	integer value	Number of seconds to wait for an ack from a cast/call. After each retry attempt this timeout is multiplied by some specified multiplier.
<b>rpc_ack_timeout_multiplier</b> = 2	integer value	Number to multiply base ack timeout by after each retry attempt.
<b>rpc_backend</b> = rabbit	string value	The messaging driver to use, defaults to rabbit. Other drivers include amqp and zmq.
<b>rpc_conn_pool_size</b> = 30	integer value	Size of RPC connection pool.
<b>rpc_message_ttl</b> = 300	integer value	Expiration timeout in seconds of a sent/received message after which it is not tracked anymore by a client/server.
<b>rpc_poll_timeout</b> = 1	integer value	The default number of seconds that poll should wait. Poll raises timeout exception when timeout expired.
<b>rpc_response_timeout</b> = 60	integer value	Seconds to wait for a response from a call.
<b>rpc_retry_attempts</b> = 3	integer value	Default number of message sending attempts in case of any problems occurred: positive value N means at most N retries, 0 means no retries, None or -1 (or any other negative values) mean to retry forever. This option is used only if acknowledgments are enabled.
<b>rpc_thread_pool_size</b> = 100	integer value	Maximum number of (green) threads to work concurrently.

Configuration option = Default value	Type	Description
<b>rpc_use_acks = False</b>	boolean value	Wait for message acknowledgements from receivers. This mechanism works only via proxy without PUB/SUB.
<b>rpc_zmq_bind_address = *</b>	string value	ZeroMQ bind address. Should be a wildcard (*), an ethernet interface, or IP. The "host" option should point or resolve to this address.
<b>rpc_zmq_bind_port_retries = 100</b>	integer value	Number of retries to find free port number before fail with ZMQBindError.
<b>rpc_zmq_contexts = 1</b>	integer value	Number of ZeroMQ contexts, defaults to 1.
<b>rpc_zmq_host = ccosp-undercloud.usersys.redhat.com</b>	string value	Name of this node. Must be a valid hostname, FQDN, or IP address. Must match "host" option, if running Nova.
<b>rpc_zmq_ipc_dir = /var/run/openstack</b>	string value	Directory for holding IPC sockets.
<b>rpc_zmq_matchmaker = redis</b>	string value	MatchMaker driver.
<b>rpc_zmq_max_port = 65536</b>	integer value	Maximal port number for random ports range.
<b>rpc_zmq_min_port = 49153</b>	port value	Minimal port number for random ports range.
<b>rpc_zmq_serialization = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>rpc_zmq_topic_backlog = None</b>	integer value	Maximum number of ingress messages to locally buffer per topic. Default is unlimited.
<b>sample_source = openstack</b>	string value	Source for samples emitted on this instance.
<b>subscribe_on = []</b>	list value	List of publisher hosts SubConsumer can subscribe on. This option has higher priority then the default publishers list taken from the matchmaker.
<b>syslog-log-facility = LOG_USER</b>	string value	Syslog facility to receive log lines. This option is ignored if log_config_append is set.

Configuration option = Default value	Type	Description
<b>transport_url = None</b>	string value	<p>The network address and optional user credentials for connecting to the messaging backend, in URL format. The expected format is:</p> <pre>driver://[user:pass@]host:port[, [userN:passN@]hostN:portN]/virtual_host?query</pre> <p>Example: rabbit://rabbitmq:password@127.0.0.1:5672//</p> <p>For full details on the fields in the URL see the documentation of oslo_messaging.TransportURL at <a href="https://docs.openstack.org/oslo.messaging/latest/reference/transport.html">https://docs.openstack.org/oslo.messaging/latest/reference/transport.html</a></p>
<b>use-journal = False</b>	boolean value	Enable journald for logging. If running in a systemd environment you may wish to enable journal support. Doing so will use the journal native protocol which includes structured metadata in addition to log messages. This option is ignored if log_config_append is set.
<b>use-json = False</b>	boolean value	Use JSON formatting for logging. This option is ignored if log_config_append is set.
<b>use-syslog = False</b>	boolean value	Use syslog for logging. Existing syslog format is DEPRECATED and will be changed later to honor RFC5424. This option is ignored if log_config_append is set.
<b>use_dynamic_connections = False</b>	boolean value	This option makes direct connections dynamic or static. It makes sense only with use_router_proxy=False which means to use direct connections for direct message types (ignored otherwise).
<b>use_pub_sub = False</b>	boolean value	Use PUB/SUB pattern for fanout methods. PUB/SUB always uses proxy.
<b>use_router_proxy = False</b>	boolean value	Use ROUTER remote proxy.
<b>use_stderr = False</b>	boolean value	Log output to standard error. This option is ignored if log_config_append is set.

Configuration option = Default value	Type	Description
<b>watch-log-file = False</b>	boolean value	Uses logging handler designed to watch file system. When log file is moved or removed this handler will open a new log file with specified path instantaneously. It makes sense only if log_file option is specified and Linux platform is used. This option is ignored if log_config_append is set.
<b>zmq_failover_connections = 2</b>	integer value	How many additional connections to a host will be made for failover reasons. This option is actual only in dynamic connections mode.
<b>zmq_immediate = True</b>	boolean value	This option configures round-robin mode in zmq socket. True means not keeping a queue when server side disconnects. False means to keep queue and messages even if server is disconnected, when the server appears we send all accumulated messages to it.
<b>zmq_linger = -1</b>	integer value	Number of seconds to wait before all pending messages will be sent after closing a socket. The default value of -1 specifies an infinite linger period. The value of 0 specifies no linger period. Pending messages shall be discarded immediately when the socket is closed. Positive values specify an upper bound for the linger period.
<b>zmq_target_expire = 300</b>	integer value	Expiration timeout in seconds of a name service record about existing target (< 0 means no timeout).
<b>zmq_target_update = 180</b>	integer value	Update period in seconds of a name service record about existing target.
<b>zmq_tcp_keepalive = -1</b>	integer value	Enable/disable TCP keepalive (KA) mechanism. The default value of -1 (or any other negative value) means to skip any overrides and leave it to OS default; 0 and 1 (or any other positive value) mean to disable and enable the option respectively.
<b>zmq_tcp_keepalive_cnt = -1</b>	integer value	The number of retransmissions to be carried out before declaring that remote end is not available. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.

Configuration option = Default value	Type	Description
<b>zmq_tcp_keepalive_idle = -1</b>	integer value	The duration between two keepalive transmissions in idle condition. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_intvl = -1</b>	integer value	The duration between two successive keepalive retransmissions, if acknowledgement to the previous keepalive transmission is not received. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.

### 2.1.2. compute

The following table outlines the options available under the **[compute]** group in the `/etc/ceilometer/ceilometer.conf` file.

Table 2.2. compute

Configuration option = Default value	Type	Description
<b>instance_discovery_method = libvirt_metadata</b>	string value	Ceilometer offers many methods to discover the instance running on a compute node: <ul style="list-style-type: none"> <li>* naive: poll nova to get all instances</li> <li>* workload_partitioning: poll nova to get instances of the compute *</li> <li>libvirt_metadata: get instances from libvirt metadata but without instance metadata (recommended for Gnocchi backend)</li> </ul>
<b>resource_cache_expiry = 3600</b>	integer value	The expiry to totally refresh the instances resource cache, since the instance may be migrated to another host, we need to clean the legacy instances info in local cache by totally refreshing the local cache. The minimum should be the value of the config option of <code>resource_update_interval</code> . This option is only used for agent polling to Nova API, so it will work only when <code>instance_discovery_method</code> is set to <i>naive</i> .

Configuration option = Default value	Type	Description
<b>resource_update_interval = 0</b>	integer value	New instances will be discovered periodically based on this option (in seconds). By default, the agent discovers instances according to pipeline polling interval. If option is greater than 0, the instance list to poll will be updated based on this option's interval. Measurements relating to the instances will match intervals defined in pipeline. This option is only used for agent polling to Nova API, so it will work only when <i>instance_discovery_method</i> is set to <i>naive</i> .

### 2.1.3. coordination

The following table outlines the options available under the **[coordination]** group in the */etc/ceilometer/ceilometer.conf* file.

Table 2.3. coordination

Configuration option = Default value	Type	Description
<b>backend_url = None</b>	string value	The backend URL to use for distributed coordination. If left empty, per-deployment central agent and per-host compute agent won't do workload partitioning and will only function correctly if a single instance of that service is running.
<b>check_watchers = 10.0</b>	floating point value	Number of seconds between checks to see if group membership has changed

### 2.1.4. dispatcher\_gnocchi

The following table outlines the options available under the **[dispatcher\_gnocchi]** group in the */etc/ceilometer/ceilometer.conf* file.

Table 2.4. dispatcher\_gnocchi

Configuration option = Default value	Type	Description
<b>archive_policy = None</b>	string value	The archive policy to use when the dispatcher create a new metric.
<b>filter_project = gnocchi</b>	string value	Gnocchi project used to filter out samples generated by Gnocchi service activity

Configuration option = Default value	Type	Description
<b>request_timeout = 6.05</b>	floating point value	Number of seconds before request to gnocchi times out
<b>resources_definition_file = gnocchi_resources.yaml</b>	string value	The Yaml file that defines mapping between samples and gnocchi resources/metrics

### 2.1.5. event

The following table outlines the options available under the **[event]** group in the `/etc/ceilometer/ceilometer.conf` file.

Table 2.5. event

Configuration option = Default value	Type	Description
<b>definitions_cfg_file = event_definitions.yaml</b>	string value	Configuration file for event definitions.
<b>drop_unmatched_notifications = False</b>	boolean value	Drop notifications if no event definition matches. (Otherwise, we convert them with just the default traits)
<b>store_raw = []</b>	multi valued	Store the raw notification for select priority levels (info and/or error). By default, raw details are not captured.

### 2.1.6. hardware

The following table outlines the options available under the **[hardware]** group in the `/etc/ceilometer/ceilometer.conf` file.

Table 2.6. hardware

Configuration option = Default value	Type	Description
<b>meter_definitions_file = snmp.yaml</b>	string value	Configuration file for defining hardware snmp meters.
<b>readonly_user_auth_protocol = None</b>	string value	SNMPd v3 authentication algorithm of all the nodes running in the cloud
<b>readonly_user_name = ro_snmp_user</b>	string value	SNMPd user name of all nodes running in the cloud.

Configuration option = Default value	Type	Description
<b>readonly_user_password = password</b>	string value	SNMPd v3 authentication password of all the nodes running in the cloud.
<b>readonly_user_priv_password = None</b>	string value	SNMPd v3 encryption password of all the nodes running in the cloud.
<b>readonly_user_priv_proto = None</b>	string value	SNMPd v3 encryption algorithm of all the nodes running in the cloud
<b>tripleo_network_name = ctlplane</b>	string value	Name of the control plane Tripleo network
<b>url_scheme = snmp://</b>	string value	URL scheme to use for hardware nodes.

### 2.1.7. ipmi

The following table outlines the options available under the **[ipmi]** group in the `/etc/ceilometer/ceilometer.conf` file.

Table 2.7. ipmi

Configuration option = Default value	Type	Description
<b>node_manager_init_retry = 3</b>	integer value	Number of retries upon Intel Node Manager initialization failure
<b>polling_retry = 3</b>	integer value	Tolerance of IPMI/NM polling failures before disable this pollster. Negative indicates retrying forever.

### 2.1.8. matchmaker\_redis

The following table outlines the options available under the **[matchmaker\_redis]** group in the `/etc/ceilometer/ceilometer.conf` file.

Table 2.8. matchmaker\_redis

Configuration option = Default value	Type	Description
<b>check_timeout = 20000</b>	integer value	Time in ms to wait before the transaction is killed.
<b>host = 127.0.0.1</b>	string value	Host to locate redis.
<b>password = ``</b>	string value	Password for Redis server (optional).



Configuration option = Default value	Type	Description
<b>port = 6379</b>	port value	Use this port to connect to redis host.
<b>sentinel_group_name = oslo-messaging-zeromq</b>	string value	Redis replica set name.
<b>sentinel_hosts = []</b>	list value	List of Redis Sentinel hosts (fault tolerance mode), e.g., [host:port, host1:port ... ]
<b>socket_timeout = 10000</b>	integer value	Timeout in ms on blocking socket operations.
<b>wait_timeout = 2000</b>	integer value	Time in ms to wait between connection attempts.

### 2.1.9. meter

The following table outlines the options available under the **[meter]** group in the **/etc/ceilometer/ceilometer.conf** file.

Table 2.9. meter

Configuration option = Default value	Type	Description
<b>meter_definitions_cfg_file = None</b>	string value	Configuration file for defining meter notifications. This option is deprecated and use <code>meter_definitions_dirs</code> to configure meter notification file. Meter definitions configuration file will be sought according to the parameter.
<b>meter_definitions_dirs = [u'/etc/ceilometer/meters.d', u'/usr/lib/python2.7/site-packages/ceilometer/data/meters.d']</b>	multi valued	List directory to find files of defining meter notifications.

### 2.1.10. notification

The following table outlines the options available under the **[notification]** group in the **/etc/ceilometer/ceilometer.conf** file.

Table 2.10. notification

Configuration option = Default value	Type	Description
<b>ack_on_event_error = True</b>	boolean value	Acknowledge message when event persistence fails.
<b>batch_size = 100</b>	integer value	Number of notification messages to wait before publishing them. Batching is advised when transformations are applied in pipeline.
<b>batch_timeout = 5</b>	integer value	Number of seconds to wait before publishing samples when batch_size is not reached (None means indefinitely)
<b>messaging_urls = []</b>	multi valued	Messaging URLs to listen for notifications. Example: rabbit://user:pass@host1:port1[,user:pass@hostN:port N]/virtual_host (DEFAULT/transport_url is used if empty). This is useful when you have dedicate messaging nodes for each service, for example, all nova notifications go to rabbit-nova:5672, while all cinder notifications go to rabbit-cinder:5672.
<b>notification_control_exchanges = [u'nova', u'glance', u'neutron', u'cinder', u'heat', u'keystone', u'sahara', u'trove', u'zaqar', u'swift', u'ceilometer', u'magnum', u'dns', u'ironic', u'aodh']</b>	multi valued	Exchanges name to listen for notifications.
<b>pipeline_processing_queues = 10</b>	integer value	Number of queues to parallelize workload across. This value should be larger than the number of active notification agents for optimal results. WARNING: Once set, lowering this value may result in lost data.
<b>pipelines = [u'meter', u'event']</b>	multi valued	Select which pipeline managers to enable to generate data
<b>workers = 1</b>	integer value	Number of workers for notification service, default value is 1.
<b>workload_partitioning = False</b>	boolean value	Enable workload partitioning, allowing multiple notification agents to be run simultaneously.

### 2.1.11. oslo\_concurrency

The following table outlines the options available under the **[oslo\_concurrency]** group in the `/etc/ceilometer/ceilometer.conf` file.

Table 2.11. oslo\_concurrency

Configuration option = Default value	Type	Description
<b>disable_process_locking = False</b>	boolean value	Enables or disables inter-process locks.
<b>lock_path = None</b>	string value	Directory to use for lock files. For security, the specified directory should only be writable by the user running the processes that need locking. Defaults to environment variable OSLO_LOCK_PATH. If external locks are used, a lock path must be set.

### 2.1.12. oslo\_messaging\_amqp

The following table outlines the options available under the **[oslo\_messaging\_amqp]** group in the `/etc/ceilometer/ceilometer.conf` file.

Table 2.12. oslo\_messaging\_amqp

Configuration option = Default value	Type	Description
<b>addressing_mode = dynamic</b>	string value	Indicates the addressing mode used by the driver. Permitted values: <i>legacy</i> - use legacy non-routable addressing <i>routable</i> - use routable addresses <i>dynamic</i> - use legacy addresses if the message bus does not support routing otherwise use routable addressing
<b>allow_insecure_clients = False</b>	boolean value	Accept clients using either SSL or plain TCP
<b>anycast_address = anycast</b>	string value	Appended to the address prefix when sending to a group of consumers. Used by the message bus to identify messages that should be delivered in a round-robin fashion across consumers.
<b>broadcast_prefix = broadcast</b>	string value	address prefix used when broadcasting to all servers
<b>connection_retry_backoff = 2</b>	integer value	Increase the <code>connection_retry_interval</code> by this many seconds after each unsuccessful failover attempt.
<b>connection_retry_interval = 1</b>	integer value	Seconds to pause before attempting to re-connect.
<b>connection_retry_interval_max = 30</b>	integer value	Maximum limit for <code>connection_retry_interval</code> + <code>connection_retry_backoff</code>

Configuration option = Default value	Type	Description
<b>container_name = None</b>	string value	Name for the AMQP container. must be globally unique. Defaults to a generated UUID
<b>default_notification_exchange = None</b>	string value	Exchange name used in notification addresses. Exchange name resolution precedence: Target.exchange if set else default_notification_exchange if set else control_exchange if set else <i>notify</i>
<b>default_notify_timeout = 30</b>	integer value	The deadline for a sent notification message delivery. Only used when caller does not provide a timeout expiry.
<b>default_reply_retry = 0</b>	integer value	The maximum number of attempts to re-send a reply message which failed due to a recoverable error.
<b>default_reply_timeout = 30</b>	integer value	The deadline for an rpc reply message delivery.
<b>default_rpc_exchange = None</b>	string value	Exchange name used in RPC addresses. Exchange name resolution precedence: Target.exchange if set else default_rpc_exchange if set else control_exchange if set else <i>rpc</i>
<b>default_send_timeout = 30</b>	integer value	The deadline for an rpc cast or call message delivery. Only used when caller does not provide a timeout expiry.
<b>default_sender_link_timeout = 600</b>	integer value	The duration to schedule a purge of idle sender links. Detach link after expiry.
<b>group_request_prefix = unicast</b>	string value	address prefix when sending to any server in group
<b>idle_timeout = 0</b>	integer value	Timeout for inactive connections (in seconds)
<b>link_retry_delay = 10</b>	integer value	Time to pause between re-connecting an AMQP 1.0 link that failed due to a recoverable error.
<b>multicast_address = multicast</b>	string value	Appended to the address prefix when sending a fanout message. Used by the message bus to identify fanout messages.
<b>notify_address_prefix = openstack.org/om/notify</b>	string value	Address prefix for all generated Notification addresses

Configuration option = Default value	Type	Description
<b>notify_server_credit = 100</b>	integer value	Window size for incoming Notification messages
<code>`password = `</code>	string value	Password for message broker authentication
<b>pre_settled = [u'rpc-cast', u'rpc-reply']</b>	multi valued	Send messages of this type pre-settled. Pre-settled messages will not receive acknowledgement from the peer. Note well: pre-settled messages may be silently discarded if the delivery fails. Permitted values: <i>rpc-call</i> - send RPC Calls pre-settled <i>rpc-reply</i> - send RPC Replies pre-settled <i>rpc-cast</i> - Send RPC Casts pre-settled <i>notify</i> - Send Notifications pre-settled
<b>pseudo_vhost = True</b>	boolean value	Enable virtual host support for those message buses that do not natively support virtual hosting (such as <i>qpidd</i> ). When set to true the virtual host name will be added to all message bus addresses, effectively creating a private <i>subnet</i> per virtual host. Set to False if the message bus supports virtual hosting using the <i>hostname</i> field in the AMQP 1.0 Open performative as the name of the virtual host.
<b>reply_link_credit = 200</b>	integer value	Window size for incoming RPC Reply messages.
<b>rpc_address_prefix = openstack.org/om/rpc</b>	string value	Address prefix for all generated RPC addresses
<b>rpc_server_credit = 100</b>	integer value	Window size for incoming RPC Request messages
<code>`sasl_config_dir = `</code>	string value	Path to directory that contains the SASL configuration
<code>`sasl_config_name = `</code>	string value	Name of configuration file (without <i>.conf</i> suffix)
<code>`sasl_default_realm = `</code>	string value	SASL realm to use if no realm present in username
<code>`sasl_mechanisms = `</code>	string value	Space separated list of acceptable SASL mechanisms
<b>server_request_prefix = exclusive</b>	string value	address prefix used when sending to a specific server
<b>ssl = False</b>	boolean value	Attempt to connect via SSL. If no other <i>ssl</i> -related parameters are given, it will use the system's CA-bundle to verify the server's certificate.

Configuration option = Default value	Type	Description
<code>`ssl_ca_file = `</code>	string value	CA certificate PEM file used to verify the server's certificate
<code>`ssl_cert_file = `</code>	string value	Self-identifying certificate PEM file for client authentication
<code>`ssl_key_file = `</code>	string value	Private key PEM file used to sign <code>ssl_cert_file</code> certificate (optional)
<b>ssl_key_password = None</b>	string value	Password for decrypting <code>ssl_key_file</code> (if encrypted)
<b>ssl_verify_vhost = False</b>	boolean value	By default SSL checks that the name in the server's certificate matches the hostname in the <code>transport_url</code> . In some configurations it may be preferable to use the virtual hostname instead, for example if the server uses the Server Name Indication TLS extension (rfc6066) to provide a certificate per virtual host. Set <code>ssl_verify_vhost</code> to True if the server's SSL certificate uses the virtual host name instead of the DNS name.
<b>trace = False</b>	boolean value	Debug: dump AMQP frames to stdout
<b>unicast_address = unicast</b>	string value	Appended to the address prefix when sending to a particular RPC/Notification server. Used by the message bus to identify messages sent to a single destination.
<code>`username = `</code>	string value	User name for message broker authentication

### 2.1.13. oslo\_messaging\_kafka

The following table outlines the options available under the **[oslo\_messaging\_kafka]** group in the `/etc/ceilometer/ceilometer.conf` file.

Table 2.13. oslo\_messaging\_kafka

Configuration option = Default value	Type	Description
<b>conn_pool_min_size = 2</b>	integer value	The pool size limit for connections expiration policy
<b>conn_pool_ttl = 1200</b>	integer value	The time-to-live in sec of idle connections in the pool

Configuration option = Default value	Type	Description
<b>consumer_group = oslo_messaging_consumer</b>	string value	Group id for Kafka consumer. Consumers in one group will coordinate message consumption
<b>kafka_consumer_timeout = 1.0</b>	floating point value	Default timeout(s) for Kafka consumers
<b>kafka_default_host = localhost</b>	string value	Default Kafka broker Host
<b>kafka_default_port = 9092</b>	port value	Default Kafka broker Port
<b>kafka_max_fetch_bytes = 1048576</b>	integer value	Max fetch bytes of Kafka consumer
<b>pool_size = 10</b>	integer value	Pool Size for Kafka Consumers
<b>producer_batch_size = 16384</b>	integer value	Size of batch for the producer async send
<b>producer_batch_timeout = 0.0</b>	floating point value	Upper bound on the delay for KafkaProducer batching in seconds

#### 2.1.14. oslo\_messaging\_notifications

The following table outlines the options available under the **[oslo\_messaging\_notifications]** group in the `/etc/ceilometer/ceilometer.conf` file.

Table 2.14. oslo\_messaging\_notifications

Configuration option = Default value	Type	Description
<b>driver = []</b>	multi valued	The Drivers(s) to handle sending notifications. Possible values are messaging, messagingv2, routing, log, test, noop
<b>retry = -1</b>	integer value	The maximum number of attempts to re-send a notification message which failed to be delivered due to a recoverable error. 0 - No retry, -1 - indefinite
<b>topics = [u'notifications']</b>	list value	AMQP topic used for OpenStack notifications.

Configuration option = Default value	Type	Description
<b>transport_url = None</b>	string value	A URL representing the messaging driver to use for notifications. If not set, we fall back to the same configuration used for RPC.

### 2.1.15. oslo\_messaging\_rabbit

The following table outlines the options available under the **[oslo\_messaging\_rabbit]** group in the `/etc/ceilometer/ceilometer.conf` file.

Table 2.15. oslo\_messaging\_rabbit

Configuration option = Default value	Type	Description
<b>amqp_auto_delete = False</b>	boolean value	Auto-delete queues in AMQP.
<b>amqp_durable_queues = False</b>	boolean value	Use durable queues in AMQP.
<b>channel_max = None</b>	integer value	Maximum number of channels to allow
<b>connection_factory = single</b>	string value	Connection factory implementation
<b>default_notification_exchange = \${control_exchange}_notification</b>	string value	Exchange name for sending notifications
<b>default_notification_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending notification, -1 means infinite retry.
<b>default_rpc_exchange = \${control_exchange}_rpc</b>	string value	Exchange name for sending RPC messages
<b>default_rpc_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending RPC message, -1 means infinite retry. If actual retry attempts in not 0 the rpc request could be processed more than one time
<b>default_serializer_type = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages



Configuration option = Default value	Type	Description
<b>fake_rabbit = False</b>	boolean value	Deprecated, use <code>rpc_backend=kombu+memory</code> or <code>rpc_backend=fake</code>
<b>frame_max = None</b>	integer value	The maximum byte size for an AMQP frame
<b>heartbeat_interval = 3</b>	integer value	How often to send heartbeats for consumer's connections
<b>heartbeat_rate = 2</b>	integer value	How often times during the <code>heartbeat_timeout_threshold</code> we check the heartbeat.
<b>heartbeat_timeout_threshold = 60</b>	integer value	Number of seconds after which the Rabbit broker is considered down if heartbeat's keep-alive fails (0 disable the heartbeat). EXPERIMENTAL
<b>host_connection_reconnect_delay = 0.25</b>	floating point value	Set delay for reconnection to some host which has connection error
<b>kombu_compression = None</b>	string value	EXPERIMENTAL: Possible values are: <code>gzip</code> , <code>bz2</code> . If not set compression will not be used. This option may not be available in future versions.
<b>kombu_failover_strategy = round-robin</b>	string value	Determines how the next RabbitMQ node is chosen in case the one we are currently connected to becomes unavailable. Takes effect only if more than one RabbitMQ node is provided in config.
<b>kombu_missing_consumer_retry_timeout = 60</b>	integer value	How long to wait a missing client before abandoning to send it its replies. This value should not be longer than <code>rpc_response_timeout</code> .
<b>kombu_reconnect_delay = 1.0</b>	floating point value	How long to wait before reconnecting in response to an AMQP consumer cancel notification.
<b>notification_listener_prefetch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to notification listener.
<b>notification_persistence = False</b>	boolean value	Persist notification messages.
<b>notification_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending notification message
<b>pool_max_overflow = 0</b>	integer value	Maximum number of connections to create above <b>pool_max_size</b> .

Configuration option = Default value	Type	Description
<b>pool_max_size = 30</b>	integer value	Maximum number of connections to keep queued.
<b>pool_recycle = 600</b>	integer value	Lifetime of a connection (since creation) in seconds or None for no recycling. Expired connections are closed on acquire.
<b>pool_stale = 60</b>	integer value	Threshold at which inactive (since release) connections are considered stale in seconds or None for no staleness. Stale connections are closed on acquire.
<b>pool_timeout = 30</b>	integer value	Default number of seconds to wait for a connections to available
<b>rabbit_ha_queues = False</b>	boolean value	Try to use HA queues in RabbitMQ (x-ha-policy: all). If you change this option, you must wipe the RabbitMQ database. In RabbitMQ 3.0, queue mirroring is no longer controlled by the x-ha-policy argument when declaring a queue. If you just want to make sure that all queues (except those with auto-generated names) are mirrored across all nodes, run: <code>"rabbitmqctl set_policy HA ^(?!amq\).* {"ha-mode": "all"}"</code>
<b>rabbit_host = localhost</b>	string value	The RabbitMQ broker address where a single node is used.
<b>rabbit_hosts = [u'\$rabbit_host:\$rabbit_port']</b>	list value	RabbitMQ HA cluster host:port pairs.
<b>rabbit_interval_max = 30</b>	integer value	Maximum interval of RabbitMQ connection retries. Default is 30 seconds.
<b>rabbit_login_method = AMQPLAIN</b>	string value	The RabbitMQ login method.
<b>rabbit_max_retries = 0</b>	integer value	Maximum number of RabbitMQ connection retries. Default is 0 (infinite retry count).
<b>rabbit_password = guest</b>	string value	The RabbitMQ password.
<b>rabbit_port = 5672</b>	port value	The RabbitMQ broker port where a single node is used.

Configuration option = Default value	Type	Description
<b>rabbit_qos_prefetch_count = 0</b>	integer value	Specifies the number of messages to prefetch. Setting to zero allows unlimited messages.
<b>rabbit_retry_backoff = 2</b>	integer value	How long to backoff for between retries when connecting to RabbitMQ.
<b>rabbit_retry_interval = 1</b>	integer value	How frequently to retry connecting with RabbitMQ.
<b>rabbit_transient_queues_ttl = 1800</b>	integer value	Positive integer representing duration in seconds for queue TTL (x-expires). Queues which are unused for the duration of the TTL are automatically deleted. The parameter affects only reply and fanout queues.
<b>rabbit_userid = guest</b>	string value	The RabbitMQ userid.
<b>rabbit_virtual_host = /</b>	string value	The RabbitMQ virtual host.
<b>rpc_listener_prefetch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to rpc listener.
<b>rpc_queue_expiration = 60</b>	integer value	Time to live for rpc queues without consumers in seconds.
<b>rpc_reply_exchange = \${control_exchange}_rpc_reply</b>	string value	Exchange name for receiving RPC replies
<b>rpc_reply_listener_prefetch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to rpc reply listener.
<b>rpc_reply_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending reply. -1 means infinite retry during rpc_timeout
<b>rpc_reply_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending reply.
<b>rpc_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending RPC message
<b>socket_timeout = 0.25</b>	floating point value	Set socket timeout in seconds for connection's socket
<b>ssl = None</b>	boolean value	Enable SSL

Configuration option = Default value	Type	Description
<code>`ssl_ca_file = `</code>	string value	SSL certification authority file (valid only if SSL enabled).
<code>`ssl_cert_file = `</code>	string value	SSL cert file (valid only if SSL enabled).
<code>`ssl_key_file = `</code>	string value	SSL key file (valid only if SSL enabled).
<b>ssl_options = None</b>	dict value	Arguments passed to <code>ssl.wrap_socket</code>
<code>`ssl_version = `</code>	string value	SSL version to use (valid only if SSL enabled). Valid values are TLSv1 and SSLv23. SSLv2, SSLv3, TLSv1_1, and TLSv1_2 may be available on some distributions.
<b>tcp_user_timeout = 0.25</b>	floating point value	Set TCP_USER_TIMEOUT in seconds for connection's socket

### 2.1.16. oslo\_messaging\_zmq

The following table outlines the options available under the **[oslo\_messaging\_zmq]** group in the `/etc/ceilometer/ceilometer.conf` file.

Table 2.16. oslo\_messaging\_zmq

Configuration option = Default value	Type	Description
<b>rpc_ack_timeout_base = 15</b>	integer value	Number of seconds to wait for an ack from a cast/call. After each retry attempt this timeout is multiplied by some specified multiplier.
<b>rpc_ack_timeout_multiplier = 2</b>	integer value	Number to multiply base ack timeout by after each retry attempt.
<b>rpc_message_ttl = 300</b>	integer value	Expiration timeout in seconds of a sent/received message after which it is not tracked anymore by a client/server.
<b>rpc_poll_timeout = 1</b>	integer value	The default number of seconds that poll should wait. Poll raises timeout exception when timeout expired.
<b>rpc_retry_attempts = 3</b>	integer value	Default number of message sending attempts in case of any problems occurred: positive value N means at most N retries, 0 means no retries, None or -1 (or any other negative values) mean to retry forever. This option is used only if acknowledgments are enabled.

Configuration option = Default value	Type	Description
<b>rpc_thread_pool_size = 100</b>	integer value	Maximum number of (green) threads to work concurrently.
<b>rpc_use_acks = False</b>	boolean value	Wait for message acknowledgements from receivers. This mechanism works only via proxy without PUB/SUB.
<b>rpc_zmq_bind_address = *</b>	string value	ZeroMQ bind address. Should be a wildcard (*), an ethernet interface, or IP. The "host" option should point or resolve to this address.
<b>rpc_zmq_bind_port_retries = 100</b>	integer value	Number of retries to find free port number before fail with ZMQBindError.
<b>rpc_zmq_contexts = 1</b>	integer value	Number of ZeroMQ contexts, defaults to 1.
<b>rpc_zmq_host = ccosp-undercloud.usersys.redhat.com</b>	string value	Name of this node. Must be a valid hostname, FQDN, or IP address. Must match "host" option, if running Nova.
<b>rpc_zmq_ipc_dir = /var/run/openstack</b>	string value	Directory for holding IPC sockets.
<b>rpc_zmq_matchmaker = redis</b>	string value	MatchMaker driver.
<b>rpc_zmq_max_port = 65536</b>	integer value	Maximal port number for random ports range.
<b>rpc_zmq_min_port = 49153</b>	port value	Minimal port number for random ports range.
<b>rpc_zmq_serialization = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>rpc_zmq_topic_backlog = None</b>	integer value	Maximum number of ingress messages to locally buffer per topic. Default is unlimited.
<b>subscribe_on = []</b>	list value	List of publisher hosts SubConsumer can subscribe on. This option has higher priority then the default publishers list taken from the matchmaker.

Configuration option = Default value	Type	Description
<b>use_dynamic_connections = False</b>	boolean value	This option makes direct connections dynamic or static. It makes sense only with <code>use_router_proxy=False</code> which means to use direct connections for direct message types (ignored otherwise).
<b>use_pub_sub = False</b>	boolean value	Use PUB/SUB pattern for fanout methods. PUB/SUB always uses proxy.
<b>use_router_proxy = False</b>	boolean value	Use ROUTER remote proxy.
<b>zmq_failover_connections = 2</b>	integer value	How many additional connections to a host will be made for failover reasons. This option is actual only in dynamic connections mode.
<b>zmq_immediate = True</b>	boolean value	This option configures round-robin mode in zmq socket. True means not keeping a queue when server side disconnects. False means to keep queue and messages even if server is disconnected, when the server appears we send all accumulated messages to it.
<b>zmq_linger = -1</b>	integer value	Number of seconds to wait before all pending messages will be sent after closing a socket. The default value of -1 specifies an infinite linger period. The value of 0 specifies no linger period. Pending messages shall be discarded immediately when the socket is closed. Positive values specify an upper bound for the linger period.
<b>zmq_target_expire = 300</b>	integer value	Expiration timeout in seconds of a name service record about existing target (< 0 means no timeout).
<b>zmq_target_update = 180</b>	integer value	Update period in seconds of a name service record about existing target.
<b>zmq_tcp_keepalive = -1</b>	integer value	Enable/disable TCP keepalive (KA) mechanism. The default value of -1 (or any other negative value) means to skip any overrides and leave it to OS default; 0 and 1 (or any other positive value) mean to disable and enable the option respectively.
<b>zmq_tcp_keepalive_cnt = -1</b>	integer value	The number of retransmissions to be carried out before declaring that remote end is not available. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.

Configuration option = Default value	Type	Description
<b>zmq_tcp_keepalive_idle = -1</b>	integer value	The duration between two keepalive transmissions in idle condition. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_intvl = -1</b>	integer value	The duration between two successive keepalive retransmissions, if acknowledgement to the previous keepalive transmission is not received. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.

### 2.1.17. polling

The following table outlines the options available under the **[polling]** group in the `/etc/ceilometer/ceilometer.conf` file.

Table 2.17. polling

Configuration option = Default value	Type	Description
<b>cfg_file = polling.yaml</b>	string value	Configuration file for polling definition.
<b>partitioning_group_prefix = None</b>	string value	Work-load partitioning group prefix. Use only if you want to run multiple polling agents with different config files. For each sub-group of the agent pool with the same <code>partitioning_group_prefix</code> a disjoint subset of pollsters should be loaded.

### 2.1.18. publisher

The following table outlines the options available under the **[publisher]** group in the `/etc/ceilometer/ceilometer.conf` file.

Table 2.18. publisher

Configuration option = Default value	Type	Description
<b>telemetry_secret = change this for valid signing</b>	string value	Secret value for signing messages. Set value empty if signing is not required to avoid computational overhead.

### 2.1.19. publisher\_notifier

The following table outlines the options available under the **[publisher\_notifier]** group in the `/etc/ceilometer/ceilometer.conf` file.

Table 2.19. publisher\_notifier

Configuration option = Default value	Type	Description
<b>event_topic = event</b>	string value	The topic that ceilometer uses for event notifications.
<b>metering_topic = metering</b>	string value	The topic that ceilometer uses for metering notifications.
<b>telemetry_driver = messagingv2</b>	string value	The driver that ceilometer uses for metering notifications.

### 2.1.20. rgw\_admin\_credentials

The following table outlines the options available under the **[rgw\_admin\_credentials]** group in the `/etc/ceilometer/ceilometer.conf` file.

Table 2.20. rgw\_admin\_credentials

Configuration option = Default value	Type	Description
<b>access_key = None</b>	string value	Access key for Radosgw Admin.
<b>secret_key = None</b>	string value	Secret key for Radosgw Admin.

### 2.1.21. service\_credentials

The following table outlines the options available under the **[service\_credentials]** group in the `/etc/ceilometer/ceilometer.conf` file.

Table 2.21. service\_credentials

Configuration option = Default value	Type	Description
<b>auth-url = None</b>	string value	Authentication URL
<b>auth_section = None</b>	string value	Config Section from which to load plugin specific options
<b>auth_type = None</b>	string value	Authentication type to load



Configuration option = Default value	Type	Description
<b>default-domain-id = None</b>	string value	Optional domain ID to use with v3 and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.
<b>default-domain-name = None</b>	string value	Optional domain name to use with v3 API and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.
<b>domain-id = None</b>	string value	Domain ID to scope to
<b>domain-name = None</b>	string value	Domain name to scope to
<b>interface = public</b>	string value	Type of endpoint in Identity service catalog to use for communication with OpenStack services.
<b>password = None</b>	string value	User's password
<b>project-domain-id = None</b>	string value	Domain ID containing project
<b>project-domain-name = None</b>	string value	Domain name containing project
<b>project-id = None</b>	string value	Project ID to scope to
<b>project-name = None</b>	string value	Project name to scope to
<b>region-name = None</b>	string value	Region name to use for OpenStack service endpoints.
<b>system-scope = None</b>	string value	Scope for system operations
<b>trust-id = None</b>	string value	Trust ID
<b>user-domain-id = None</b>	string value	User's domain id
<b>user-domain-name = None</b>	string value	User's domain name
<b>user-id = None</b>	string value	User id
<b>username = None</b>	string value	Username

### 2.1.22. service\_types

The following table outlines the options available under the **[service\_types]** group in the `/etc/ceilometer/ceilometer.conf` file.

Table 2.22. service\_types

Configuration option = Default value	Type	Description
<b>cinder = volumev3</b>	string value	Cinder service type.
<b>glance = image</b>	string value	Glance service type.
<b>neutron = network</b>	string value	Neutron service type.
<b>neutron_lbaas_version = v2</b>	string value	Neutron load balancer version.
<b>nova = compute</b>	string value	Nova service type.
<b>radosgw = None</b>	string value	Radosgw service type.
<b>swift = object-store</b>	string value	Swift service type.

### 2.1.23. vmware

The following table outlines the options available under the **[vmware]** group in the `/etc/ceilometer/ceilometer.conf` file.

Table 2.23. vmware

Configuration option = Default value	Type	Description
<b>api_retry_count = 10</b>	integer value	Number of times a VMware vSphere API may be retried.
<b>ca_file = None</b>	string value	CA bundle file to use in verifying the vCenter server certificate.
<b>host_ip = 127.0.0.1</b>	unknown value	IP address of the VMware vSphere host.
<code>`host_password = `</code>	string value	Password of VMware vSphere.
<b>host_port = 443</b>	port value	Port of the VMware vSphere host.
<code>`host_username = `</code>	string value	Username of VMware vSphere.

Configuration option = Default value	Type	Description
<b>insecure = False</b>	boolean value	If true, the vCenter server certificate is not verified. If false, then the default CA truststore is used for verification. This option is ignored if "ca_file" is set.
<b>task_poll_interval = 0.5</b>	floating point value	Sleep time in seconds for polling an ongoing async task.
<b>wSDL_location = None</b>	string value	Optional vim service WSDL location e.g <a href="http://&lt;server&gt;/vimService.wsdl">http://&lt;server&gt;/vimService.wsdl</a> . Optional over-ride to default location for bug work-arounds.

### 2.1.24. xenapi

The following table outlines the options available under the **[xenapi]** group in the `/etc/ceilometer/ceilometer.conf` file.

Table 2.24. xenapi

Configuration option = Default value	Type	Description
<b>connection_password = None</b>	string value	Password for connection to XenServer/Xen Cloud Platform.
<b>connection_url = None</b>	string value	URL for connection to XenServer/Xen Cloud Platform.
<b>connection_username = root</b>	string value	Username for connection to XenServer/Xen Cloud Platform.

## CHAPTER 3. CINDER

The following chapter contains information about the configuration options in the **cinder** service.

### 3.1. CINDER.CONF

This section contains options for the `/etc/cinder/cinder.conf` file.

#### 3.1.1. DEFAULT

The following table outlines the options available under the **[DEFAULT]** group in the `/etc/cinder/cinder.conf` file.

Table 3.1. DEFAULT

Configuration option = Default value	Type	Description
<b>allocated_capacity_weight_multiplier = -1.0</b>	floating point value	Multiplier used for weighing allocated capacity. Positive numbers mean to stack vs spread.
<b>allow_availability_zone_fallback = False</b>	boolean value	If the requested Cinder availability zone is unavailable, fall back to the value of <code>default_availability_zone</code> , then <code>storage_availability_zone</code> , instead of failing.
<b>allow_inuse_volume_type_modification = False</b>	boolean value	DEPRECATED: Allow the ability to modify the extra-spec settings of an in-use volume-type.
<b>allowed_direct_url_schemes = []</b>	list value	A list of url schemes that can be downloaded directly via the <code>direct_url</code> . Currently supported schemes: <code>[file, cinder]</code> .
<b>api_paste_config = api-paste.ini</b>	string value	File name for the <code>paste.deploy</code> config for api service
<b>api_rate_limit = True</b>	boolean value	Enables or disables rate limit of the API.
<b>auth_strategy = keystone</b>	string value	The strategy to use for auth. Supports <code>noauth</code> or <code>keystone</code> .
<b>az_cache_duration = 3600</b>	integer value	Cache volume availability zones in memory for the provided duration in seconds

Configuration option = Default value	Type	Description
<b>backdoor_port = None</b>	string value	Enable eventlet backdoor. Acceptable values are 0, <port>, and <start>:<end>, where 0 results in listening on a random tcp port number; <port> results in listening on the specified port number (and not enabling backdoor if that port is in use); and <start>:<end> results in listening on the smallest unused port number within the specified range of port numbers. The chosen port is displayed in the service's log file.
<b>backdoor_socket = None</b>	string value	Enable eventlet backdoor, using the provided path as a unix socket that can receive connections. This option is mutually exclusive with <i>backdoor_port</i> in that only one should be provided. If both are provided then the existence of this option overrides the usage of that option.
<b>backend_availability_zone = None</b>	string value	Availability zone for this volume backend. If not set, the <i>storage_availability_zone</i> option value is used as the default for all backends.
<b>backup_api_class = cinder.backup.api.API</b>	string value	The full class name of the volume backup API class
<b>backup_ceph_chunk_size = 134217728</b>	integer value	The chunk size, in bytes, that a backup is broken into before transfer to the Ceph object store.
<b>backup_ceph_conf = /etc/ceph/ceph.conf</b>	string value	Ceph configuration file to use.
<b>backup_ceph_image_journals = False</b>	boolean value	If True, apply JOURNALING and EXCLUSIVE_LOCK feature bits to the backup RBD objects to allow mirroring
<b>backup_ceph_pool = backups</b>	string value	The Ceph pool where volume backups are stored.
<b>backup_ceph_stripe_count = 0</b>	integer value	RBD stripe count to use when creating a backup image.
<b>backup_ceph_stripe_unit = 0</b>	integer value	RBD stripe unit to use when creating a backup image.
<b>backup_ceph_user = cinder</b>	string value	The Ceph user to connect with. Default here is to use the same user as for Cinder volumes. If not using cephx this should be set to None.
<b>backup_compression_algorithm = zlib</b>	string value	Compression algorithm (None to disable)

Configuration option = Default value	Type	Description
<b>backup_container = None</b>	string value	Custom directory to use for backups.
<b>backup_driver = cinder.backup.drivers.swift.SwiftBackupDriver</b>	string value	Driver to use for backups.
<b>backup_enable_progress_timer = True</b>	boolean value	Enable or Disable the timer to send the periodic progress notifications to Ceilometer when backing up the volume to the backend storage. The default value is True to enable the timer.
<b>backup_file_size = 1999994880</b>	integer value	The maximum size in bytes of the files used to hold backups. If the volume being backed up exceeds this size, then it will be backed up into multiple files. backup_file_size must be a multiple of backup_sha_block_size_bytes.
<b>backup_gcs_block_size = 32768</b>	integer value	The size in bytes that changes are tracked for incremental backups. backup_gcs_object_size has to be multiple of backup_gcs_block_size.
<b>backup_gcs_bucket = None</b>	string value	The GCS bucket to use.
<b>backup_gcs_bucket_location = US</b>	string value	Location of GCS bucket.
<b>backup_gcs_credential_file = None</b>	string value	Absolute path of GCS service account credential file.
<b>backup_gcs_enable_progress_timer = True</b>	boolean value	Enable or Disable the timer to send the periodic progress notifications to Ceilometer when backing up the volume to the GCS backend storage. The default value is True to enable the timer.
<b>backup_gcs_num_retries = 3</b>	integer value	Number of times to retry.
<b>backup_gcs_object_size = 52428800</b>	integer value	The size in bytes of GCS backup objects.
<b>backup_gcs_project_id = None</b>	string value	Owner project id for GCS bucket.
<b>backup_gcs_proxy_url = None</b>	uri value	URL for http proxy access.

Configuration option = Default value	Type	Description
<b>backup_gcs_reader_chunk_size = 2097152</b>	integer value	GCS object will be downloaded in chunks of bytes.
<b>backup_gcs_retry_error_codes = [u'429']</b>	list value	List of GCS error codes.
<b>backup_gcs_storage_class = NEARLINE</b>	string value	Storage class of GCS bucket.
<b>backup_gcs_user_agent = gcscinder</b>	string value	Http user-agent string for gcs api.
<b>backup_gcs_writer_chunk_size = 2097152</b>	integer value	GCS object will be uploaded in chunks of bytes. Pass in a value of -1 if the file is to be uploaded as a single chunk.
<b>backup_manager = cinder.backup.manager.BackupManager</b>	string value	Full class name for the Manager for volume backup
<b>backup_metadata_version = 2</b>	integer value	Backup metadata version to be used when backing up volume metadata. If this number is bumped, make sure the service doing the restore supports the new version.
<b>backup_mount_options = None</b>	string value	Mount options passed to the NFS client. See NFS man page for details.
<b>backup_mount_point_base = \$state_path/backup_mount</b>	string value	Base dir containing mount point for NFS share.
<b>backup_name_template = backup-%s</b>	string value	Template string to be used to generate backup names
<b>backup_native_threads_pool_size = 60</b>	integer value	Size of the native threads pool for the backups. Most backup drivers rely heavily on this, it can be decreased for specific drivers that don't.
<b>backup_object_number_per_notification = 10</b>	integer value	The number of chunks or objects, for which one Ceilometer notification will be sent
<b>backup_posix_path = \$state_path/backup</b>	string value	Path specifying where to store backups.

Configuration option = Default value	Type	Description
<b>backup_service_inithost_offload = True</b>	boolean value	Offload pending backup delete during backup service startup. If false, the backup service will remain down until all pending backups are deleted.
<b>backup_sha_block_size_bytes = 32768</b>	integer value	The size in bytes that changes are tracked for incremental backups. backup_file_size has to be multiple of backup_sha_block_size_bytes.
<b>backup_share = None</b>	string value	NFS share in hostname:path, ipv4addr:path, or "[ipv6addr]:path" format.
<b>backup_swift_auth = per_user</b>	string value	Swift authentication mechanism (per_user or single_user).
<b>backup_swift_auth_insecure = False</b>	boolean value	Bypass verification of server certificate when making SSL connection to Swift.
<b>backup_swift_auth_url = None</b>	uri value	The URL of the Keystone endpoint
<b>backup_swift_auth_version = 1</b>	string value	Swift authentication version. Specify "1" for auth 1.0, or "2" for auth 2.0 or "3" for auth 3.0
<b>backup_swift_block_size = 32768</b>	integer value	The size in bytes that changes are tracked for incremental backups. backup_swift_object_size has to be multiple of backup_swift_block_size.
<b>backup_swift_ca_cert_file = None</b>	string value	Location of the CA certificate file to use for swift client requests.
<b>backup_swift_container = volumebackups</b>	string value	The default Swift container to use
<b>backup_swift_enable_progress_timer = True</b>	boolean value	Enable or Disable the timer to send the periodic progress notifications to Ceilometer when backing up the volume to the Swift backend storage. The default value is True to enable the timer.
<b>backup_swift_key = None</b>	string value	Swift key for authentication
<b>backup_swift_object_size = 52428800</b>	integer value	The size in bytes of Swift backup objects
<b>backup_swift_project = None</b>	string value	Swift project/account name. Required when connecting to an auth 3.0 system



Configuration option = Default value	Type	Description
<b>backup_swift_project_domain = None</b>	string value	Swift project domain name. Required when connecting to an auth 3.0 system
<b>backup_swift_retry_attempts = 3</b>	integer value	The number of retries to make for Swift operations
<b>backup_swift_retry_backoff = 2</b>	integer value	The backoff time in seconds between Swift retries
<b>backup_swift_tenant = None</b>	string value	Swift tenant/account name. Required when connecting to an auth 2.0 system
<b>backup_swift_url = None</b>	uri value	The URL of the Swift endpoint
<b>backup_swift_user = None</b>	string value	Swift user name
<b>backup_swift_user_domain = None</b>	string value	Swift user domain name. Required when connecting to an auth 3.0 system
<b>backup_timer_interval = 120</b>	integer value	Interval, in seconds, between two progress notifications reporting the backup status
<b>backup_tsm_compression = True</b>	boolean value	Enable or Disable compression for backups
<b>backup_tsm_password = password</b>	string value	TSM password for the running username
<b>backup_tsm_volume_prefix = backup</b>	string value	Volume prefix for the backup id when backing up to TSM
<b>backup_use_same_host = False</b>	boolean value	Backup services use same backend.
<b>backup_use_temp_snapshot_hot = False</b>	boolean value	If this is set to True, a temporary snapshot will be created for performing non-disruptive backups. Otherwise a temporary volume will be cloned in order to perform a backup.
<b>capacity_weight_multiplier = 1.0</b>	floating point value	Multiplier used for weighing free capacity. Negative numbers mean to stack vs spread.
<code>`chap_password = `</code>	string value	Password for specified CHAP account name.
<code>`chap_username = `</code>	string value	CHAP user name.

Configuration option = Default value	Type	Description
<b>chiscsi_conf = /etc/chelsio-iscsi/chiscsi.conf</b>	string value	Chiscsi (CXT) global defaults configuration file
<b>cinder_internal_tenant_project_id = None</b>	string value	ID of the project which will be used as the Cinder internal tenant.
<b>cinder_internal_tenant_user_id = None</b>	string value	ID of the user to be used in volume operations as the Cinder internal tenant.
<b>client_socket_timeout = 900</b>	integer value	Timeout for client connections' socket operations. If an incoming connection is idle for this number of seconds it will be closed. A value of 0 means wait forever.
<b>cloned_volume_same_az = True</b>	boolean value	Ensure that the new volumes are the same AZ as snapshot or source volume
<b>cluster = None</b>	string value	Name of this cluster. Used to group volume hosts that share the same backend configurations to work in HA Active-Active mode. Active-Active is not yet supported.
<b>compute_api_class = cinder.compute.nova.API</b>	string value	The full class name of the compute API class to use
<b>config-dir = [u'~/project/project.conf.d/', u'~/project.conf.d/', u'/etc/project/project.conf.d/', u'/etc/project.conf.d/']</b>	list value	Path to a config directory to pull <b>*.conf</b> files from. This file set is sorted, so as to provide a predictable parse order if individual options are over-ridden. The set is parsed after the file(s) specified via previous <code>--config-file</code> , arguments hence over-ridden options in the directory take precedence.
<b>config-file = [u'~/project/project.conf', u'~/project.conf', u'/etc/project/project.conf', u'/etc/project.conf']</b>	unknown value	Path to a config file to use. Multiple config files can be specified, with values in later files taking precedence. Defaults to <code>%(default)s</code> .
<b>conn_pool_min_size = 2</b>	integer value	The pool size limit for connections expiration policy
<b>conn_pool_ttl = 1200</b>	integer value	The time-to-live in sec of idle connections in the pool

Configuration option = Default value	Type	Description
<b>consistencygroup_api_class = cinder.consistencygroup.api.API</b>	string value	The full class name of the consistencygroup API class
<b>control_exchange = openstack</b>	string value	The default exchange under which topics are scoped. May be overridden by an exchange name specified in the transport_url option.
<b>datacore_api_timeout = 300</b>	integer value	Seconds to wait for a response from a DataCore API call.
<b>datacore_disk_failed_delay = 15</b>	integer value	Seconds to wait for DataCore virtual disk to come out of the "Failed" state.
<b>datacore_disk_pools = []</b>	list value	List of DataCore disk pools that can be used by volume driver.
<b>datacore_disk_type = single</b>	string value	DataCore virtual disk type (single/mirrored). Mirrored virtual disks require two storage servers in the server group.
<b>datacore_iscsi_chap_enabled = False</b>	boolean value	Configure CHAP authentication for iSCSI connections.
<b>datacore_iscsi_chap_storage = None</b>	string value	iSCSI CHAP authentication password storage file.
<b>datacore_iscsi_unallowed_targets = []</b>	list value	List of iSCSI targets that cannot be used to attach volume. To prevent the DataCore iSCSI volume driver from using some front-end targets in volume attachment, specify this option and list the iqn and target machine for each target as the value, such as <iqn:target name>, <iqn:target name>, <iqn:target name>.
<b>datacore_storage_profile = None</b>	string value	DataCore virtual disk storage profile.
<b>db_driver = cinder.db</b>	string value	Driver to use for database access
<b>debug = False</b>	boolean value	If set to true, the logging level will be set to DEBUG instead of the default INFO level.
<b>default_availability_zone = None</b>	string value	Default availability zone for new volumes. If not set, the storage_availability_zone option value is used as the default for new volumes.

Configuration option = Default value	Type	Description
<b>default_group_type = None</b>	string value	Default group type to use
<b>default_log_levels = [u'amqp=WARN', u'amqplib=WARN', u'boto=WARN', u'qpid=WARN', u'sqlalchemy=WARN', u'suds=INFO', u'oslo.messaging=INFO', u'oslo_messaging=INFO', u'iso8601=WARN', u'requests.packages.urlli b3.connectionpool=WAR N', u'urllib3.connectionpool= WARN', u'websocket=WARN', u'requests.packages.urlli b3.util.retry=WARN', u'urllib3.util.retry=WARN', u'keystonemiddleware=W ARN', u'routes.middleware=WA RN', u'stevedore=WARN', u'taskflow=WARN', u'keystoneauth=WARN', u'oslo.cache=INFO', u'dogpile.core.dogpile=IN FO']</b>	list value	List of package logging levels in logger=LEVEL pairs. This option is ignored if log_config_append is set.
<b>default_volume_type = None</b>	string value	Default volume type to use
<b>driver_client_cert = None</b>	string value	The path to the client certificate for verification, if the driver supports it.
<b>driver_client_cert_key = None</b>	string value	The path to the client certificate key for verification, if the driver supports it.
<b>driver_data_namespace = None</b>	string value	Namespace for driver private data values to be saved in.
<b>driver_ssl_cert_path = None</b>	string value	Can be used to specify a non default path to a CA_BUNDLE file or directory with certificates of trusted CAs, which will be used to validate the backend

Configuration option = Default value	Type	Description
<b>driver_ssl_cert_verify = False</b>	boolean value	If set to True the http client will validate the SSL certificate of the backend endpoint.
<b>driver_use_ssl = False</b>	boolean value	Tell driver to use SSL for connection to backend storage if the driver supports it.
<b>enable_force_upload = False</b>	boolean value	Enables the Force option on upload_to_image. This enables running upload_volume on in-use volumes for backends that support it.
<b>enable_new_services = True</b>	boolean value	Services to be added to the available pool on create
<b>enable_unsupported_driver = False</b>	boolean value	Set this to True when you want to allow an unsupported driver to start. Drivers that haven't maintained a working CI system and testing are marked as unsupported until CI is working again. This also marks a driver as deprecated and may be removed in the next release.
<b>enable_v2_api = True</b>	boolean value	DEPRECATED: Deploy v2 of the Cinder API.
<b>enable_v3_api = True</b>	boolean value	Deploy v3 of the Cinder API.
<b>enabled_backends = None</b>	list value	A list of backend names to use. These backend names should be backed by a unique [CONFIG] group with its options
<b>enforce_multipath_for_image_xfer = False</b>	boolean value	If this is set to True, attachment of volumes for image transfer will be aborted when multipathd is not running. Otherwise, it will fallback to single path.
<b>executor_thread_pool_size = 64</b>	integer value	Size of executor thread pool when executor is threading or eventlet.
<b>fatal_deprecations = False</b>	boolean value	Enables or disables fatal status of deprecations.
<b>fatal_exception_format_errors = False</b>	boolean value	Make exception message format errors fatal.
<b>filter_function = None</b>	string value	String representation for an equation that will be used to filter hosts. Only used when the driver filter is set to be used by the Cinder scheduler.

Configuration option = Default value	Type	Description
<b>glance_api_insecure = False</b>	boolean value	Allow to perform insecure SSL (https) requests to glance (https will be used but cert validation will not be performed).
<b>glance_api_servers = None</b>	list value	A list of the URLs of glance API servers available to cinder ([http[s]://][hostname ip]:port). If protocol is not specified it defaults to http.
<b>glance_api_ssl_compression = False</b>	boolean value	Enables or disables negotiation of SSL layer compression. In some cases disabling compression can improve data throughput, such as when high network bandwidth is available and you use compressed image formats like qcow2.
<b>glance_ca_certificates_file = None</b>	string value	Location of ca certificates file to use for glance client requests.
<b>glance_catalog_info = image:glance:publicURL</b>	string value	Info to match when looking for glance in the service catalog. Format is: separated values of the form: <service_type>:<service_name>:<endpoint_type> - Only used if glance_api_servers are not provided.
<b>glance_core_properties = [u'checksum', u'container_format', u'disk_format', u'image_name', u'image_id', u'min_disk', u'min_ram', u'name', u'size']</b>	list value	Default core properties of image
<b>glance_num_retries = 0</b>	integer value	Number retries when downloading an image from glance
<b>glance_request_timeout = None</b>	integer value	http/https timeout value for glance operations. If no value (None) is supplied here, the glanceclient default value is used.
<b>glusterfs_backup_mount_point = \$state_path/backup_mount</b>	string value	Base dir containing mount point for gluster share.
<b>glusterfs_backup_share = None</b>	string value	GlusterFS share in <hostname ipv4addr ipv6addr>:<gluster_vol_name> format. Eg: 1.2.3.4:backup_vol

Configuration option = Default value	Type	Description
<b>goodness_function = None</b>	string value	String representation for an equation that will be used to determine the goodness of a host. Only used when using the goodness weigher is set to be used by the Cinder scheduler.
<b>graceful_shutdown_timeo ut = 60</b>	integer value	Specify a timeout after which a gracefully shutdown server will exit. Zero value means endless wait.
<b>group_api_class = cinder.group.api.API</b>	string value	The full class name of the group API class
<b>host = ccosp- undercloud.usersys.redh at.com</b>	unknown value	Name of this node. This can be an opaque identifier. It is not necessarily a host name, FQDN, or IP address.
<b>iet_conf = /etc/iet/ietd.conf</b>	string value	IET configuration file
<b>image_conversion_dir = \$state_path/conversion</b>	string value	Directory used for temporary storage during image conversion
<b>image_upload_use_cinde r_backend = False</b>	boolean value	If set to True, upload-to-image in raw format will create a cloned volume and register its location to the image service, instead of uploading the volume content. The cinder backend and locations support must be enabled in the image service.
<b>image_upload_use_intern al_tenant = False</b>	boolean value	If set to True, the image volume created by upload-to-image will be placed in the internal tenant. Otherwise, the image volume is created in the current context's tenant.
<b>image_volume_cache_en abled = False</b>	boolean value	Enable the image volume cache for this backend.
<b>image_volume_cache_ma x_count = 0</b>	integer value	Max number of entries allowed in the image volume cache. 0 ⇒ unlimited.
<b>image_volume_cache_ma x_size_gb = 0</b>	integer value	Max size of the image volume cache for this backend in GB. 0 ⇒ unlimited.
<code>`instance_format = [instance: %(uuid)s] `</code>	string value	The format for an instance that is passed with the log message.
<code>`instance_uuid_format = [instance: %(uuid)s] `</code>	string value	The format for an instance UUID that is passed with the log message.

Configuration option = Default value	Type	Description
<b>instorage_mcs_allow_tenant_qos = False</b>	boolean value	Allow tenants to specify QOS on create
<b>instorage_mcs_iscsi_chap_enabled = True</b>	boolean value	Configure CHAP authentication for iSCSI connections (Default: Enabled)
<b>instorage_mcs_localcopy_rate = 50</b>	integer value	Specifies the InStorage LocalCopy copy rate to be used when creating a full volume copy. The default is rate is 50, and the valid rates are 1-100.
<b>instorage_mcs_localcopy_timeout = 120</b>	integer value	Maximum number of seconds to wait for LocalCopy to be prepared.
<b>instorage_mcs_vol_autoexpand = True</b>	boolean value	Storage system autoexpand parameter for volumes (True/False)
<b>instorage_mcs_vol_compression = False</b>	boolean value	Storage system compression option for volumes
<b>instorage_mcs_vol_grain_size = 256</b>	integer value	Storage system grain size parameter for volumes (32/64/128/256)
<b>instorage_mcs_vol_intier = True</b>	boolean value	Enable InTier for volumes
<b>instorage_mcs_vol_iogrp = 0</b>	string value	The I/O group in which to allocate volumes. It can be a comma-separated list in which case the driver will select an io_group based on least number of volumes associated with the io_group.
<b>instorage_mcs_vol_rsize = 2</b>	integer value	Storage system space-efficiency parameter for volumes (percentage)
<b>instorage_mcs_vol_warning = 0</b>	integer value	Storage system threshold for volume capacity warnings (percentage)
<b>instorage_mcs_volpool_name = [u'volpool']</b>	list value	Comma separated list of storage system storage pools for volumes.
<b>instorage_san_secondary_ip = None</b>	string value	Specifies secondary management IP or hostname to be used if san_ip is invalid or becomes inaccessible.
<b>iscsi_iotype = fileio</b>	string value	Sets the behavior of the iSCSI target to either perform blockio or fileio optionally, auto can be set and Cinder will autodetect type of backing device



Configuration option = Default value	Type	Description
<b>iscsi_secondary_ip_addresses = []</b>	list value	The list of secondary IP addresses of the iSCSI daemon
<code>`iscsi_target_flags = `</code>	string value	Sets the target-specific flags for the iSCSI target. Only used for tgtadm to specify backing device flags using bsoflags option. The specified string is passed as is to the underlying tool.
<b>iscsi_write_cache = on</b>	string value	Sets the behavior of the iSCSI target to either perform write-back(on) or write-through(off). This parameter is valid if target_helper is set to tgtadm.
<b>iser_helper = tgtadm</b>	string value	The name of the iSER target user-land tool to use
<b>iser_ip_address = \$my_ip</b>	string value	The IP address that the iSER daemon is listening on
<b>iser_port = 3260</b>	port value	The port that the iSER daemon is listening on
<b>iser_target_prefix = iqn.2010-10.org.openstack:</b>	string value	Prefix for iSER volumes
<b>keystone_catalog_info = identity:Identity Service:publicURL</b>	string value	Info to match when looking for keystone in the service catalog. Format is: separated values of the form: <service_type>:<service_name>:<endpoint_type> - Only used if backup_swift_auth_url is unset
<b>log-config-append = None</b>	string value	The name of a logging configuration file. This file is appended to any existing logging configuration files. For details about logging configuration files, see the Python logging module documentation. Note that when logging configuration files are used then all logging configuration is set in the configuration file and other logging configuration options are ignored (for example, logging_context_format_string).
<b>log-date-format = %Y-%m-%d %H:%M:%S</b>	string value	Defines the format string for %(asctime)s in log records. Default: %(default)s . This option is ignored if log_config_append is set.
<b>log-dir = None</b>	string value	(Optional) The base directory used for relative log_file paths. This option is ignored if log_config_append is set.

Configuration option = Default value	Type	Description
<b>log-file = None</b>	string value	(Optional) Name of log file to send logging output to. If no default is set, logging will go to stderr as defined by use_stderr. This option is ignored if log_config_append is set.
<b>log_options = True</b>	boolean value	Enables or disables logging values of all registered options when starting a service (at DEBUG level).
<b>logging_context_format_string = %(asctime)s.% (msecs)03d %(process)d %(levelname)s %(name)s [%(request_id)s % (user_identity)s] % (instance)s%(message)s</b>	string value	Format string to use for log messages with context.
<b>logging_debug_format_s uffix = %(funcName)s % (pathname)s:%(lineno)d</b>	string value	Additional data to append to log message when logging level for the message is DEBUG.
<b>logging_default_format_s tring = %(asctime)s.% (msecs)03d %(process)d %(levelname)s %(name)s [-] %(instance)s% (message)s</b>	string value	Format string to use for log messages when context is undefined.
<b>logging_exception_prefix = %(asctime)s.% (msecs)03d %(process)d ERROR %(name)s % (instance)s</b>	string value	Prefix each line of exception output with this format.
<b>logging_user_identity_for mat = %(user)s % (tenant)s %(domain)s % (user_domain)s % (project_domain)s</b>	string value	Defines the format string for %(user_identity)s that is used in logging_context_format_string.
<b>max_age = 0</b>	integer value	Number of seconds between subsequent usage refreshes
<b>max_header_line = 16384</b>	integer value	Maximum line size of message headers to be accepted. max_header_line may need to be increased when using large tokens (typically those generated when keystone is configured to use PKI tokens with big service catalogs).

Configuration option = Default value	Type	Description
<b>max_over_subscription_ratio = 20.0</b>	string value	Representation of the over subscription ratio when thin provisioning is enabled. Default ratio is 20.0, meaning provisioned capacity can be 20 times of the total physical capacity. If the ratio is 10.5, it means provisioned capacity can be 10.5 times of the total physical capacity. A ratio of 1.0 means provisioned capacity cannot exceed the total physical capacity. If ratio is <i>auto</i> , Cinder will automatically calculate the ratio based on the provisioned capacity and the used space. If not set to <i>auto</i> , the ratio has to be a minimum of 1.0.
<b>message_reap_interval = 86400</b>	integer value	interval between periodic task runs to clean expired messages in seconds.
<b>message_ttl = 2592000</b>	integer value	message minimum life in seconds.
<b>migration_create_volume_timeout_secs = 300</b>	integer value	Timeout for creating the volume to migrate to when performing volume migration (seconds)
<b>monkey_patch = False</b>	boolean value	Enable monkey patching
<b>monkey_patch_modules = []</b>	list value	List of modules/decorators to monkey patch
<b>my_ip = 10.65.176.99</b>	unknown value	IP address of this host
<b>no_snapshot_gb_quota = False</b>	boolean value	Whether snapshots count against gigabyte quota
<b>num_iscsi_scan_tries = 3</b>	integer value	The maximum number of times to rescan iSER target to find volume
<b>num_shell_tries = 3</b>	integer value	Number of times to attempt to run flakey shell commands
<b>num_volume_device_scan_tries = 3</b>	integer value	The maximum number of times to rescan targets to find volume
<b>osapi_max_limit = 1000</b>	integer value	The maximum number of items that a collection resource returns in a single response
<b>osapi_volume_ext_list = []</b>	list value	Specify list of extensions to load when using <code>osapi_volume_extension</code> option with <code>cinder.api.contrib.select_extensions</code>

Configuration option = Default value	Type	Description
<b>osapi_volume_extension =</b> [u'cinder.api.contrib.standard_extensions']	multi valued	osapi volume extension to load
<b>osapi_volume_listen =</b> 0.0.0.0	string value	IP address on which OpenStack Volume API listens
<b>osapi_volume_listen_port =</b> 8776	port value	Port on which OpenStack Volume API listens
<b>osapi_volume_use_ssl =</b> False	boolean value	Wraps the socket in a SSL context if True is set. A certificate file and key file must be specified.
<b>osapi_volume_workers =</b> None	integer value	Number of workers for OpenStack Volume API service. The default is equal to the number of CPUs available.
<b>per_volume_size_limit =</b> 1	integer value	Max size allowed per volume, in gigabytes
<b>periodic_fuzzy_delay =</b> 60	integer value	Range, in seconds, to randomly delay when starting the periodic task scheduler to reduce stampeding. (Disable by setting to 0)
<b>periodic_interval =</b> 60	integer value	Interval, in seconds, between running periodic tasks
<b>public_endpoint =</b> None	string value	Public url to use for versions endpoint. The default is None, which will use the request's host_url attribute to populate the URL base. If Cinder is operating behind a proxy, you will want to change this to represent the proxy's URL.
<b>publish_errors =</b> False	boolean value	Enables or disables publication of error events.
<b>query_volume_filters =</b> [u'name', u'status', u'metadata', u'availability_zone', u'bootable', u'group_id']	list value	Volume filter options which non-admin user could use to query volumes. Default values are: [ <i>name, status, metadata, availability_zone, bootable, group_id</i> ]
<b>quota_backup_gigabytes =</b> 1000	integer value	Total amount of storage, in gigabytes, allowed for backups per project
<b>quota_backups =</b> 10	integer value	Number of volume backups allowed per project

Configuration option = Default value	Type	Description
<b>quota_consistencygroups = 10</b>	integer value	Number of consistencygroups allowed per project
<b>quota_driver = cinder.quota.DbQuotaDriver</b>	string value	Default driver to use for quota checks
<b>quota_gigabytes = 1000</b>	integer value	Total amount of storage, in gigabytes, allowed for volumes and snapshots per project
<b>quota_groups = 10</b>	integer value	Number of groups allowed per project
<b>quota_snapshots = 10</b>	integer value	Number of volume snapshots allowed per project
<b>quota_volumes = 10</b>	integer value	Number of volumes allowed per project
<b>rate_limit_burst = 0</b>	integer value	Maximum number of logged messages per rate_limit_interval.
<b>rate_limit_except_level = CRITICAL</b>	string value	Log level name used by rate limiting: CRITICAL, ERROR, INFO, WARNING, DEBUG or empty string. Logs with level greater or equal to rate_limit_except_level are not filtered. An empty string means that all levels are filtered.
<b>rate_limit_interval = 0</b>	integer value	Interval, number of seconds, of log rate limiting.
<b>replication_device = None</b>	dict value	Multi opt of dictionaries to represent a replication target device. This option may be specified multiple times in a single config section to specify multiple replication target devices. Each entry takes the standard dict config form: replication_device = target_device_id:<required>,key1:value1,key2:value2...
<b>report_discard_supported = False</b>	boolean value	Report to clients of Cinder that the backend supports discard (aka. trim/unmap). This will not actually change the behavior of the backend or the client directly, it will only notify that it can be used.
<b>report_interval = 10</b>	integer value	Interval, in seconds, between nodes reporting state to datastore
<b>reservation_clean_interval = \$reservation_expire</b>	integer value	Interval between periodic task runs to clean expired reservations in seconds.
<b>reservation_expire = 86400</b>	integer value	Number of seconds until a reservation expires

Configuration option = Default value	Type	Description
<b>reserved_percentage = 0</b>	integer value	The percentage of backend capacity is reserved
<b>resource_query_filters_file = /etc/cinder/resource_filters.json</b>	string value	Json file indicating user visible filter parameters for list queries.
<b>restore_discard_excess_bytes = True</b>	boolean value	If True, always discard excess bytes when restoring volumes i.e. pad with zeroes.
<b>rootwrap_config = /etc/cinder/rootwrap.conf</b>	string value	Path to the rootwrap configuration file to use for running commands as root
<b>rpc_ack_timeout_base = 15</b>	integer value	Number of seconds to wait for an ack from a cast/call. After each retry attempt this timeout is multiplied by some specified multiplier.
<b>rpc_ack_timeout_multiplier = 2</b>	integer value	Number to multiply base ack timeout by after each retry attempt.
<b>rpc_backend = rabbit</b>	string value	The messaging driver to use, defaults to rabbit. Other drivers include amqp and zmq.
<b>rpc_conn_pool_size = 30</b>	integer value	Size of RPC connection pool.
<b>rpc_message_ttl = 300</b>	integer value	Expiration timeout in seconds of a sent/received message after which it is not tracked anymore by a client/server.
<b>rpc_poll_timeout = 1</b>	integer value	The default number of seconds that poll should wait. Poll raises timeout exception when timeout expired.
<b>rpc_response_timeout = 60</b>	integer value	Seconds to wait for a response from a call.
<b>rpc_retry_attempts = 3</b>	integer value	Default number of message sending attempts in case of any problems occurred: positive value N means at most N retries, 0 means no retries, None or -1 (or any other negative values) mean to retry forever. This option is used only if acknowledgments are enabled.
<b>rpc_thread_pool_size = 100</b>	integer value	Maximum number of (green) threads to work concurrently.

Configuration option = Default value	Type	Description
<b>rpc_use_acks = False</b>	boolean value	Wait for message acknowledgements from receivers. This mechanism works only via proxy without PUB/SUB.
<b>rpc_zmq_bind_address = *</b>	string value	ZeroMQ bind address. Should be a wildcard (*), an ethernet interface, or IP. The "host" option should point or resolve to this address.
<b>rpc_zmq_bind_port_retries = 100</b>	integer value	Number of retries to find free port number before fail with ZMQBindError.
<b>rpc_zmq_contexts = 1</b>	integer value	Number of ZeroMQ contexts, defaults to 1.
<b>rpc_zmq_host = ccosp-undercloud.usersys.redhat.com</b>	string value	Name of this node. Must be a valid hostname, FQDN, or IP address. Must match "host" option, if running Nova.
<b>rpc_zmq_ipc_dir = /var/run/openstack</b>	string value	Directory for holding IPC sockets.
<b>rpc_zmq_matchmaker = redis</b>	string value	MatchMaker driver.
<b>rpc_zmq_max_port = 65536</b>	integer value	Maximal port number for random ports range.
<b>rpc_zmq_min_port = 49153</b>	port value	Minimal port number for random ports range.
<b>rpc_zmq_serialization = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>rpc_zmq_topic_backlog = None</b>	integer value	Maximum number of ingress messages to locally buffer per topic. Default is unlimited.
<b>run_external_periodic_tasks = True</b>	boolean value	Some periodic tasks can be run in a separate process. Should we run them here?
<b>scheduler_default_filters = [u'AvailabilityZoneFilter', u'CapacityFilter', u'CapabilitiesFilter']</b>	list value	Which filter class names to use for filtering hosts when not specified in the request.

Configuration option = Default value	Type	Description
<b>scheduler_default_weighters = [u'CapacityWeigher']</b>	list value	Which weigher class names to use for weighing hosts.
<b>scheduler_driver = cinder.scheduler.filter_scheduler.FilterScheduler</b>	string value	Default scheduler driver to use
<b>scheduler_host_manager = cinder.scheduler.host_manager.HostManager</b>	string value	The scheduler host manager class to use
<code>^scheduler_json_config_location = ^</code>	string value	Absolute path to scheduler configuration JSON file.
<b>scheduler_manager = cinder.scheduler.manager.SchedulerManager</b>	string value	Full class name for the Manager for scheduler
<b>scheduler_max_attempts = 3</b>	integer value	Maximum number of attempts to schedule a volume
<b>scheduler_weight_handler = cinder.scheduler.weights.OrderedHostWeightHandler</b>	string value	Which handler to use for selecting the host/pool after weighing
<b>scst_target_driver = iscsi</b>	string value	SCST target implementation can choose from multiple SCST target drivers.
<b>scst_target_iqn_name = None</b>	string value	Certain ISCSI targets have predefined target names, SCST target driver uses this name.
<b>service_down_time = 60</b>	integer value	Maximum time since last check-in for a service to be considered up
<b>snapshot_name_template = snapshot-%s</b>	string value	Template string to be used to generate snapshot names
<b>snapshot_same_host = True</b>	boolean value	Create volume from snapshot at the host where snapshot resides
<b>ssh_hosts_key_file = \$state_path/ssh_known_hosts</b>	string value	File containing SSH host keys for the systems with which Cinder needs to communicate. OPTIONAL: Default=\$state_path/ssh_known_hosts



Configuration option = Default value	Type	Description
<b>state_path = /var/lib/cinder</b>	string value	Top-level directory for maintaining cinder's state
<b>storage_availability_zone = nova</b>	string value	Availability zone of this node. Can be overridden per volume backend with the option "backend_availability_zone".
<b>storage_protocol = iscsi</b>	string value	Protocol for transferring data between host and storage back-end.
<b>storpool_replication = 3</b>	integer value	The default StorPool chain replication value. Used when creating a volume with no specified type if storpool_template is not set. Also used for calculating the apparent free space reported in the stats.
<b>storpool_template = None</b>	string value	The StorPool template for volumes with no type.
<b>strict_ssh_host_key_policy = False</b>	boolean value	Option to enable strict host key checking. When set to "True" Cinder will only connect to systems with a host key present in the configured "ssh_hosts_key_file". When set to "False" the host key will be saved upon first connection and used for subsequent connections. Default=False
<b>subscribe_on = []</b>	list value	List of publisher hosts SubConsumer can subscribe on. This option has higher priority then the default publishers list taken from the matchmaker.
<b>swift_catalog_info = object-store:swift:publicURL</b>	string value	Info to match when looking for swift in the service catalog. Format is: separated values of the form: <service_type>:<service_name>:<endpoint_type> - Only used if backup_swift_url is unset
<b>syslog-log-facility = LOG_USER</b>	string value	Syslog facility to receive log lines. This option is ignored if log_config_append is set.
<b>target_helper = tgtadm</b>	string value	iSCSI target user-land tool to use. tgtadm is default, use lioadm for LIO iSCSI support, scstadmin for SCST target support, ietadm for iSCSI Enterprise Target, iscsictl for Chelsio iSCSI Target or fake for testing.
<b>target_ip_address = \$my_ip</b>	string value	The IP address that the iSCSI daemon is listening on

Configuration option = Default value	Type	Description
<b>target_port = 3260</b>	port value	The port that the iSCSI daemon is listening on
<b>target_prefix = iqn.2010-10.org.openstack:</b>	string value	Prefix for iSCSI volumes
<b>target_protocol = iscsi</b>	string value	Determines the iSCSI protocol for new iSCSI volumes, created with tgtadm or lioadm target helpers. In order to enable RDMA, this parameter should be set with the value "iser". The supported iSCSI protocol values are "iscsi" and "iser".
<b>tcp_keepalive = True</b>	boolean value	Sets the value of TCP_KEEPALIVE (True/False) for each server socket.
<b>tcp_keepalive_count = None</b>	integer value	Sets the value of TCP_KEEPCNT for each server socket. Not supported on OS X.
<b>tcp_keepalive_interval = None</b>	integer value	Sets the value of TCP_KEEPINTVL in seconds for each server socket. Not supported on OS X.
<b>tcp_keepidle = 600</b>	integer value	Sets the value of TCP_KEEPIDLE in seconds for each server socket. Not supported on OS X.
<b>trace_flags = None</b>	list value	List of options that control which trace info is written to the DEBUG log level to assist developers. Valid values are method and api.
<b>transfer_api_class = cinder.transfer.api.API</b>	string value	The full class name of the volume transfer API class
<b>transport_url = None</b>	string value	The network address and optional user credentials for connecting to the messaging backend, in URL format. The expected format is:  driver://[user:pass@]host:port[, [userN:passN@]hostN:portN]/virtual_host?query  Example: rabbit://rabbitmq:password@127.0.0.1:5672//  For full details on the fields in the URL see the documentation of oslo_messaging.TransportURL at <a href="https://docs.openstack.org/oslo.messaging/latest/reference/transport.html">https://docs.openstack.org/oslo.messaging/latest/reference/transport.html</a>
<b>until_refresh = 0</b>	integer value	Count of reservations until usage is refreshed

Configuration option = Default value	Type	Description
<b>use-journal = False</b>	boolean value	Enable journald for logging. If running in a systemd environment you may wish to enable journal support. Doing so will use the journal native protocol which includes structured metadata in addition to log messages. This option is ignored if log_config_append is set.
<b>use-json = False</b>	boolean value	Use JSON formatting for logging. This option is ignored if log_config_append is set.
<b>use-syslog = False</b>	boolean value	Use syslog for logging. Existing syslog format is DEPRECATED and will be changed later to honor RFC5424. This option is ignored if log_config_append is set.
<b>use_chap_auth = False</b>	boolean value	Option to enable/disable CHAP authentication for targets.
<b>use_default_quota_class = True</b>	boolean value	Enables or disables use of default quota class with default quota.
<b>use_dynamic_connections = False</b>	boolean value	This option makes direct connections dynamic or static. It makes sense only with use_router_proxy=False which means to use direct connections for direct message types (ignored otherwise).
<b>use_forwarded_for = False</b>	boolean value	Treat X-Forwarded-For as the canonical remote address. Only enable this if you have a sanitizing proxy.
<b>use_multipath_for_image_xfer = False</b>	boolean value	Do we attach/detach volumes in cinder using multipath for volume to image and image to volume transfers?
<b>use_pub_sub = False</b>	boolean value	Use PUB/SUB pattern for fanout methods. PUB/SUB always uses proxy.
<b>use_router_proxy = False</b>	boolean value	Use ROUTER remote proxy.
<b>use_stderr = False</b>	boolean value	Log output to standard error. This option is ignored if log_config_append is set.
<b>volume_api_class = cinder.volume.api.API</b>	string value	The full class name of the volume API class to use

Configuration option = Default value	Type	Description
<b>volume_backend_name = None</b>	string value	The backend name for a given driver implementation
<b>volume_clear = zero</b>	string value	Method used to wipe old volumes
<b>volume_clear_ionice = None</b>	string value	The flag to pass to ionice to alter the i/o priority of the process used to zero a volume after deletion, for example "-c3" for idle only priority.
<b>volume_clear_size = 0</b>	integer value	Size in MiB to wipe at start of old volumes. 1024 MiB at max. 0 ⇒ all
<b>volume_copy_blkio_cgrou p_name = cinder- volume-copy</b>	string value	The blkio cgroup name to be used to limit bandwidth of volume copy
<b>volume_copy_bps_limit = 0</b>	integer value	The upper limit of bandwidth of volume copy. 0 ⇒ unlimited
<b>volume_dd_blocksize = 1M</b>	string value	The default block size used when copying/clearing volumes
<b>volume_manager = cinder.volume.manager.V olumeManager</b>	string value	Full class name for the Manager for volume
<b>volume_name_template = volume-%s</b>	string value	Template string to be used to generate volume names
<b>volume_number_multiplie r = -1.0</b>	floating point value	Multiplier used for weighing volume number. Negative numbers mean to spread vs stack.
<b>volume_service_inithost_ offload = False</b>	boolean value	Offload pending volume delete during volume service startup
<b>volume_transfer_key_len gth = 16</b>	integer value	The number of characters in the autogenerated auth key.
<b>volume_transfer_salt_len gth = 8</b>	integer value	The number of characters in the salt.
<b>volume_usage_audit_peri od = month</b>	string value	Time period for which to generate volume usages. The options are hour, day, month, or year.

Configuration option = Default value	Type	Description
<b>volumes_dir = \$state_path/volumes</b>	string value	Volume configuration file storage directory
<b>watch-log-file = False</b>	boolean value	Uses logging handler designed to watch file system. When log file is moved or removed this handler will open a new log file with specified path instantaneously. It makes sense only if log_file option is specified and Linux platform is used. This option is ignored if log_config_append is set.
<b>wsgi_default_pool_size = 100</b>	integer value	Size of the pool of greenthreads used by wsgi
<b>wsgi_keep_alive = True</b>	boolean value	If False, closes the client socket connection explicitly.
<b>wsgi_log_format = % (client_ip)s "% (request_line)s" status: % (status_code)s len: % (body_length)s time: % (wall_seconds).7f</b>	string value	A python format string that is used as the template to generate log lines. The following values can beformatted into it: client_ip, date_time, request_line, status_code, body_length, wall_seconds.
<b>zmq_failover_connection s = 2</b>	integer value	How many additional connections to a host will be made for failover reasons. This option is actual only in dynamic connections mode.
<b>zmq_immediate = True</b>	boolean value	This option configures round-robin mode in zmq socket. True means not keeping a queue when server side disconnects. False means to keep queue and messages even if server is disconnected, when the server appears we send all accumulated messages to it.
<b>zmq_linger = -1</b>	integer value	Number of seconds to wait before all pending messages will be sent after closing a socket. The default value of -1 specifies an infinite linger period. The value of 0 specifies no linger period. Pending messages shall be discarded immediately when the socket is closed. Positive values specify an upper bound for the linger period.
<b>zmq_target_expire = 300</b>	integer value	Expiration timeout in seconds of a name service record about existing target (< 0 means no timeout).
<b>zmq_target_update = 180</b>	integer value	Update period in seconds of a name service record about existing target.

Configuration option = Default value	Type	Description
<b>zmq_tcp_keepalive = -1</b>	integer value	Enable/disable TCP keepalive (KA) mechanism. The default value of -1 (or any other negative value) means to skip any overrides and leave it to OS default; 0 and 1 (or any other positive value) mean to disable and enable the option respectively.
<b>zmq_tcp_keepalive_cnt = -1</b>	integer value	The number of retransmissions to be carried out before declaring that remote end is not available. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_idle = -1</b>	integer value	The duration between two keepalive transmissions in idle condition. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_intvl = -1</b>	integer value	The duration between two successive keepalive retransmissions, if acknowledgement to the previous keepalive transmission is not received. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zoning_mode = None</b>	string value	FC Zoning mode configured, only <i>fabric</i> is supported now.

### 3.1.2. backend

The following table outlines the options available under the **[backend]** group in the `/etc/cinder/cinder.conf` file.

Table 3.2. backend

Configuration option = Default value	Type	Description
<b>backend_host = None</b>	string value	Backend override of host value.

### 3.1.3. backend\_defaults

The following table outlines the options available under the **[backend\_defaults]** group in the `/etc/cinder/cinder.conf` file.

Table 3.3. backend\_defaults

Configuration option = Default value	Type	Description
<b>auto_calc_max_oversubscription_ratio = False</b>	boolean value	K2 driver will calculate max_oversubscription_ratio on setting this option as True.
<b>backend_availability_zone = None</b>	string value	Availability zone for this volume backend. If not set, the storage_availability_zone option value is used as the default for all backends.
<b>backend_native_threads_pool_size = 20</b>	integer value	Size of the native threads pool for the backend. Increase for backends that heavily rely on this, like the RBD driver.
<b>backup_use_temp_snapshots = False</b>	boolean value	If this is set to True, a temporary snapshot will be created for performing non-disruptive backups. Otherwise a temporary volume will be cloned in order to perform a backup.
<b>chap = disabled</b>	string value	CHAP authentication mode, effective only for iscsi (disabled enabled)
<code>`chap_password = `</code>	string value	Password for specified CHAP account name.
<code>`chap_username = `</code>	string value	CHAP user name.
<b>check_max_pool_luns_threshold = False</b>	boolean value	Report free_capacity_gb as 0 when the limit to maximum number of pool LUNs is reached. By default, the value is False.
<b>chiscsi_conf = /etc/chelsio-iscsi/chiscsi.conf</b>	string value	Chiscsi (CXT) global defaults configuration file
<b>cinder_dell_emc_config_file = /etc/cinder/cinder_dell_emc_config.xml</b>	string value	Use this file for cinder emc plugin config data.
<b>cinder_eternus_config_file = /etc/cinder/cinder_fujitsu_eternus_dx.xml</b>	string value	config file for cinder eternus_dx volume driver
<b>cinder_huawei_conf_file = /etc/cinder/cinder_huawei_conf.xml</b>	string value	The configuration file for the Cinder Huawei driver.

Configuration option = Default value	Type	Description
<b>clone_volume_timeout = 680</b>	integer value	Create clone volume timeout.
<b>connection_type = iscsi</b>	string value	Connection type to the IBM Storage Array
<b>coprhd_emulate_snapshot = False</b>	boolean value	True   False to indicate if the storage array in CoprHD is VMAX or VPLEX
<b>coprhd_hostname = None</b>	string value	Hostname for the CoprHD Instance
<b>coprhd_password = None</b>	string value	Password for accessing the CoprHD Instance
<b>coprhd_port = 4443</b>	port value	Port for the CoprHD Instance
<b>coprhd_project = None</b>	string value	Project to utilize within the CoprHD Instance
<b>coprhd_scaleio_rest_gateway_host = None</b>	string value	Rest Gateway IP or FQDN for Scaleio
<b>coprhd_scaleio_rest_gateway_port = 4984</b>	port value	Rest Gateway Port for Scaleio
<b>coprhd_scaleio_rest_server_password = None</b>	string value	Rest Gateway Password
<b>coprhd_scaleio_rest_server_username = None</b>	string value	Username for Rest Gateway
<b>coprhd_tenant = None</b>	string value	Tenant to utilize within the CoprHD Instance
<b>coprhd_username = None</b>	string value	Username for accessing the CoprHD Instance
<b>coprhd_varray = None</b>	string value	Virtual Array to utilize within the CoprHD Instance
<b>cycle_period_seconds = 300</b>	integer value	This defines an optional cycle period that applies to Global Mirror relationships with a cycling mode of multi. A Global Mirror relationship using the multi cycling_mode performs a complete cycle at most once each period. The default is 300 seconds, and the valid seconds are 60-86400.
<b>datara_503_interval = 5</b>	integer value	Interval between 503 retries
<b>datara_503_timeout = 120</b>	integer value	Timeout for HTTP 503 retry messages



Configuration option = Default value	Type	Description
<b>datera_api_port = 7717</b>	string value	Datera API port.
<b>datera_api_version = 2</b>	string value	Datera API version.
<b>datera_debug = False</b>	boolean value	True to set function arg and return logging
<b>datera_debug_replica_count_override = False</b>	boolean value	ONLY FOR DEBUG/TESTING PURPOSES True to set replica_count to 1
<b>datera_disable_profiler = False</b>	boolean value	Set to True to disable profiling in the Datera driver
<b>datera_tenant_id = None</b>	string value	If set to <i>Map</i> --> OpenStack project ID will be mapped implicitly to Datera tenant ID If set to <i>None</i> --> Datera tenant ID will not be used during volume provisioning If set to anything else --> Datera tenant ID will be the provided value
<b>default_timeout = 31536000</b>	integer value	Default timeout for CLI operations in minutes. For example, LUN migration is a typical long running operation, which depends on the LUN size and the load of the array. An upper bound in the specific deployment can be set to avoid unnecessary long wait. By default, it is 365 days long.
<b>dell_sc_api_port = 3033</b>	port value	Dell API port
<b>dell_sc_server_folder = openstack</b>	string value	Name of the server folder to use on the Storage Center
<b>dell_sc_ssn = 64702</b>	integer value	Storage Center System Serial Number
<b>dell_sc_verify_cert = False</b>	boolean value	Enable HTTPS SC certificate verification
<b>dell_sc_volume_folder = openstack</b>	string value	Name of the volume folder to use on the Storage Center
<b>dell_server_os = Red Hat Linux 6.x</b>	string value	Server OS type to use when creating a new server on the Storage Center.
<b>destroy_empty_storage_group = False</b>	boolean value	To destroy storage group when the last LUN is removed from it. By default, the value is False.

Configuration option = Default value	Type	Description
<b>disable_discovery = False</b>	boolean value	Disabling iSCSI discovery (sendtargets) for multipath connections on K2 driver.
<b>disco_choice_client = None</b>	string value	Use soap client or rest client for communicating with DISCO. Possible values are "soap" or "rest".
<b>disco_client = 127.0.0.1</b>	IP address value	The IP of DMS client socket server
<b>disco_client_port = 9898</b>	port value	The port to connect DMS client socket server
<b>disco_clone_check_timeout = 3600</b>	integer value	How long we check whether a clone is finished before we give up
<b>disco_rest_ip = None</b>	IP address value	The IP address of the REST server
<b>disco_restore_check_timeout = 3600</b>	integer value	How long we check whether a restore is finished before we give up
<b>disco_retry_interval = 1</b>	integer value	How long we wait before retrying to get an item detail
<b>disco_snapshot_check_timeout = 3600</b>	integer value	How long we check whether a snapshot is finished before we give up
<b>disco_src_api_port = 8080</b>	port value	The port of DISCO source API
<b>disco_volume_name_prefix = openstack-</b>	string value	Prefix before volume name to differentiate DISCO volume created through openstack and the other ones
<b>disco_wsd_path = /etc/cinder/DISCOService.wsd</b>	string value	Path to the wsd file to communicate with DISCO request manager
<b>dpl_pool = `</b>	string value	DPL pool uuid in which DPL volumes are stored.
<b>dpl_port = 8357</b>	port value	DPL port number.
<b>drbdmanage_devs_on_controller = True</b>	boolean value	If set, the c-vol node will receive a useable /dev/drbdX device, even if the actual data is stored on other nodes only. This is useful for debugging, maintenance, and to be able to do the iSCSI export from the c-vol node.

Configuration option = Default value	Type	Description
<b>drbdmanage_disk_options</b> = {"c-min-rate": "4M"}	string value	Disk options to set on new resources. See <a href="http://www.drbd.org/en/doc/users-guide-90/re-drbdconf">http://www.drbd.org/en/doc/users-guide-90/re-drbdconf</a> for all the details.
<b>drbdmanage_net_options</b> = {"connect-int": "4", "allow-two-primaries": "yes", "ko-count": "30", "max-buffers": "20000", "ping-timeout": "100"}	string value	Net options to set on new resources. See <a href="http://www.drbd.org/en/doc/users-guide-90/re-drbdconf">http://www.drbd.org/en/doc/users-guide-90/re-drbdconf</a> for all the details.
<b>drbdmanage_redundancy</b> = 1	integer value	Number of nodes that should replicate the data.
<b>drbdmanage_resize_plugin</b> = drbdmanage.plugins.plugins.wait_for.WaitForVolumeSize	string value	Volume resize completion wait plugin.
<b>drbdmanage_resize_policy</b> = {"timeout": "60"}	string value	Volume resize completion wait policy.
<b>drbdmanage_resource_options</b> = {"auto-promote-timeout": "300"}	string value	Resource options to set on new resources. See <a href="http://www.drbd.org/en/doc/users-guide-90/re-drbdconf">http://www.drbd.org/en/doc/users-guide-90/re-drbdconf</a> for all the details.
<b>drbdmanage_resource_plugin</b> = drbdmanage.plugins.plugins.wait_for.WaitForResource	string value	Resource deployment completion wait plugin.
<b>drbdmanage_resource_policy</b> = {"ratio": "0.51", "timeout": "60"}	string value	Resource deployment completion wait policy.
<b>drbdmanage_snapshot_plugin</b> = drbdmanage.plugins.plugins.wait_for.WaitForSnapshot	string value	Snapshot completion wait plugin.
<b>drbdmanage_snapshot_policy</b> = {"count": "1", "timeout": "60"}	string value	Snapshot completion wait policy.

Configuration option = Default value	Type	Description
<b>driver_client_cert = None</b>	string value	The path to the client certificate for verification, if the driver supports it.
<b>driver_client_cert_key = None</b>	string value	The path to the client certificate key for verification, if the driver supports it.
<b>driver_data_namespace = None</b>	string value	Namespace for driver private data values to be saved in.
<b>driver_ssl_cert_path = None</b>	string value	Can be used to specify a non default path to a CA_BUNDLE file or directory with certificates of trusted CAs, which will be used to validate the backend
<b>driver_ssl_cert_verify = False</b>	boolean value	If set to True the http client will validate the SSL certificate of the backend endpoint.
<b>driver_use_ssl = False</b>	boolean value	Tell driver to use SSL for connection to backend storage if the driver supports it.
<code>`ds8k_devadd_unitadd_mapping = `</code>	string value	Mapping between IODevice address and unit address.
<b>ds8k_host_type = auto</b>	string value	Set to zLinux if your OpenStack version is prior to Liberty and you're connecting to zLinux systems. Otherwise set to auto. Valid values for this parameter are: <i>auto, AMDLinuxRHEL, AMDLinuxSuse, AppleOSX, Fujitsu, Hp, HpTru64, HpVms, LinuxDT, LinuxRF, LinuxRHEL, LinuxSuse, Novell, SGI, SVC, SanFsAIX, SanFsLinux, Sun, VMWare, Win2000, Win2003, Win2008, Win2012, iLinux, nSeries, pLinux, pSeries, pSeriesPowerswap, zLinux, iSeries.</i>
<b>ds8k_ssid_prefix = FF</b>	string value	Set the first two digits of SSID.
<b>dsware_isthin = False</b>	boolean value	The flag of thin storage allocation.
<code>`dsware_manager = `</code>	string value	Fusionstorage manager ip addr for cinder-volume.
<b>enable_unsupported_driver = False</b>	boolean value	Set this to True when you want to allow an unsupported driver to start. Drivers that haven't maintained a working CI system and testing are marked as unsupported until CI is working again. This also marks a driver as deprecated and may be removed in the next release.

Configuration option = Default value	Type	Description
<b>enforce_multipath_for_image_xfer = False</b>	boolean value	If this is set to True, attachment of volumes for image transfer will be aborted when multipathd is not running. Otherwise, it will fallback to single path.
<b>eqlx_cli_max_retries = 5</b>	integer value	Maximum retry count for reconnection. Default is 5.
<b>eqlx_group_name = group-0</b>	string value	Group name to use for creating volumes. Defaults to "group-0".
<b>eqlx_pool = default</b>	string value	Pool in which volumes will be created. Defaults to "default".
<b>excluded_domain_ip = None</b>	IP address value	Domain IP to be excluded from iSCSI returns.
<b>expiry_thres_minutes = 720</b>	integer value	This option specifies the threshold for last access time for images in the NFS image cache. When a cache cleaning cycle begins, images in the cache that have not been accessed in the last M minutes, where M is the value of this parameter, will be deleted from the cache to create free space on the NFS share.
<b>extra_capabilities = {}</b>	string value	User defined capabilities, a JSON formatted string specifying key/value pairs. The key/value pairs can be used by the CapabilitiesFilter to select between backends when requests specify volume types. For example, specifying a service level or the geographical location of a backend, then creating a volume type to allow the user to select by these different properties.
<b>filter_function = None</b>	string value	String representation for an equation that will be used to filter hosts. Only used when the driver filter is set to be used by the Cinder scheduler.
<b>flashsystem_connection_protocol = FC</b>	string value	Connection protocol should be FC. (Default is FC.)
<b>flashsystem_iscsi_portid = 0</b>	integer value	Default iSCSI Port ID of FlashSystem. (Default port is 0.)
<b>flashsystem_multihostmap_enabled = True</b>	boolean value	Allows vdisk to multi host mapping. (Default is True)
<b>flashsystem_multipath_enabled = False</b>	boolean value	This option no longer has any affect. It is deprecated and will be removed in the next release.

Configuration option = Default value	Type	Description
<b>force_delete_lun_in_storage_group = False</b>	boolean value	Delete a LUN even if it is in Storage Groups. By default, the value is False.
<code>`fusionstorageagent = `</code>	string value	Fusionstorage agent ip addr range.
<b>goodness_function = None</b>	string value	String representation for an equation that will be used to determine the goodness of a host. Only used when using the goodness weigher is set to be used by the Cinder scheduler.
<b>gpfs_hosts = []</b>	list value	Comma-separated list of IP address or hostnames of GPFS nodes.
<b>gpfs_hosts_key_file = \$state_path/ssh_known_hosts</b>	string value	File containing SSH host keys for the gpfs nodes with which driver needs to communicate. Default=\$state_path/ssh_known_hosts
<b>gpfs_images_dir = None</b>	string value	Specifies the path of the Image service repository in GPFS. Leave undefined if not storing images in GPFS.
<b>gpfs_images_share_mode = None</b>	string value	Specifies the type of image copy to be used. Set this when the Image service repository also uses GPFS so that image files can be transferred efficiently from the Image service to the Block Storage service. There are two valid values: "copy" specifies that a full copy of the image is made; "copy_on_write" specifies that copy-on-write optimization strategy is used and unmodified blocks of the image file are shared efficiently.
<b>gpfs_max_clone_depth = 0</b>	integer value	Specifies an upper limit on the number of indirections required to reach a specific block due to snapshots or clones. A lengthy chain of copy-on-write snapshots or clones can have a negative impact on performance, but improves space utilization. 0 indicates unlimited clone depth.
<b>gpfs_mount_point_base = None</b>	string value	Specifies the path of the GPFS directory where Block Storage volume and snapshot files are stored.
<code>`gpfs_private_key = `</code>	string value	Filename of private key to use for SSH authentication.
<b>gpfs_sparse_volumes = True</b>	boolean value	Specifies that volumes are created as sparse files which initially consume no space. If set to False, the volume is created as a fully allocated file, in which case, creation may take a significantly longer time.

Configuration option = Default value	Type	Description
<b>gpfs_ssh_port = 22</b>	port value	SSH port to use.
<b>gpfs_storage_pool = system</b>	string value	Specifies the storage pool that volumes are assigned to. By default, the system storage pool is used.
<b>gpfs_strict_host_key_policy = False</b>	boolean value	Option to enable strict gpfs host key checking while connecting to gpfs nodes. Default=False
<b>gpfs_user_login = root</b>	string value	Username for GPFS nodes.
<code>`gpfs_user_password = `</code>	string value	Password for GPFS node user.
<b>hgst_net = Net 1 (IPv4)</b>	string value	Space network name to use for data transfer
<b>hgst_redundancy = 0</b>	string value	Should spaces be redundantly stored (1/0)
<b>hgst_space_group = disk</b>	string value	Group to own created spaces
<b>hgst_space_mode = 0600</b>	string value	UNIX mode for created spaces
<b>hgst_space_user = root</b>	string value	User to own created spaces
<b>hgst_storage_servers = os:gbd0</b>	string value	Comma separated list of Space storage servers:devices. ex: os1_stor:gbd0,os2_stor:gbd0
<code>`hpe3par_api_url = `</code>	string value	3PAR WSAPI Server Url like <a href="https://&lt;3par ip&gt;:8080/api/v1">https://&lt;3par ip&gt;:8080/api/v1</a>
<b>hpe3par_cpg = [u'OpenStack']</b>	list value	List of the CPG(s) to use for volume creation
<code>`hpe3par_cpg_snap = `</code>	string value	The CPG to use for Snapshots for volumes. If empty the userCPG will be used.
<b>hpe3par_debug = False</b>	boolean value	Enable HTTP debugging to 3PAR
<b>hpe3par_iscsi_chap_enabled = False</b>	boolean value	Enable CHAP authentication for iSCSI connections.
<b>hpe3par_iscsi_ips = []</b>	list value	List of target iSCSI addresses to use.
<code>`hpe3par_password = `</code>	string value	3PAR password for the user specified in hpe3par_username

Configuration option = Default value	Type	Description
<code>`hpe3par_snapshot_expiration = `</code>	string value	The time in hours when a snapshot expires and is deleted. This must be larger than expiration
<code>`hpe3par_snapshot_retention = `</code>	string value	The time in hours to retain a snapshot. You can't delete it before this expires.
<code>`hpe3par_username = `</code>	string value	3PAR username with the <i>edit</i> role
<b>hpelefthand_api_url = None</b>	uri value	HPE LeftHand WSAPI Server Url like <a href="https://&lt;LeftHand ip&gt;:8081/lhos">https://&lt;LeftHand ip&gt;:8081/lhos</a>
<b>hpelefthand_clustername = None</b>	string value	HPE LeftHand cluster name
<b>hpelefthand_debug = False</b>	boolean value	Enable HTTP debugging to LeftHand
<b>hpelefthand_iscsi_chap_enabled = False</b>	boolean value	Configure CHAP authentication for iSCSI connections (Default: Disabled)
<b>hpelefthand_password = None</b>	string value	HPE LeftHand Super user password
<b>hpelefthand_ssh_port = 16022</b>	port value	Port number of SSH service.
<b>hpelefthand_username = None</b>	string value	HPE LeftHand Super user username
<b>hpmsa_api_protocol = https</b>	string value	HPMSA API interface protocol.
<b>hpmsa_backend_name = A</b>	string value	Pool or Vdisk name to use for volume creation.
<b>hpmsa_backend_type = virtual</b>	string value	linear (for Vdisk) or virtual (for Pool).
<b>hpmsa_iscsi_ips = []</b>	list value	List of comma-separated target iSCSI IP addresses.
<b>hpmsa_verify_certificate = False</b>	boolean value	Whether to verify HPMSA array SSL certificate.
<b>hpmsa_verify_certificate_path = None</b>	string value	HPMSA array SSL certificate path.



Configuration option = Default value	Type	Description
<b>hypermetro_devices = None</b>	string value	The remote device hypermetro will use.
<b>iet_conf = /etc/iet/ietd.conf</b>	string value	IET configuration file
<b>ignore_pool_full_threshold = False</b>	boolean value	Force LUN creation even if the full threshold of pool is reached. By default, the value is False.
<b>image_upload_use_cinder_backend = False</b>	boolean value	If set to True, upload-to-image in raw format will create a cloned volume and register its location to the image service, instead of uploading the volume content. The cinder backend and locations support must be enabled in the image service.
<b>image_upload_use_internal_tenant = False</b>	boolean value	If set to True, the image volume created by upload-to-image will be placed in the internal tenant. Otherwise, the image volume is created in the current context's tenant.
<b>image_volume_cache_enabled = False</b>	boolean value	Enable the image volume cache for this backend.
<b>image_volume_cache_max_count = 0</b>	integer value	Max number of entries allowed in the image volume cache. 0 ⇒ unlimited.
<b>image_volume_cache_max_size_gb = 0</b>	integer value	Max size of the image volume cache for this backend in GB. 0 ⇒ unlimited.
<b>infinidat_iscsi_netspaces = []</b>	list value	List of names of network spaces to use for iSCSI connectivity
<b>infinidat_pool_name = None</b>	string value	Name of the pool from which volumes are allocated
<b>infinidat_storage_protocol = fc</b>	string value	Protocol for transferring data between host and storage back-end.
<b>infinidat_use_compression = False</b>	boolean value	Specifies whether to turn on compression for newly created volumes.
<b>initiator_auto_deregistration = False</b>	boolean value	Automatically deregister initiators after the related storage group is destroyed. By default, the value is False.
<b>initiator_auto_registration = False</b>	boolean value	Automatically register initiators. By default, the value is False.

Configuration option = Default value	Type	Description
<b>initiator_check = False</b>	boolean value	Use this value to enable the initiator_check.
<b>interval = 3</b>	integer value	Use this value to specify length of the interval in seconds.
<b>io_port_list = None</b>	list value	Comma separated iSCSI or FC ports to be used in Nova or Cinder.
<b>iscsi_initiators = None</b>	string value	Mapping between hostname and its iSCSI initiator IP addresses.
<b>iscsi_iotype = fileio</b>	string value	Sets the behavior of the iSCSI target to either perform blockio or fileio optionally, auto can be set and Cinder will autodetect type of backing device
<b>iscsi_secondary_ip_addresses = []</b>	list value	The list of secondary IP addresses of the iSCSI daemon
<code>`iscsi_target_flags = `</code>	string value	Sets the target-specific flags for the iSCSI target. Only used for tgtadm to specify backing device flags using bsflags option. The specified string is passed as is to the underlying tool.
<b>iscsi_write_cache = on</b>	string value	Sets the behavior of the iSCSI target to either perform write-back(on) or write-through(off). This parameter is valid if target_helper is set to tgtadm.
<b>iser_helper = tgtadm</b>	string value	The name of the iSER target user-land tool to use
<b>iser_ip_address = \$my_ip</b>	string value	The IP address that the iSER daemon is listening on
<b>iser_port = 3260</b>	port value	The port that the iSER daemon is listening on
<b>iser_target_prefix = iqn.2010- 10.org.openstack:</b>	string value	Prefix for iSER volumes
<b>lenovo_api_protocol = https</b>	string value	Lenovo api interface protocol.
<b>lenovo_backend_name = A</b>	string value	Pool or Vdisk name to use for volume creation.
<b>lenovo_backend_type = virtual</b>	string value	linear (for VDisk) or virtual (for Pool).

Configuration option = Default value	Type	Description
<b>lenovo_iscsi_ips = []</b>	list value	List of comma-separated target iSCSI IP addresses.
<b>lenovo_verify_certificate = False</b>	boolean value	Whether to verify Lenovo array SSL certificate.
<b>lenovo_verify_certificate_path = None</b>	string value	Lenovo array SSL certificate path.
<code>`lss_range_for_cg = `</code>	string value	Reserve LSSs for consistency group.
<b>lvm_conf_file = /etc/cinder/lvm.conf</b>	string value	LVM conf file to use for the LVM driver in Cinder; this setting is ignored if the specified file does not exist (You can also specify <i>None</i> to not use a conf file even if one exists).
<b>lvm_max_over_subscription_ratio = None</b>	floating point value	max_over_subscription_ratio setting for the LVM driver. If set to None (the default), the general max_over_subscription_ratio is used.
<b>lvm_mirrors = 0</b>	integer value	If >0, create LVs with multiple mirrors. Note that this requires lvm_mirrors + 2 PVs with available space
<b>lvm_suppress_fd_warnings = False</b>	boolean value	Suppress leaked file descriptor warnings in LVM commands.
<b>lvm_type = auto</b>	string value	Type of LVM volumes to deploy; (default, thin, or auto). Auto defaults to thin if thin is supported.
<code>`management_ips = `</code>	string value	List of Management IP addresses (separated by commas)
<b>max_luns_per_storage_group = 255</b>	integer value	Default max number of LUNs in a storage group. By default, the value is 255.
<b>max_over_subscription_ratio = 20.0</b>	string value	Representation of the over subscription ratio when thin provisioning is enabled. Default ratio is 20.0, meaning provisioned capacity can be 20 times of the total physical capacity. If the ratio is 10.5, it means provisioned capacity can be 10.5 times of the total physical capacity. A ratio of 1.0 means provisioned capacity cannot exceed the total physical capacity. If ratio is <i>auto</i> , Cinder will automatically calculate the ratio based on the provisioned capacity and the used space. If not set to auto, the ratio has to be a minimum of 1.0.

Configuration option = Default value	Type	Description
<b>metro_domain_name = None</b>	string value	The remote metro device domain name.
<b>metro_san_address = None</b>	string value	The remote metro device request url.
<b>metro_san_password = None</b>	string value	The remote metro device san password.
<b>metro_san_user = None</b>	string value	The remote metro device san user.
<b>metro_storage_pools = None</b>	string value	The remote metro device pool names.
<code>`nas_host = `</code>	string value	IP address or Hostname of NAS system.
<b>nas_login = admin</b>	string value	User name to connect to NAS system.
<b>nas_mount_options = None</b>	string value	Options used to mount the storage backend file system where Cinder volumes are stored.
<code>`nas_password = `</code>	string value	Password to connect to NAS system.
<code>`nas_private_key = `</code>	string value	Filename of private key to use for SSH authentication.
<b>nas_secure_file_operations = auto</b>	string value	Allow network-attached storage systems to operate in a secure environment where root level access is not permitted. If set to False, access is as the root user and insecure. If set to True, access is not as root. If set to auto, a check is done to determine if this is a new installation: True is used if so, otherwise False. Default is auto.
<b>nas_secure_file_permissions = auto</b>	string value	Set more secure file permissions on network-attached storage volume files to restrict broad other/world access. If set to False, volumes are created with open permissions. If set to True, volumes are created with permissions for the cinder user and group (660). If set to auto, a check is done to determine if this is a new installation: True is used if so, otherwise False. Default is auto.

Configuration option = Default value	Type	Description
<code>`nas_share_path = `</code>	string value	Path to the share to use for storing Cinder volumes. For example: <code>"/srv/export1"</code> for an NFS server export available at <code>10.0.5.10:/srv/export1</code> .
<b><code>nas_ssh_port = 22</code></b>	port value	SSH port to use to connect to NAS system.
<b><code>nas_volume_prov_type = thin</code></b>	string value	Provisioning type that will be used when creating volumes.
<b><code>navisecli_path = None</code></b>	string value	Navisecli Path.
<b><code>netapp_api_trace_pattern = (.*)</code></b>	string value	A regular expression to limit the API tracing. This option is honored only if enabling <b>api</b> tracing with the <b>trace_flags</b> option. By default, all APIs will be traced.
<b><code>netapp_controller_ips = None</code></b>	string value	This option is only utilized when the storage family is configured to <code>eseries</code> . This option is used to restrict provisioning to the specified controllers. Specify the value of this option to be a comma separated list of controller hostnames or IP addresses to be used for provisioning.
<b><code>netapp_copyoffload_tool_path = None</code></b>	string value	This option specifies the path of the NetApp copy offload tool binary. Ensure that the binary has execute permissions set which allow the effective user of the <code>cinder-volume</code> process to execute the file.
<b><code>netapp_enable_multiattach = False</code></b>	boolean value	This option specifies whether the driver should allow operations that require multiple attachments to a volume. An example would be live migration of servers that have volumes attached. When enabled, this backend is limited to 256 total volumes in order to guarantee volumes can be accessed by more than one host.
<b><code>netapp_host_type = None</code></b>	string value	This option defines the type of operating system for all initiators that can access a LUN. This information is used when mapping LUNs to individual hosts or groups of hosts.
<b><code>netapp_login = None</code></b>	string value	Administrative user account name used to access the storage system or proxy server.

Configuration option = Default value	Type	Description
<b>netapp_lun_ostype = None</b>	string value	This option defines the type of operating system that will access a LUN exported from Data ONTAP; it is assigned to the LUN at the time it is created.
<b>netapp_lun_space_reservation = enabled</b>	string value	This option determines if storage space is reserved for LUN allocation. If enabled, LUNs are thick provisioned. If space reservation is disabled, storage space is allocated on demand.
<b>netapp_password = None</b>	string value	Password for the administrative user account specified in the netapp_login option.
<b>netapp_pool_name_search_pattern = (.+)</b>	string value	This option is used to restrict provisioning to the specified pools. Specify the value of this option to be a regular expression which will be applied to the names of objects from the storage backend which represent pools in Cinder. This option is only utilized when the storage protocol is configured to use iSCSI or FC.
<b>netapp_replication_aggregate_map = None</b>	dict value	Multi opt of dictionaries to represent the aggregate mapping between source and destination back ends when using whole back end replication. For every source aggregate associated with a cinder pool (NetApp FlexVol), you would need to specify the destination aggregate on the replication target device. A replication target device is configured with the configuration option replication_device. Specify this option as many times as you have replication devices. Each entry takes the standard dict config form: netapp_replication_aggregate_map = backend_id: <name_of_replication_device_section>,src_aggr_name1:dest_aggr_name1,src_aggr_name2:dest_aggr_name2,...
<b>netapp_sa_password = None</b>	string value	Password for the NetApp E-Series storage array.
<b>netapp_server_hostname = None</b>	string value	The hostname (or IP address) for the storage system or proxy server.
<b>netapp_server_port = None</b>	integer value	The TCP port to use for communication with the storage system or proxy server. If not specified, Data ONTAP drivers will use 80 for HTTP and 443 for HTTPS; E-Series will use 8080 for HTTP and 8443 for HTTPS.

Configuration option = Default value	Type	Description
<b>netapp_size_multiplier = 1.2</b>	floating point value	The quantity to be multiplied by the requested volume size to ensure enough space is available on the virtual storage server (Vserver) to fulfill the volume creation request. Note: this option is deprecated and will be removed in favor of "reserved_percentage" in the Mitaka release.
<b>netapp_snapmirror_quiesce_timeout = 3600</b>	integer value	The maximum time in seconds to wait for existing SnapMirror transfers to complete before aborting during a failover.
<b>netapp_storage_family = ontap_cluster</b>	string value	The storage family type used on the storage system; valid values are ontap_cluster for using clustered Data ONTAP, or eseries for using E-Series.
<b>netapp_storage_protocol = None</b>	string value	The storage protocol to be used on the data path with the storage system.
<b>netapp_transport_type = http</b>	string value	The transport protocol used when communicating with the storage system or proxy server.
<b>netapp_vserver = None</b>	string value	This option specifies the virtual storage server (Vserver) name on the storage cluster on which provisioning of block storage volumes should occur.
<b>netapp_webservice_path = /devmgr/v2</b>	string value	This option is used to specify the path to the E-Series proxy application on a proxy server. The value is combined with the value of the netapp_transport_type, netapp_server_hostname, and netapp_server_port options to create the URL used by the driver to connect to the proxy application.
<b>nexenta_blocksize = 4096</b>	integer value	Block size for datasets
<b>nexenta_chunksize = 32768</b>	integer value	NexentaEdge iSCSI LUN object chunk size
<b>nexenta_client_address = `</b>	string value	NexentaEdge iSCSI Gateway client address for non-VIP service
<b>nexenta_dataset_compression = on</b>	string value	Compression value for new ZFS folders.
<b>nexenta_dataset_dedup = off</b>	string value	Deduplication value for new ZFS folders.

Configuration option = Default value	Type	Description
<code>`nexenta_dataset_description` = `</code>	string value	Human-readable description for the folder.
<code>`nexenta_host` = `</code>	string value	IP address of Nexenta SA
<code>`nexenta_iscsi_service` = `</code>	string value	NexentaEdge iSCSI service name
<b>nexenta_iscsi_target_port al_port = 3260</b>	integer value	Nexenta target portal port
<code>`nexenta_lun_container` = `</code>	string value	NexentaEdge logical path of bucket for LUNs
<b>nexenta_mount_point_base = \$state_path/mnt</b>	string value	Base directory that contains NFS share mount points
<b>nexenta_nbd_symlinks_dir = /dev/disk/by-path</b>	string value	NexentaEdge logical path of directory to store symbolic links to NBDs
<b>nexenta_nms_cache_volroot = True</b>	boolean value	If set True cache NexentaStor appliance volroot option value.
<b>nexenta_ns5_blocksize = 32</b>	integer value	Block size for datasets
<b>nexenta_password = nexenta</b>	string value	Password to connect to Nexenta SA
<code>`nexenta_rest_address` = `</code>	string value	IP address of NexentaEdge management REST API endpoint
<b>nexenta_rest_password = nexenta</b>	string value	Password to connect to NexentaEdge
<b>nexenta_rest_port = 0</b>	integer value	HTTP(S) port to connect to Nexenta REST API server. If it is equal zero, 8443 for HTTPS and 8080 for HTTP is used
<b>nexenta_rest_protocol = auto</b>	string value	Use http or https for REST connection (default auto)
<b>nexenta_rest_user = admin</b>	string value	User name to connect to NexentaEdge
<b>nexenta_rrmgr_compression = 0</b>	integer value	Enable stream compression, level 1..9. 1 - gives best speed; 9 - gives best compression.



Configuration option = Default value	Type	Description
<b>nexenta_rrmgr_connections = 2</b>	integer value	Number of TCP connections.
<b>nexenta_rrmgr_tcp_buf_size = 4096</b>	integer value	TCP Buffer size in KiloBytes.
<b>nexenta_shares_config = /etc/cinder/nfs_shares</b>	string value	File with the list of available nfs shares
<b>nexenta_sparse = False</b>	boolean value	Enables or disables the creation of sparse datasets
<b>nexenta_sparsed_volumes = True</b>	boolean value	Enables or disables the creation of volumes as sparsed files that take no space. If disabled (False), volume is created as a regular file, which takes a long time.
<b>nexenta_target_group_prefix = cinder/</b>	string value	Prefix for iSCSI target groups on SA
<b>nexenta_target_prefix = iqn.1986-03.com.sun:02:cinder-</b>	string value	IQN prefix for iSCSI targets
<b>nexenta_use_https = True</b>	boolean value	Use secure HTTP for REST connection (default True)
<b>nexenta_user = admin</b>	string value	User name to connect to Nexenta SA
<b>nexenta_volume = cinder</b>	string value	SA Pool that holds all volumes
<b>nexenta_volume_group = iscsi</b>	string value	Volume group for ns5
<b>nfs_mount_attempts = 3</b>	integer value	The number of attempts to mount NFS shares before raising an error. At least one attempt will be made to mount an NFS share, regardless of the value specified.
<b>nfs_mount_options = None</b>	string value	Mount options passed to the NFS client. See section of the NFS man page for details.
<b>nfs_mount_point_base = \$state_path/mnt</b>	string value	Base dir containing mount points for NFS shares.
<b>nfs_qcow2_volumes = False</b>	boolean value	Create volumes as QCOW2 files rather than raw files.

Configuration option = Default value	Type	Description
<b>nfs_shares_config = /etc/cinder/nfs_shares</b>	string value	File with the list of available NFS shares.
<b>nfs_snapshot_support = False</b>	boolean value	Enable support for snapshots on the NFS driver. Platforms using libvirt <1.2.7 will encounter issues with this feature.
<b>nfs_sparsed_volumes = True</b>	boolean value	Create volumes as sparsed files which take no space. If set to False volume is created as regular file. In such case volume creation takes a lot of time.
<b>nimble_pool_name = default</b>	string value	Nimble Controller pool name
<b>nimble_subnet_label = *</b>	string value	Nimble Subnet Label
<b>nimble_verify_cert_path = None</b>	string value	Path to Nimble Array SSL certificate
<b>nimble_verify_certificate = False</b>	boolean value	Whether to verify Nimble SSL Certificate
<b>num_iser_scan_tries = 3</b>	integer value	The maximum number of times to rescan iSER target to find volume
<b>num_shell_tries = 3</b>	integer value	Number of times to attempt to run flakey shell commands
<b>num_volume_device_scan_tries = 3</b>	integer value	The maximum number of times to rescan targets to find volume
<b>pool_id_filter = []</b>	list value	Pool id permit to use.
<b>pool_type = default</b>	string value	Pool type, like sata-2copy.
<b>proxy = cinder.volume.drivers.ibm.ibm_storage.proxy.IBMStorageProxy</b>	string value	Proxy driver that connects to the IBM Storage Array
<b>pure_api_token = None</b>	string value	REST API authorization token.
<b>pure_automatic_max_oversubscription_ratio = True</b>	boolean value	Automatically determine an oversubscription ratio based on the current total data reduction values. If used this calculated value will override the max_oversubscription_ratio config option.

Configuration option = Default value	Type	Description
<b>pure_eradicate_on_delete = False</b>	boolean value	When enabled, all Pure volumes, snapshots, and protection groups will be eradicated at the time of deletion in Cinder. Data will NOT be recoverable after a delete with this set to True! When disabled, volumes and snapshots will go into pending eradication state and can be recovered.
<b>pure_replica_interval_default = 3600</b>	integer value	Snapshot replication interval in seconds.
<b>pure_replica_retention_logging_term_default = 7</b>	integer value	Retain snapshots per day on target for this time (in days.)
<b>pure_replica_retention_logging_term_per_day_default = 3</b>	integer value	Retain how many snapshots for each day.
<b>pure_replica_retention_short_term_default = 14400</b>	integer value	Retain all snapshots on target for this time (in seconds.)
<b>qnap_management_url = None</b>	uri value	The URL to management QNAP Storage
<b>qnap_poolname = None</b>	string value	The pool name in the QNAP Storage
<b>qnap_storage_protocol = iscsi</b>	string value	Communication protocol to access QNAP storage
<b>quobyte_client_cfg = None</b>	string value	Path to a Quobyte Client configuration file.
<b>quobyte_mount_point_base = \$state_path/mnt</b>	string value	Base dir containing the mount point for the Quobyte volume.
<b>quobyte_qcow2_volumes = True</b>	boolean value	Create volumes as QCOW2 files rather than raw files.
<b>quobyte_sparsed_volumes = True</b>	boolean value	Create volumes as sparse files which take no space. If set to False, volume is created as regular file.
<b>quobyte_volume_url = None</b>	string value	Quobyte URL to the Quobyte volume using e.g. a DNS SRV record (preferred) or a host list (alternatively) like quobyte://<DIR host1>, <DIR host2>/<volume name>

Configuration option = Default value	Type	Description
<b>rados_connect_timeout = -1</b>	integer value	Timeout value (in seconds) used when connecting to ceph cluster. If value < 0, no timeout is set and default librados value is used.
<b>rados_connection_interval = 5</b>	integer value	Interval value (in seconds) between connection retries to ceph cluster.
<b>rados_connection_retries = 3</b>	integer value	Number of retries if connection to ceph cluster failed.
<code>`rbd_ceph_conf = `</code>	string value	Path to the ceph configuration file
<b>rbd_cluster_name = ceph</b>	string value	The name of ceph cluster
<b>rbd_exclusive_cinder_pool = True</b>	boolean value	Set to True if the pool is used exclusively by Cinder. On exclusive use driver won't query images' provisioned size as they will match the value calculated by the Cinder core code for <code>allocated_capacity_gb</code> . This reduces the load on the Ceph cluster as well as on the volume service.
<b>rbd_flatten_volume_from_snapshot = False</b>	boolean value	Flatten volumes created from snapshots to remove dependency from volume to snapshot
<code>`rbd_keyring_conf = `</code>	string value	Path to the ceph keyring file
<b>rbd_max_clone_depth = 5</b>	integer value	Maximum number of nested volume clones that are taken before a flatten occurs. Set to 0 to disable cloning.
<b>rbd_pool = rbd</b>	string value	The RADOS pool where rbd volumes are stored
<b>rbd_secret_uuid = None</b>	string value	The libvirt uuid of the secret for the rbd_user volumes
<b>rbd_store_chunk_size = 4</b>	integer value	Volumes will be chunked into objects of this size (in megabytes).
<b>rbd_user = None</b>	string value	The RADOS client name for accessing rbd volumes - only set when using cephx authentication
<b>remove_empty_host = False</b>	boolean value	To remove the host from Unity when the last LUN is detached from it. By default, it is False.

Configuration option = Default value	Type	Description
<b>replication_connect_timeout = 5</b>	integer value	Timeout value (in seconds) used when connecting to ceph cluster to do a demotion/promotion of volumes. If value < 0, no timeout is set and default librados value is used.
<b>replication_device = None</b>	dict value	Multi opt of dictionaries to represent a replication target device. This option may be specified multiple times in a single config section to specify multiple replication target devices. Each entry takes the standard dict config form: replication_device = target_device_id:<required>,key1:value1,key2:value2...
<b>report_discard_supported = False</b>	boolean value	Report to clients of Cinder that the backend supports discard (aka. trim/unmap). This will not actually change the behavior of the backend or the client directly, it will only notify that it can be used.
<b>report_dynamic_total_capacity = True</b>	boolean value	Set to True for driver to report total capacity as a dynamic value -used + current free- and to False to report a static value -quota max bytes if defined and global size of cluster if not-.
<b>reserved_percentage = 0</b>	integer value	The percentage of backend capacity is reserved
<b>retries = 200</b>	integer value	Use this value to specify number of retries.
<b>san_api_port = None</b>	port value	Port to use to access the SAN API
<code>`san_clustername = `</code>	string value	Cluster name to use for creating volumes
<code>`san_ip = `</code>	string value	IP address of SAN controller
<b>san_is_local = False</b>	boolean value	Execute commands locally instead of over SSH; use if the volume service is running on the SAN device
<b>san_login = admin</b>	string value	Username for SAN controller
<code>`san_password = `</code>	string value	Password for SAN controller
<code>`san_private_key = `</code>	string value	Filename of private key to use for SSH authentication
<b>san_rest_port = 8443</b>	port value	REST server port number.
<b>san_ssh_port = 22</b>	port value	SSH port to use with SAN

Configuration option = Default value	Type	Description
<b>san_thin_provision = True</b>	boolean value	Use thin provisioning for SAN volumes?
<b>scaleio_server_certificate_path = None</b>	string value	Server certificate path
<b>scaleio_verify_server_certificate = False</b>	boolean value	verify server certificate
<b>scst_target_driver = iscsi</b>	string value	SCST target implementation can choose from multiple SCST target drivers.
<b>scst_target_iqn_name = None</b>	string value	Certain ISCSI targets have predefined target names, SCST target driver uses this name.
<code>`secondary_san_ip = `</code>	string value	IP address of secondary DSM controller
<b>secondary_san_login = Admin</b>	string value	Secondary DSM user name
<code>`secondary_san_password = `</code>	string value	Secondary DSM user password name
<b>secondary_sc_api_port = 3033</b>	port value	Secondary Dell API port
<b>sf_account_prefix = None</b>	string value	Create SolidFire accounts with this prefix. Any string can be used here, but the string "hostname" is special and will create a prefix using the cinder node hostname (previous default behavior). The default is NO prefix.
<b>sf_allow_template_caching = False</b>	boolean value	This option is deprecated and will be removed in the next OpenStack release. Please use the general cinder image-caching feature instead.
<b>sf_allow_tenant_qos = False</b>	boolean value	Allow tenants to specify QOS on create
<b>sf_api_port = 443</b>	port value	SolidFire API port. Useful if the device api is behind a proxy on a different port.
<b>sf_emulate_512 = True</b>	boolean value	Set 512 byte emulation on volume creation;
<b>sf_enable_vag = False</b>	boolean value	Utilize volume access groups on a per-tenant basis.

Configuration option = Default value	Type	Description
<b>sf_svip = None</b>	string value	Overrides default cluster SVIP with the one specified. This is required for deployments that have implemented the use of VLANs for iSCSI networks in their cloud.
<b>sf_template_account_name = openstack-vtemplate</b>	string value	Account name on the SolidFire Cluster to use as owner of template/cache volumes (created if does not exist).
<b>sf_volume_prefix = UUID-</b>	string value	Create SolidFire volumes with this prefix. Volume names are of the form <sf_volume_prefix><cinder-volume-id>. The default is to use a prefix of <i>UUID-</i> .
<b>sheepdog_store_address = 127.0.0.1</b>	string value	IP address of sheep daemon.
<b>sheepdog_store_port = 7000</b>	port value	Port of sheep daemon.
<b>sio_allow_non_padded_volumes = False</b>	boolean value	Allow volumes to be created in Storage Pools when zero padding is disabled. This option should not be enabled if multiple tenants will utilize volumes from a shared Storage Pool.
<b>sio_max_over_subscription_ratio = 10.0</b>	floating point value	max_over_subscription_ratio setting for the ScaleIO driver. This replaces the general max_over_subscription_ratio which has no effect in this driver. Maximum value allowed for ScaleIO is 10.0.
<b>sio_protection_domain_id = None</b>	string value	DEPRECATED: Protection Domain ID.
<b>sio_protection_domain_name = None</b>	string value	DEPRECATED: Protection Domain name.
<b>sio_rest_server_port = 443</b>	string value	REST server port.
<b>sio_round_volume_capacity = True</b>	boolean value	Round up volume capacity.
<b>sio_server_api_version = None</b>	string value	ScaleIO API version.

Configuration option = Default value	Type	Description
<b>sio_server_certificate_path = None</b>	string value	Server certificate path.
<b>sio_storage_pool_id = None</b>	string value	DEPRECATED: Storage Pool ID.
<b>sio_storage_pool_name = None</b>	string value	DEPRECATED: Storage Pool name.
<b>sio_storage_pools = None</b>	string value	Storage Pools.
<b>sio_unmap_volume_before_deletion = False</b>	boolean value	Unmap volume before deletion.
<b>sio_verify_server_certificate = False</b>	boolean value	Verify server certificate.
<b>smbfs_default_volume_format = vhd</b>	string value	Default format that will be used when creating volumes if no volume format is specified.
<b>smbfs_mount_point_base = C:\OpenStack\mnt</b>	string value	Base dir containing mount points for smbfs shares.
<b>smbfs_pool_mappings = {}</b>	dict value	Mappings between share locations and pool names. If not specified, the share names will be used as pool names. Example: //addr/share:pool_name, //addr/share2:pool_name2
<b>smbfs_shares_config = C:\OpenStack\smbfs_shares.txt</b>	string value	File with the list of available smbfs shares.
<b>ssh_conn_timeout = 30</b>	integer value	SSH connection timeout in seconds
<b>ssh_max_pool_conn = 5</b>	integer value	Maximum ssh connections in the pool
<b>ssh_min_pool_conn = 1</b>	integer value	Minimum ssh connections in the pool
<b>storage_protocol = iscsi</b>	string value	Protocol for transferring data between host and storage back-end.
<b>storage_vnx_authentication_type = global</b>	string value	VNX authentication scope type. By default, the value is global.



Configuration option = Default value	Type	Description
<b>storage_vnx_pool_names</b> = None	list value	Comma-separated list of storage pool names to be used.
<b>storage_vnx_security_file_dir</b> = None	string value	Directory path that contains the VNX security file. Make sure the security file is generated first.
<b>storwize_peer_pool</b> = None	string value	Specifies the name of the peer pool for hyperswap volume, the peer pool must exist on the other site.
<b>storwize_preferred_host_site</b> = {}	dict value	Specifies the site information for host. One WWPN or multi WWPNs used in the host can be specified. For example: storwize_preferred_host_site=site1:wwpn1,site2:wwpn2&wwpn3 or storwize_preferred_host_site=site1:iqn1,site2:iqn2
<b>storwize_san_secondary_ip</b> = None	string value	Specifies secondary management IP or hostname to be used if san_ip is invalid or becomes inaccessible.
<b>storwize_svc_allow_tenant_qos</b> = False	boolean value	Allow tenants to specify QOS on create
<b>storwize_svc_flashcopy_rate</b> = 50	integer value	Specifies the Storwize FlashCopy copy rate to be used when creating a full volume copy. The default is rate is 50, and the valid rates are 1-100.
<b>storwize_svc_flashcopy_timeout</b> = 120	integer value	Maximum number of seconds to wait for FlashCopy to be prepared.
<b>storwize_svc_iscsi_chap_enabled</b> = True	boolean value	Configure CHAP authentication for iSCSI connections (Default: Enabled)
<b>storwize_svc_mirror_pool</b> = None	string value	Specifies the name of the pool in which mirrored copy is stored. Example: "pool2"
<b>storwize_svc_multihostmap_enabled</b> = True	boolean value	This option no longer has any affect. It is deprecated and will be removed in the next release.
<b>storwize_svc_multipath_enabled</b> = False	boolean value	Connect with multipath (FC only; iSCSI multipath is controlled by Nova)
<b>storwize_svc_stretched_cluster_partner</b> = None	string value	If operating in stretched cluster mode, specify the name of the pool in which mirrored copies are stored.Example: "pool2"

Configuration option = Default value	Type	Description
<b>storwize_svc_vol_autoexpand = True</b>	boolean value	Storage system autoexpand parameter for volumes (True/False)
<b>storwize_svc_vol_compression = False</b>	boolean value	Storage system compression option for volumes
<b>storwize_svc_vol_easytier = True</b>	boolean value	Enable Easy Tier for volumes
<b>storwize_svc_vol_grainsize = 256</b>	integer value	Storage system grain size parameter for volumes (32/64/128/256)
<b>storwize_svc_vol_iogrp = 0</b>	string value	The I/O group in which to allocate volumes. It can be a comma-separated list in which case the driver will select an io_group based on least number of volumes associated with the io_group.
<b>storwize_svc_vol_nofmtdisk = False</b>	boolean value	Specifies that the volume not be formatted during creation.
<b>storwize_svc_vol_rsize = 2</b>	integer value	Storage system space-efficiency parameter for volumes (percentage)
<b>storwize_svc_vol_warning = 0</b>	integer value	Storage system threshold for volume capacity warnings (percentage)
<b>storwize_svc_volpool_name = [u'volpool']</b>	list value	Comma separated list of storage system storage pools for volumes.
<b>suppress_requests_ssl_warnings = False</b>	boolean value	Suppress requests library SSL certificate warnings.
<b>synology_admin_port = 5000</b>	port value	Management port for Synology storage.
<b>synology_device_id = None</b>	string value	Device id for skip one time password check for logging in Synology storage if OTP is enabled.
<b>synology_one_time_pass = None</b>	string value	One time password of administrator for logging in Synology storage if OTP is enabled.
<code>synology_password = `</code>	string value	Password of administrator for logging in Synology storage.

Configuration option = Default value	Type	Description
<code>`synology_pool_name = `</code>	string value	Volume on Synology storage to be used for creating lun.
<b>synology_ssl_verify = True</b>	boolean value	Do certificate validation or not if <code>\$driver_use_ssl</code> is True
<b>synology_username = admin</b>	string value	Administrator of Synology storage.
<b>target_helper = tgtadm</b>	string value	iSCSI target user-land tool to use. <code>tgtadm</code> is default, use <code>lioadm</code> for LIO iSCSI support, <code>scstadmin</code> for SCST target support, <code>ietadm</code> for iSCSI Enterprise Target, <code>iscsictl</code> for Chelsio iSCSI Target or <code>fake</code> for testing.
<b>target_ip_address = \$my_ip</b>	string value	The IP address that the iSCSI daemon is listening on
<b>target_port = 3260</b>	port value	The port that the iSCSI daemon is listening on
<b>target_prefix = iqn.2010-10.org.openstack:</b>	string value	Prefix for iSCSI volumes
<b>target_protocol = iscsi</b>	string value	Determines the iSCSI protocol for new iSCSI volumes, created with <code>tgtadm</code> or <code>lioadm</code> target helpers. In order to enable RDMA, this parameter should be set with the value "iser". The supported iSCSI protocol values are "iscsi" and "iser".
<b>thres_avl_size_perc_start = 20</b>	integer value	If the percentage of available space for an NFS share has dropped below the value specified by this option, the NFS image cache will be cleaned.
<b>thres_avl_size_perc_stop = 60</b>	integer value	When the percentage of available space on an NFS share has reached the percentage specified by this option, the driver will stop clearing files from the NFS image cache that have not been accessed in the last M minutes, where M is the value of the <code>expiry_thres_minutes</code> configuration option.
<b>tintri_api_version = v310</b>	string value	API version for the storage system
<b>tintri_image_cache_expiry_days = 30</b>	integer value	Delete unused image snapshots older than mentioned days

Configuration option = Default value	Type	Description
<b>tintri_image_shares_config = None</b>	string value	Path to image nfs shares file
<b>tintri_server_hostname = None</b>	string value	The hostname (or IP address) for the storage system
<b>tintri_server_password = None</b>	string value	Password for the storage system
<b>tintri_server_username = None</b>	string value	User name for the storage system
<b>trace_flags = None</b>	list value	List of options that control which trace info is written to the DEBUG log level to assist developers. Valid values are method and api.
<b>unique_fqdn_network = True</b>	boolean value	Whether or not our private network has unique FQDN on each initiator or not. For example networks with QA systems usually have multiple servers/VMs with the same FQDN. When true this will create host entries on K2 using the FQDN, when false it will use the reversed IQN/WWNN.
<b>unity_io_ports = []</b>	list value	A comma-separated list of iSCSI or FC ports to be used. Each port can be Unix-style glob expressions.
<b>unity_storage_pool_names = []</b>	list value	A comma-separated list of storage pool names to be used.
<b>use_chap_auth = False</b>	boolean value	Option to enable/disable CHAP authentication for targets.
<b>use_multipath_for_image_xfer = False</b>	boolean value	Do we attach/detach volumes in cinder using multipath for volume to image and image to volume transfers?
<b>vmax_array = None</b>	string value	Serial number of the array to connect to.
<b>vmax_port_groups = None</b>	list value	List of port groups containing frontend ports configured prior for server connection.
<b>vmax_service_level = None</b>	string value	Service level to use for provisioning storage.
<b>vmax_srp = None</b>	string value	Storage resource pool on array to use for provisioning.

Configuration option = Default value	Type	Description
<b>vmware_workload = None</b>	string value	Workload
<b>vmware_adapter_type = lsiLogic</b>	string value	Default adapter type to be used for attaching volumes.
<b>vmware_api_retry_count = 10</b>	integer value	Number of times VMware vCenter server API must be retried upon connection related issues.
<b>vmware_ca_file = None</b>	string value	CA bundle file to use in verifying the vCenter server certificate.
<b>vmware_cluster_name = None</b>	multi valued	Name of a vCenter compute cluster where volumes should be created.
<b>vmware_connection_pool _size = 10</b>	integer value	Maximum number of connections in http connection pool.
<b>vmware_host_ip = None</b>	string value	IP address for connecting to VMware vCenter server.
<b>vmware_host_password = None</b>	string value	Password for authenticating with VMware vCenter server.
<b>vmware_host_port = 443</b>	port value	Port number for connecting to VMware vCenter server.
<b>vmware_host_username = None</b>	string value	Username for authenticating with VMware vCenter server.
<b>vmware_host_version = None</b>	string value	Optional string specifying the VMware vCenter server version. The driver attempts to retrieve the version from VMware vCenter server. Set this configuration only if you want to override the vCenter server version.
<b>vmware_image_transfer_t imeout_secs = 7200</b>	integer value	Timeout in seconds for VMDK volume transfer between Cinder and Glance.
<b>vmware_insecure = False</b>	boolean value	If true, the vCenter server certificate is not verified. If false, then the default CA truststore is used for verification. This option is ignored if "vmware_ca_file" is set.

Configuration option = Default value	Type	Description
<b>vmware_lazy_create = True</b>	boolean value	If true, the backend volume in vCenter server is created lazily when the volume is created without any source. The backend volume is created when the volume is attached, uploaded to image service or during backup.
<b>vmware_max_objects_retrieval = 100</b>	integer value	Max number of objects to be retrieved per batch. Query results will be obtained in batches from the server and not in one shot. Server may still limit the count to something less than the configured value.
<b>vmware_snapshot_format = template</b>	string value	Volume snapshot format in vCenter server.
<b>vmware_task_poll_interval = 2.0</b>	floating point value	The interval (in seconds) for polling remote tasks invoked on VMware vCenter server.
<b>vmware_tmp_dir = /tmp</b>	string value	Directory where virtual disks are stored during volume backup and restore.
<b>vmware_volume_folder = Volumes</b>	string value	Name of the vCenter inventory folder that will contain Cinder volumes. This folder will be created under "OpenStack/<project_folder>", where project_folder is of format "Project (<volume_project_id>)".
<b>vmware_wsdl_location = None</b>	string value	Optional VIM service WSDL Location e.g <a href="http://&lt;server&gt;/vimService.wsdl">http://&lt;server&gt;/vimService.wsdl</a> . Optional over-ride to default location for bug work-arounds.
<b>volume_backend_name = None</b>	string value	The backend name for a given driver implementation
<b>volume_clear = zero</b>	string value	Method used to wipe old volumes
<b>volume_clear_ionice = None</b>	string value	The flag to pass to ionice to alter the i/o priority of the process used to zero a volume after deletion, for example "-c3" for idle only priority.
<b>volume_clear_size = 0</b>	integer value	Size in MiB to wipe at start of old volumes. 1024 MiB at max. 0 ⇒ all
<b>volume_copy_blkio_cgroup_name = cinder-volume-copy</b>	string value	The blkio cgroup name to be used to limit bandwidth of volume copy

Configuration option = Default value	Type	Description
<b>volume_copy_bps_limit = 0</b>	integer value	The upper limit of bandwidth of volume copy. 0 ⇒ unlimited
<b>volume_dd_blocksize = 1M</b>	string value	The default block size used when copying/clearing volumes
<b>volume_driver = cinder.volume.drivers.lvm.LVMVolumeDriver</b>	string value	Driver to use for volume creation
<b>volume_group = cinder-volumes</b>	string value	Name for the VG that will contain exported volumes
<b>volumes_dir = \$state_path/volumes</b>	string value	Volume configuration file storage directory
<b>vsstorage_default_volume_format = raw</b>	string value	Default format that will be used when creating volumes if no volume format is specified.
<b>vsstorage_mount_options = None</b>	list value	Mount options passed to the vsstorage client. See section of the pstorage-mount man page for details.
<b>vsstorage_mount_point_base = \$state_path/mnt</b>	string value	Base dir containing mount points for vsstorage shares.
<b>vsstorage_shares_config = /etc/cinder/vsstorage_shares</b>	string value	File with the list of available vsstorage shares.
<b>vsstorage_sparsed_volumes = True</b>	boolean value	Create volumes as sparsed files which take no space rather than regular files when using raw format, in which case volume creation takes lot of time.
<b>vsstorage_used_ratio = 0.95</b>	floating point value	Percent of ACTUAL usage of the underlying volume before no new volumes can be allocated to the volume destination.
<b>windows_iscsi_lun_path = C:\iSCSIVirtualDisks</b>	string value	Path to store VHD backed volumes
<b>xtremio_array_busy_retry_count = 5</b>	integer value	Number of retries in case array is busy

Configuration option = Default value	Type	Description
<b>xtremio_array_busy_retry_interval = 5</b>	integer value	Interval between retries in case array is busy
<b>xtremio_clean_unused_ig = False</b>	boolean value	Should the driver remove initiator groups with no volumes after the last connection was terminated. Since the behavior till now was to leave the IG be, we default to False (not deleting IGs without connected volumes); setting this parameter to True will remove any IG after terminating its connection to the last volume.
<code>xtremio_cluster_name = `</code>	string value	XMS cluster id in multi-cluster environment
<b>xtremio_volumes_per_glance_cache = 100</b>	integer value	Number of volumes created from each cached glance image
<b>zadara_default_snap_policy = False</b>	boolean value	VPSA - Attach snapshot policy for volumes
<b>zadara_password = None</b>	string value	VPSA - Password
<b>zadara_ssl_cert_verify = True</b>	boolean value	If set to True the http client will validate the SSL certificate of the VPSA endpoint.
<b>zadara_use_iser = True</b>	boolean value	VPSA - Use ISER instead of iSCSI
<b>zadara_user = None</b>	string value	VPSA - Username
<b>zadara_vol_encrypt = False</b>	boolean value	VPSA - Default encryption policy for volumes
<b>zadara_vol_name_template = OS_%s</b>	string value	VPSA - Default template for VPSA volume names
<b>zadara_vpsa_host = None</b>	string value	VPSA - Management Host name or IP address
<b>zadara_vpsa_poolname = None</b>	string value	VPSA - Storage Pool assigned for volumes
<b>zadara_vpsa_port = None</b>	port value	VPSA - Port number
<b>zadara_vpsa_use_ssl = False</b>	boolean value	VPSA - Use SSL connection
<b>zfssa_cache_directory = os-cinder-cache</b>	string value	Name of directory inside zfssa_nfs_share where cache volumes are stored.



Configuration option = Default value	Type	Description
<b>zfssa_cache_project = os-cinder-cache</b>	string value	Name of ZFSSA project where cache volumes are stored.
<b>zfssa_data_ip = None</b>	string value	Data path IP address
<b>zfssa_enable_local_cache = True</b>	boolean value	Flag to enable local caching: True, False.
<b>zfssa_https_port = 443</b>	string value	HTTPS port number
<code>`zfssa_initiator = `</code>	string value	iSCSI initiator IQNs. (comma separated)
<code>`zfssa_initiator_config = `</code>	string value	iSCSI initiators configuration.
<code>`zfssa_initiator_group = `</code>	string value	iSCSI initiator group.
<code>`zfssa_initiator_password = `</code>	string value	Secret of the iSCSI initiator CHAP user.
<code>`zfssa_initiator_user = `</code>	string value	iSCSI initiator CHAP user (name).
<b>zfssa_lun_compression = off</b>	string value	Data compression.
<b>zfssa_lun_logbias = latency</b>	string value	Synchronous write bias.
<b>zfssa_lun_sparse = False</b>	boolean value	Flag to enable sparse (thin-provisioned): True, False.
<b>zfssa_lun_volblocksize = 8k</b>	string value	Block size.
<b>zfssa_manage_policy = loose</b>	string value	Driver policy for volume manage.
<code>`zfssa_nfs_mount_options = `</code>	string value	Options to be passed while mounting share over nfs
<code>`zfssa_nfs_pool = `</code>	string value	Storage pool name.
<b>zfssa_nfs_project = NFSProject</b>	string value	Project name.
<b>zfssa_nfs_share = nfs_share</b>	string value	Share name.

Configuration option = Default value	Type	Description
<b>zfssa_nfs_share_compression = off</b>	string value	Data compression.
<b>zfssa_nfs_share_logbias = latency</b>	string value	Synchronous write bias-latency, throughput.
<b>zfssa_pool = None</b>	string value	Storage pool name.
<b>zfssa_project = None</b>	string value	Project name.
<code>`zfssa_replication_ip = `</code>	string value	IP address used for replication data. (maybe the same as data ip)
<b>zfssa_rest_timeout = None</b>	integer value	REST connection timeout. (seconds)
<b>zfssa_target_group = tgt-grp</b>	string value	iSCSI target group name.
<b>zfssa_target_interfaces = None</b>	string value	Network interfaces of iSCSI targets. (comma separated)
<code>`zfssa_target_password = `</code>	string value	Secret of the iSCSI target CHAP user.
<b>zfssa_target_portal = None</b>	string value	iSCSI target portal (Data-IP:Port, w.x.y.z:3260).
<code>`zfssa_target_user = `</code>	string value	iSCSI target CHAP user (name).

### 3.1.4. barbican

The following table outlines the options available under the **[barbican]** group in the `/etc/cinder/cinder.conf` file.

Table 3.4. barbican

Configuration option = Default value	Type	Description
<b>auth_endpoint = http://localhost/identity/v3</b>	string value	Use this endpoint to connect to Keystone
<b>barbican_api_version = None</b>	string value	Version of the Barbican API, for example: "v1"

Configuration option = Default value	Type	Description
<b>barbican_endpoint = None</b>	string value	Use this endpoint to connect to Barbican, for example: "http://localhost:9311/"
<b>number_of_retries = 60</b>	integer value	Number of times to retry poll for key creation completion
<b>retry_delay = 1</b>	integer value	Number of seconds to wait before retrying poll for key creation completion
<b>verify_ssl = True</b>	boolean value	Specifies if insecure TLS (https) requests. If False, the server's certificate will not be validated

### 3.1.5. brcd\_fabric\_example

The following table outlines the options available under the **[brcd\_fabric\_example]** group in the `/etc/cinder/cinder.conf` file.

Table 3.5. brcd\_fabric\_example

Configuration option = Default value	Type	Description
<code>`fc_fabric_address = `</code>	string value	Management IP of fabric.
<code>`fc_fabric_password = `</code>	string value	Password for user.
<b>fc_fabric_port = 22</b>	port value	Connecting port
<code>`fc_fabric_ssh_cert_path = `</code>	string value	Local SSH certificate Path.
<code>`fc_fabric_user = `</code>	string value	Fabric user ID.
<b>fc_southbound_protocol = HTTP</b>	string value	South bound connector for the fabric.
<b>fc_virtual_fabric_id = None</b>	string value	Virtual Fabric ID.
<b>zone_activate = True</b>	boolean value	Overridden zoning activation state.
<b>zone_name_prefix = openstack</b>	string value	Overridden zone name prefix.

Configuration option = Default value	Type	Description
<b>zoning_policy = initiator-target</b>	string value	Overridden zoning policy.

### 3.1.6. cisco\_fabric\_example

The following table outlines the options available under the **[cisco\_fabric\_example]** group in the **/etc/cinder/cinder.conf** file.

Table 3.6. cisco\_fabric\_example

Configuration option = Default value	Type	Description
<code>`cisco_fc_fabric_address = `</code>	string value	Management IP of fabric
<code>`cisco_fc_fabric_password = `</code>	string value	Password for user
<b>cisco_fc_fabric_port = 22</b>	port value	Connecting port
<code>`cisco_fc_fabric_user = `</code>	string value	Fabric user ID
<b>cisco_zone_activate = True</b>	boolean value	overridden zoning activation state
<b>cisco_zone_name_prefix = None</b>	string value	overridden zone name prefix
<b>cisco_zoning_policy = initiator-target</b>	string value	overridden zoning policy
<b>cisco_zoning_vsan = None</b>	string value	VSAN of the Fabric

### 3.1.7. coordination

The following table outlines the options available under the **[coordination]** group in the **/etc/cinder/cinder.conf** file.

Table 3.7. coordination

Configuration option = Default value	Type	Description
<b>backend_url = file://\$state_path</b>	string value	The backend URL to use for distributed coordination.

### 3.1.8. cors

The following table outlines the options available under the **[cors]** group in the `/etc/cinder/cinder.conf` file.

Table 3.8. cors

Configuration option = Default value	Type	Description
<b>allow_credentials = True</b>	boolean value	Indicate that the actual request can include user credentials
<b>allow_headers = [u'X-Auth-Token', u'X-Identity-Status', u'X-Roles', u'X-Service-Catalog', u'X-User-Id', u'X-Tenant-Id', u'X-OpenStack-Request-ID', u'X-Trace-Info', u'X-Trace-HMAC', u'OpenStack-API-Version']</b>	list value	Indicate which header field names may be used during the actual request.
<b>allow_methods = [u'GET', u'PUT', u'POST', u'DELETE', u'PATCH', u'HEAD']</b>	list value	Indicate which methods can be used during the actual request.
<b>allowed_origin = None</b>	list value	Indicate whether this resource may be shared with the domain received in the requests "origin" header. Format: "<protocol>://<host>[:<port>]", no trailing slash. Example: <a href="https://horizon.example.com">https://horizon.example.com</a>
<b>expose_headers = [u'X-Auth-Token', u'X-Subject-Token', u'X-Service-Token', u'X-OpenStack-Request-ID', u'OpenStack-API-Version']</b>	list value	Indicate which headers are safe to expose to the API. Defaults to HTTP Simple Headers.
<b>max_age = 3600</b>	integer value	Maximum cache age of CORS preflight requests.

### 3.1.9. database

The following table outlines the options available under the **[database]** group in the `/etc/cinder/cinder.conf` file.

Table 3.9. database

Configuration option = Default value	Type	Description
<b>backend = sqlalchemy</b>	string value	The back end to use for the database.
<b>connection = None</b>	string value	The SQLAlchemy connection string to use to connect to the database.
<b>connection_debug = 0</b>	integer value	Verbosity of SQL debugging information: 0=None, 100=Everything.
<b>connection_recycle_time = 3600</b>	integer value	Connections which have been present in the connection pool longer than this number of seconds will be replaced with a new one the next time they are checked out from the pool.
<b>connection_trace = False</b>	boolean value	Add Python stack traces to SQL as comment strings.
<b>db_inc_retry_interval = True</b>	boolean value	If True, increases the interval between retries of a database operation up to <code>db_max_retry_interval</code> .
<b>db_max_retries = 20</b>	integer value	Maximum retries in case of connection error or deadlock error before error is raised. Set to -1 to specify an infinite retry count.
<b>db_max_retry_interval = 10</b>	integer value	If <code>db_inc_retry_interval</code> is set, the maximum seconds between retries of a database operation.
<b>db_retry_interval = 1</b>	integer value	Seconds between retries of a database transaction.
<b>max_overflow = 50</b>	integer value	If set, use this value for <code>max_overflow</code> with SQLAlchemy.
<b>max_pool_size = 5</b>	integer value	Maximum number of SQL connections to keep open in a pool. Setting a value of 0 indicates no limit.
<b>max_retries = 10</b>	integer value	Maximum number of database connection retries during startup. Set to -1 to specify an infinite retry count.
<b>min_pool_size = 1</b>	integer value	Minimum number of SQL connections to keep open in a pool.

Configuration option = Default value	Type	Description
<b>mysql_enable_ndb = False</b>	boolean value	If True, transparently enables support for handling MySQL Cluster (NDB).
<b>mysql_sql_mode = TRADITIONAL</b>	string value	The SQL mode to be used for MySQL sessions. This option, including the default, overrides any server-set SQL mode. To use whatever SQL mode is set by the server configuration, set this to no value. Example: <code>mysql_sql_mode=</code>
<b>pool_timeout = None</b>	integer value	If set, use this value for <code>pool_timeout</code> with SQLAlchemy.
<b>retry_interval = 10</b>	integer value	Interval between retries of opening a SQL connection.
<b>slave_connection = None</b>	string value	The SQLAlchemy connection string to use to connect to the slave database.
<b>sqlite_synchronous = True</b>	boolean value	If True, SQLite uses synchronous mode.
<b>use_db_reconnect = False</b>	boolean value	Enable the experimental use of database reconnect on connection lost.

### 3.1.10. fc-zone-manager

The following table outlines the options available under the **[fc-zone-manager]** group in the `/etc/cinder/cinder.conf` file.

Table 3.10. fc-zone-manager

Configuration option = Default value	Type	Description
<b>brcd_sb_connector = HTTP</b>	string value	South bound connector for zoning operation
<b>cisco_sb_connector = cinder.zonemanager.drivers.cisco.cisco_fc_zone_client_cli.CiscoFCZoneClientCLI</b>	string value	Southbound connector for zoning operation

Configuration option = Default value	Type	Description
<b>enable_unsupported_driver = False</b>	boolean value	Set this to True when you want to allow an unsupported zone manager driver to start. Drivers that haven't maintained a working CI system and testing are marked as unsupported until CI is working again. This also marks a driver as deprecated and may be removed in the next release.
<b>fc_fabric_names = None</b>	string value	Comma separated list of Fibre Channel fabric names. This list of names is used to retrieve other SAN credentials for connecting to each SAN fabric
<b>fc_san_lookup_service = cinder.zonemanager.drivers.brocade.brcd_fc_san_lookup_service.BrcdFCSanLookupService</b>	string value	FC SAN Lookup Service
<b>zone_driver = cinder.zonemanager.drivers.brocade.brcd_fc_zone_driver.BrcdFCZoneDriver</b>	string value	FC Zone Driver responsible for zone management
<b>zoning_policy = initiator-target</b>	string value	Zoning policy configured by user; valid values include "initiator-target" or "initiator"

### 3.1.11. healthcheck

The following table outlines the options available under the **[healthcheck]** group in the `/etc/cinder/cinder.conf` file.

Table 3.11. healthcheck

Configuration option = Default value	Type	Description
<b>backends = []</b>	list value	Additional backends that can perform health checks and report that information back as part of a request.
<b>detailed = False</b>	boolean value	Show more detailed information as part of the response
<b>disable_by_file_path = None</b>	string value	Check the presence of a file to determine if an application is running on a port. Used by DisableByFileHealthcheck plugin.



Configuration option = Default value	Type	Description
<b>disable_by_file_paths = []</b>	list value	Check the presence of a file based on a port to determine if an application is running on a port. Expects a "port:path" list of strings. Used by DisableByFilesPortsHealthcheck plugin.
<b>path = /healthcheck</b>	string value	The path to respond to healthcheck requests on.

### 3.1.12. key\_manager

The following table outlines the options available under the **[key\_manager]** group in the **/etc/cinder/cinder.conf** file.

Table 3.12. key\_manager

Configuration option = Default value	Type	Description
<b>auth_type = None</b>	string value	The type of authentication credential to create. Possible values are <i>token</i> , <i>password</i> , <i>keystone_token</i> , and <i>keystone_password</i> . Required if no context is passed to the credential factory.
<b>auth_url = None</b>	string value	Use this endpoint to connect to Keystone.
<b>backend = barbican</b>	string value	Specify the key manager implementation. Options are "barbican" and "vault". Default is "barbican". Will support the values earlier set using <code>[key_manager]/api_class</code> for some time.
<b>domain_id = None</b>	string value	Domain ID for domain scoping. Optional for <i>keystone_token</i> and <i>keystone_password</i> auth_type.
<b>domain_name = None</b>	string value	Domain name for domain scoping. Optional for <i>keystone_token</i> and <i>keystone_password</i> auth_type.
<b>fixed_key = None</b>	string value	Fixed key returned by key manager, specified in hex
<b>password = None</b>	string value	Password for authentication. Required for <i>password</i> and <i>keystone_password</i> auth_type.
<b>project_domain_id = None</b>	string value	Project's domain ID for project. Optional for <i>keystone_token</i> and <i>keystone_password</i> auth_type.
<b>project_domain_name = None</b>	string value	Project's domain name for project. Optional for <i>keystone_token</i> and <i>keystone_password</i> auth_type.

Configuration option = Default value	Type	Description
<b>project_id = None</b>	string value	Project ID for project scoping. Optional for <i>keystone_token</i> and <i>keystone_password</i> auth_type.
<b>project_name = None</b>	string value	Project name for project scoping. Optional for <i>keystone_token</i> and <i>keystone_password</i> auth_type.
<b>reauthenticate = True</b>	boolean value	Allow fetching a new token if the current one is going to expire. Optional for <i>keystone_token</i> and <i>keystone_password</i> auth_type.
<b>token = None</b>	string value	Token for authentication. Required for <i>token</i> and <i>keystone_token</i> auth_type if no context is passed to the credential factory.
<b>trust_id = None</b>	string value	Trust ID for trust scoping. Optional for <i>keystone_token</i> and <i>keystone_password</i> auth_type.
<b>user_domain_id = None</b>	string value	User's domain ID for authentication. Optional for <i>keystone_token</i> and <i>keystone_password</i> auth_type.
<b>user_domain_name = None</b>	string value	User's domain name for authentication. Optional for <i>keystone_token</i> and <i>keystone_password</i> auth_type.
<b>user_id = None</b>	string value	User ID for authentication. Optional for <i>keystone_token</i> and <i>keystone_password</i> auth_type.
<b>username = None</b>	string value	Username for authentication. Required for <i>password</i> auth_type. Optional for the <i>keystone_password</i> auth_type.

### 3.1.13. keystone\_authtoken

The following table outlines the options available under the **[keystone\_authtoken]** group in the `/etc/cinder/cinder.conf` file.

Table 3.13. keystone\_authtoken

Configuration option = Default value	Type	Description
<b>auth_section = None</b>	string value	Config Section from which to load plugin specific options
<b>auth_type = None</b>	string value	Authentication type to load

Configuration option = Default value	Type	Description
<b>auth_uri = None</b>	string value	Complete "public" Identity API endpoint. This endpoint should not be an "admin" endpoint, as it should be accessible by all end users. Unauthenticated clients are redirected to this endpoint to authenticate. Although this endpoint should ideally be unversioned, client support in the wild varies. If you're using a versioned v2 endpoint here, then this should <b>not</b> be the same endpoint the service user utilizes for validating tokens, because normal end users may not be able to reach that endpoint. This option is deprecated in favor of <code>www_authenticate_uri</code> and will be removed in the S release.
<b>auth_version = None</b>	string value	API version of the admin Identity API endpoint.
<b>cache = None</b>	string value	Request environment key where the Swift cache object is stored. When <code>auth_token</code> middleware is deployed with a Swift cache, use this option to have the middleware share a caching backend with swift. Otherwise, use the <b>memcached_servers</b> option instead.
<b>cafile = None</b>	string value	A PEM encoded Certificate Authority to use when verifying HTTPs connections. Defaults to system CAs.
<b>certfile = None</b>	string value	Required if identity server requires client certificate
<b>check_revocations_for_cached = False</b>	boolean value	If true, the revocation list will be checked for cached tokens. This requires that PKI tokens are configured on the identity server.
<b>delay_auth_decision = False</b>	boolean value	Do not handle authorization requests within the middleware, but delegate the authorization decision to downstream WSGI components.
<b>enforce_token_bind = permissive</b>	string value	Used to control the use and type of token binding. Can be set to: "disabled" to not check token binding. "permissive" (default) to validate binding information if the bind type is of a form known to the server and ignore it if not. "strict" like "permissive" but if the bind type is unknown the token will be rejected. "required" any form of token binding is needed to be allowed. Finally the name of a binding method that must be present in tokens.

Configuration option = Default value	Type	Description
<b>hash_algorithms = [u'md5']</b>	list value	Hash algorithms to use for hashing PKI tokens. This may be a single algorithm or multiple. The algorithms are those supported by Python standard <code>hashlib.new()</code> . The hashes will be tried in the order given, so put the preferred one first for performance. The result of the first hash will be stored in the cache. This will typically be set to multiple values only while migrating from a less secure algorithm to a more secure one. Once all the old tokens are expired this option should be set to a single value for better performance.
<b>http_connect_timeout = None</b>	integer value	Request timeout value for communicating with Identity API server.
<b>http_request_max_retries = 3</b>	integer value	How many times are we trying to reconnect when communicating with Identity API Server.
<b>include_service_catalog = True</b>	boolean value	(Optional) Indicate whether to set the X-Service-Catalog header. If False, middleware will not ask for service catalog on token validation and will not set the X-Service-Catalog header.
<b>insecure = False</b>	boolean value	Verify HTTPS connections.
<b>keyfile = None</b>	string value	Required if identity server requires client certificate
<b>memcache_pool_conn_get_timeout = 10</b>	integer value	(Optional) Number of seconds that an operation will wait to get a memcached client connection from the pool.
<b>memcache_pool_dead_retry = 300</b>	integer value	(Optional) Number of seconds memcached server is considered dead before it is tried again.
<b>memcache_pool_maxsize = 10</b>	integer value	(Optional) Maximum total number of open connections to every memcached server.
<b>memcache_pool_socket_timeout = 3</b>	integer value	(Optional) Socket timeout in seconds for communicating with a memcached server.
<b>memcache_pool_unused_timeout = 60</b>	integer value	(Optional) Number of seconds a connection to memcached is held unused in the pool before it is closed.
<b>memcache_secret_key = None</b>	string value	(Optional, mandatory if <code>memcache_security_strategy</code> is defined) This string is used for key derivation.

Configuration option = Default value	Type	Description
<b>memcache_security_strategy = None</b>	string value	(Optional) If defined, indicate whether token data should be authenticated or authenticated and encrypted. If MAC, token data is authenticated (with HMAC) in the cache. If ENCRYPT, token data is encrypted and authenticated in the cache. If the value is not one of these options or empty, auth_token will raise an exception on initialization.
<b>memcache_use_advanced_pool = False</b>	boolean value	(Optional) Use the advanced (eventlet safe) memcached client pool. The advanced pool will only work under python 2.x.
<b>memcached_servers = None</b>	list value	Optionally specify a list of memcached server(s) to use for caching. If left undefined, tokens will instead be cached in-process.
<b>region_name = None</b>	string value	The region in which the identity server can be found.
<b>revocation_cache_time = 10</b>	integer value	Determines the frequency at which the list of revoked tokens is retrieved from the Identity service (in seconds). A high number of revocation events combined with a low cache duration may significantly reduce performance. Only valid for PKI tokens. This option has been deprecated in the Ocata release and will be removed in the P release.
<b>service_token_roles = [u'service']</b>	list value	A choice of roles that must be present in a service token. Service tokens are allowed to request that an expired token can be used and so this check should tightly control that only actual services should be sending this token. Roles here are applied as an ANY check so any role in this list must be present. For backwards compatibility reasons this currently only affects the allow_expired check.
<b>service_token_roles_required = False</b>	boolean value	For backwards compatibility reasons we must let valid service tokens pass that don't pass the service_token_roles check as valid. Setting this true will become the default in a future release and should be enabled if possible.
<b>signing_dir = None</b>	string value	Directory used to cache files related to PKI tokens. This option has been deprecated in the Ocata release and will be removed in the P release.

Configuration option = Default value	Type	Description
<b>token_cache_time = 300</b>	integer value	In order to prevent excessive effort spent validating tokens, the middleware caches previously-seen tokens for a configurable duration (in seconds). Set to -1 to disable caching completely.
<b>www_authenticate_uri = None</b>	string value	Complete "public" Identity API endpoint. This endpoint should not be an "admin" endpoint, as it should be accessible by all end users. Unauthenticated clients are redirected to this endpoint to authenticate. Although this endpoint should ideally be unversioned, client support in the wild varies. If you're using a versioned v2 endpoint here, then this should <b>not</b> be the same endpoint the service user utilizes for validating tokens, because normal end users may not be able to reach that endpoint.

### 3.1.14. matchmaker\_redis

The following table outlines the options available under the **[matchmaker\_redis]** group in the **/etc/cinder/cinder.conf** file.

Table 3.14. matchmaker\_redis

Configuration option = Default value	Type	Description
<b>check_timeout = 20000</b>	integer value	Time in ms to wait before the transaction is killed.
<b>host = 127.0.0.1</b>	string value	Host to locate redis.
<code>`password = `</code>	string value	Password for Redis server (optional).
<b>port = 6379</b>	port value	Use this port to connect to redis host.
<b>sentinel_group_name = oslo-messaging-zeromq</b>	string value	Redis replica set name.
<b>sentinel_hosts = []</b>	list value	List of Redis Sentinel hosts (fault tolerance mode), e.g., [host:port, host1:port ... ]
<b>socket_timeout = 10000</b>	integer value	Timeout in ms on blocking socket operations.
<b>wait_timeout = 2000</b>	integer value	Time in ms to wait between connection attempts.

### 3.1.15. nova

The following table outlines the options available under the **[nova]** group in the `/etc/cinder/cinder.conf` file.

Table 3.15. nova

Configuration option = Default value	Type	Description
<b>auth_section = None</b>	string value	Config Section from which to load plugin specific options
<b>auth_type = None</b>	string value	Authentication type to load
<b>cafile = None</b>	string value	PEM encoded Certificate Authority to use when verifying HTTPs connections.
<b>certfile = None</b>	string value	PEM encoded client certificate cert file
<b>insecure = False</b>	boolean value	Verify HTTPS connections.
<b>interface = public</b>	string value	Type of the nova endpoint to use. This endpoint will be looked up in the keystone catalog and should be one of public, internal or admin.
<b>keyfile = None</b>	string value	PEM encoded client certificate key file
<b>region_name = None</b>	string value	Name of nova region to use. Useful if keystone manages more than one region.
<b>timeout = None</b>	integer value	Timeout value for http requests
<b>token_auth_url = None</b>	string value	The authentication URL for the nova connection when using the current users token

### 3.1.16. oslo\_concurrency

The following table outlines the options available under the **[oslo\_concurrency]** group in the `/etc/cinder/cinder.conf` file.

Table 3.16. oslo\_concurrency

Configuration option = Default value	Type	Description
<b>disable_process_locking = False</b>	boolean value	Enables or disables inter-process locks.

Configuration option = Default value	Type	Description
<b>lock_path = None</b>	string value	Directory to use for lock files. For security, the specified directory should only be writable by the user running the processes that need locking. Defaults to environment variable OSLO_LOCK_PATH. If external locks are used, a lock path must be set.

### 3.1.17. oslo\_messaging\_amqp

The following table outlines the options available under the **[oslo\_messaging\_amqp]** group in the **/etc/cinder/cinder.conf** file.

Table 3.17. oslo\_messaging\_amqp

Configuration option = Default value	Type	Description
<b>addressing_mode = dynamic</b>	string value	Indicates the addressing mode used by the driver. Permitted values: <i>legacy</i> - use legacy non-routable addressing <i>routable</i> - use routable addresses <i>dynamic</i> - use legacy addresses if the message bus does not support routing otherwise use routable addressing
<b>allow_insecure_clients = False</b>	boolean value	Accept clients using either SSL or plain TCP
<b>anycast_address = anycast</b>	string value	Appended to the address prefix when sending to a group of consumers. Used by the message bus to identify messages that should be delivered in a round-robin fashion across consumers.
<b>broadcast_prefix = broadcast</b>	string value	address prefix used when broadcasting to all servers
<b>connection_retry_backoff = 2</b>	integer value	Increase the connection_retry_interval by this many seconds after each unsuccessful failover attempt.
<b>connection_retry_interval = 1</b>	integer value	Seconds to pause before attempting to re-connect.
<b>connection_retry_interval_max = 30</b>	integer value	Maximum limit for connection_retry_interval + connection_retry_backoff
<b>container_name = None</b>	string value	Name for the AMQP container. must be globally unique. Defaults to a generated UUID



Configuration option = Default value	Type	Description
<b>default_notification_exchange = None</b>	string value	Exchange name used in notification addresses. Exchange name resolution precedence: Target.exchange if set else default_notification_exchange if set else control_exchange if set else <i>notify</i>
<b>default_notify_timeout = 30</b>	integer value	The deadline for a sent notification message delivery. Only used when caller does not provide a timeout expiry.
<b>default_reply_retry = 0</b>	integer value	The maximum number of attempts to re-send a reply message which failed due to a recoverable error.
<b>default_reply_timeout = 30</b>	integer value	The deadline for an rpc reply message delivery.
<b>default_rpc_exchange = None</b>	string value	Exchange name used in RPC addresses. Exchange name resolution precedence: Target.exchange if set else default_rpc_exchange if set else control_exchange if set else <i>rpc</i>
<b>default_send_timeout = 30</b>	integer value	The deadline for an rpc cast or call message delivery. Only used when caller does not provide a timeout expiry.
<b>default_sender_link_timeout = 600</b>	integer value	The duration to schedule a purge of idle sender links. Detach link after expiry.
<b>group_request_prefix = unicast</b>	string value	address prefix when sending to any server in group
<b>idle_timeout = 0</b>	integer value	Timeout for inactive connections (in seconds)
<b>link_retry_delay = 10</b>	integer value	Time to pause between re-connecting an AMQP 1.0 link that failed due to a recoverable error.
<b>multicast_address = multicast</b>	string value	Appended to the address prefix when sending a fanout message. Used by the message bus to identify fanout messages.
<b>notify_address_prefix = openstack.org/om/notify</b>	string value	Address prefix for all generated Notification addresses
<b>notify_server_credit = 100</b>	integer value	Window size for incoming Notification messages
<b>`password = `</b>	string value	Password for message broker authentication

Configuration option = Default value	Type	Description
<b>pre_settled = [u'rpc-cast', u'rpc-reply']</b>	multi valued	Send messages of this type pre-settled. Pre-settled messages will not receive acknowledgement from the peer. Note well: pre-settled messages may be silently discarded if the delivery fails. Permitted values: <i>rpc-call</i> - send RPC Calls pre-settled <i>rpc-reply</i> - send RPC Replies pre-settled <i>rpc-cast</i> - Send RPC Casts pre-settled <i>notify</i> - Send Notifications pre-settled
<b>pseudo_vhost = True</b>	boolean value	Enable virtual host support for those message buses that do not natively support virtual hosting (such as <i>qpidd</i> ). When set to true the virtual host name will be added to all message bus addresses, effectively creating a private <i>subnet</i> per virtual host. Set to False if the message bus supports virtual hosting using the <i>hostname</i> field in the AMQP 1.0 Open performative as the name of the virtual host.
<b>reply_link_credit = 200</b>	integer value	Window size for incoming RPC Reply messages.
<b>rpc_address_prefix = openstack.org/om/rpc</b>	string value	Address prefix for all generated RPC addresses
<b>rpc_server_credit = 100</b>	integer value	Window size for incoming RPC Request messages
<code>`sasl_config_dir = `</code>	string value	Path to directory that contains the SASL configuration
<code>`sasl_config_name = `</code>	string value	Name of configuration file (without <i>.conf</i> suffix)
<code>`sasl_default_realm = `</code>	string value	SASL realm to use if no realm present in username
<code>`sasl_mechanisms = `</code>	string value	Space separated list of acceptable SASL mechanisms
<b>server_request_prefix = exclusive</b>	string value	address prefix used when sending to a specific server
<b>ssl = False</b>	boolean value	Attempt to connect via SSL. If no other <i>ssl</i> -related parameters are given, it will use the system's CA-bundle to verify the server's certificate.
<code>`ssl_ca_file = `</code>	string value	CA certificate PEM file used to verify the server's certificate
<code>`ssl_cert_file = `</code>	string value	Self-identifying certificate PEM file for client authentication

Configuration option = Default value	Type	Description
<code>`ssl_key_file = `</code>	string value	Private key PEM file used to sign <code>ssl_cert_file</code> certificate (optional)
<b>ssl_key_password = None</b>	string value	Password for decrypting <code>ssl_key_file</code> (if encrypted)
<b>ssl_verify_vhost = False</b>	boolean value	By default SSL checks that the name in the server's certificate matches the hostname in the <code>transport_url</code> . In some configurations it may be preferable to use the virtual hostname instead, for example if the server uses the Server Name Indication TLS extension (rfc6066) to provide a certificate per virtual host. Set <code>ssl_verify_vhost</code> to True if the server's SSL certificate uses the virtual host name instead of the DNS name.
<b>trace = False</b>	boolean value	Debug: dump AMQP frames to stdout
<b>unicast_address = unicast</b>	string value	Appended to the address prefix when sending to a particular RPC/Notification server. Used by the message bus to identify messages sent to a single destination.
<code>`username = `</code>	string value	User name for message broker authentication

### 3.1.18. oslo\_messaging\_kafka

The following table outlines the options available under the `[oslo_messaging_kafka]` group in the `/etc/cinder/cinder.conf` file.

Table 3.18. oslo\_messaging\_kafka

Configuration option = Default value	Type	Description
<b>conn_pool_min_size = 2</b>	integer value	The pool size limit for connections expiration policy
<b>conn_pool_ttl = 1200</b>	integer value	The time-to-live in sec of idle connections in the pool
<b>consumer_group = oslo_messaging_consumer</b>	string value	Group id for Kafka consumer. Consumers in one group will coordinate message consumption
<b>kafka_consumer_timeout = 1.0</b>	floating point value	Default timeout(s) for Kafka consumers

Configuration option = Default value	Type	Description
<b>kafka_default_host = localhost</b>	string value	Default Kafka broker Host
<b>kafka_default_port = 9092</b>	port value	Default Kafka broker Port
<b>kafka_max_fetch_bytes = 1048576</b>	integer value	Max fetch bytes of Kafka consumer
<b>pool_size = 10</b>	integer value	Pool Size for Kafka Consumers
<b>producer_batch_size = 16384</b>	integer value	Size of batch for the producer async send
<b>producer_batch_timeout = 0.0</b>	floating point value	Upper bound on the delay for KafkaProducer batching in seconds

### 3.1.19. oslo\_messaging\_notifications

The following table outlines the options available under the **[oslo\_messaging\_notifications]** group in the `/etc/cinder/cinder.conf` file.

Table 3.19. oslo\_messaging\_notifications

Configuration option = Default value	Type	Description
<b>driver = []</b>	multi valued	The Drivers(s) to handle sending notifications. Possible values are messaging, messagingv2, routing, log, test, noop
<b>retry = -1</b>	integer value	The maximum number of attempts to re-send a notification message which failed to be delivered due to a recoverable error. 0 - No retry, -1 - indefinite
<b>topics = [u'notifications']</b>	list value	AMQP topic used for OpenStack notifications.
<b>transport_url = None</b>	string value	A URL representing the messaging driver to use for notifications. If not set, we fall back to the same configuration used for RPC.

### 3.1.20. oslo\_messaging\_rabbit

The following table outlines the options available under the **[oslo\_messaging\_rabbit]** group in the `/etc/cinder/cinder.conf` file.

Table 3.20. oslo\_messaging\_rabbit

Configuration option = Default value	Type	Description
<b>amqp_auto_delete = False</b>	boolean value	Auto-delete queues in AMQP.
<b>amqp_durable_queues = False</b>	boolean value	Use durable queues in AMQP.
<b>channel_max = None</b>	integer value	Maximum number of channels to allow
<b>connection_factory = single</b>	string value	Connection factory implementation
<b>default_notification_exchange = \${control_exchange}_notification</b>	string value	Exchange name for sending notifications
<b>default_notification_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending notification, -1 means infinite retry.
<b>default_rpc_exchange = \${control_exchange}_rpc</b>	string value	Exchange name for sending RPC messages
<b>default_rpc_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending RPC message, -1 means infinite retry. If actual retry attempts in not 0 the rpc request could be processed more than one time
<b>default_serializer_type = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>fake_rabbit = False</b>	boolean value	Deprecated, use rpc_backend=kombu+memory or rpc_backend=fake
<b>frame_max = None</b>	integer value	The maximum byte size for an AMQP frame
<b>heartbeat_interval = 3</b>	integer value	How often to send heartbeats for consumer's connections
<b>heartbeat_rate = 2</b>	integer value	How often times during the heartbeat_timeout_threshold we check the heartbeat.

Configuration option = Default value	Type	Description
<b>heartbeat_timeout_thresh old = 60</b>	integer value	Number of seconds after which the Rabbit broker is considered down if heartbeat's keep-alive fails (0 disable the heartbeat). EXPERIMENTAL
<b>host_connection_reconn ect_delay = 0.25</b>	floating point value	Set delay for reconnection to some host which has connection error
<b>kombu_compression = None</b>	string value	EXPERIMENTAL: Possible values are: gzip, bz2. If not set compression will not be used. This option may not be available in future versions.
<b>kombu_failover_strategy = round-robin</b>	string value	Determines how the next RabbitMQ node is chosen in case the one we are currently connected to becomes unavailable. Takes effect only if more than one RabbitMQ node is provided in config.
<b>kombu_missing_consum er_retry_timeout = 60</b>	integer value	How long to wait a missing client before abandoning to send it its replies. This value should not be longer than <code>rpc_response_timeout</code> .
<b>kombu_reconnect_delay = 1.0</b>	floating point value	How long to wait before reconnecting in response to an AMQP consumer cancel notification.
<b>notification_listener_pref etch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to notification listener.
<b>notification_persistence = False</b>	boolean value	Persist notification messages.
<b>notification_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending notification message
<b>pool_max_overflow = 0</b>	integer value	Maximum number of connections to create above <b>pool_max_size</b> .
<b>pool_max_size = 30</b>	integer value	Maximum number of connections to keep queued.
<b>pool_recycle = 600</b>	integer value	Lifetime of a connection (since creation) in seconds or None for no recycling. Expired connections are closed on acquire.
<b>pool_stale = 60</b>	integer value	Threshold at which inactive (since release) connections are considered stale in seconds or None for no staleness. Stale connections are closed on acquire.

Configuration option = Default value	Type	Description
<b>pool_timeout = 30</b>	integer value	Default number of seconds to wait for a connections to available
<b>rabbit_ha_queues = False</b>	boolean value	Try to use HA queues in RabbitMQ (x-ha-policy: all). If you change this option, you must wipe the RabbitMQ database. In RabbitMQ 3.0, queue mirroring is no longer controlled by the x-ha-policy argument when declaring a queue. If you just want to make sure that all queues (except those with auto-generated names) are mirrored across all nodes, run: <code>"rabbitmqctl set_policy HA ^(?!amq\).* {"ha-mode": "all"}"</code>
<b>rabbit_host = localhost</b>	string value	The RabbitMQ broker address where a single node is used.
<b>rabbit_hosts = [u'\$rabbit_host:\$rabbit_port']</b>	list value	RabbitMQ HA cluster host:port pairs.
<b>rabbit_interval_max = 30</b>	integer value	Maximum interval of RabbitMQ connection retries. Default is 30 seconds.
<b>rabbit_login_method = AMQPLAIN</b>	string value	The RabbitMQ login method.
<b>rabbit_max_retries = 0</b>	integer value	Maximum number of RabbitMQ connection retries. Default is 0 (infinite retry count).
<b>rabbit_password = guest</b>	string value	The RabbitMQ password.
<b>rabbit_port = 5672</b>	port value	The RabbitMQ broker port where a single node is used.
<b>rabbit_qos_prefetch_count = 0</b>	integer value	Specifies the number of messages to prefetch. Setting to zero allows unlimited messages.
<b>rabbit_retry_backoff = 2</b>	integer value	How long to backoff for between retries when connecting to RabbitMQ.
<b>rabbit_retry_interval = 1</b>	integer value	How frequently to retry connecting with RabbitMQ.
<b>rabbit_transient_queues_ttl = 1800</b>	integer value	Positive integer representing duration in seconds for queue TTL (x-expires). Queues which are unused for the duration of the TTL are automatically deleted. The parameter affects only reply and fanout queues.

Configuration option = Default value	Type	Description
<b>rabbit_userid = guest</b>	string value	The RabbitMQ userid.
<b>rabbit_virtual_host = /</b>	string value	The RabbitMQ virtual host.
<b>rpc_listener_prefetch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to rpc listener.
<b>rpc_queue_expiration = 60</b>	integer value	Time to live for rpc queues without consumers in seconds.
<b>rpc_reply_exchange = \${control_exchange}_rpc_reply</b>	string value	Exchange name for receiving RPC replies
<b>rpc_reply_listener_prefetch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to rpc reply listener.
<b>rpc_reply_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending reply. -1 means infinite retry during rpc_timeout
<b>rpc_reply_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending reply.
<b>rpc_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending RPC message
<b>socket_timeout = 0.25</b>	floating point value	Set socket timeout in seconds for connection's socket
<b>ssl = None</b>	boolean value	Enable SSL
<code>`ssl_ca_file = `</code>	string value	SSL certification authority file (valid only if SSL enabled).
<code>`ssl_cert_file = `</code>	string value	SSL cert file (valid only if SSL enabled).
<code>`ssl_key_file = `</code>	string value	SSL key file (valid only if SSL enabled).
<b>ssl_options = None</b>	dict value	Arguments passed to <code>ssl.wrap_socket</code>
<code>`ssl_version = `</code>	string value	SSL version to use (valid only if SSL enabled). Valid values are TLSv1 and SSLv23. SSLv2, SSLv3, TLSv1_1, and TLSv1_2 may be available on some distributions.



Configuration option = Default value	Type	Description
<b>tcp_user_timeout = 0.25</b>	floating point value	Set TCP_USER_TIMEOUT in seconds for connection's socket

### 3.1.21. oslo\_messaging\_zmq

The following table outlines the options available under the **[oslo\_messaging\_zmq]** group in the **/etc/cinder/cinder.conf** file.

Table 3.21. oslo\_messaging\_zmq

Configuration option = Default value	Type	Description
<b>rpc_ack_timeout_base = 15</b>	integer value	Number of seconds to wait for an ack from a cast/call. After each retry attempt this timeout is multiplied by some specified multiplier.
<b>rpc_ack_timeout_multiplier = 2</b>	integer value	Number to multiply base ack timeout by after each retry attempt.
<b>rpc_message_ttl = 300</b>	integer value	Expiration timeout in seconds of a sent/received message after which it is not tracked anymore by a client/server.
<b>rpc_poll_timeout = 1</b>	integer value	The default number of seconds that poll should wait. Poll raises timeout exception when timeout expired.
<b>rpc_retry_attempts = 3</b>	integer value	Default number of message sending attempts in case of any problems occurred: positive value N means at most N retries, 0 means no retries, None or -1 (or any other negative values) mean to retry forever. This option is used only if acknowledgments are enabled.
<b>rpc_thread_pool_size = 100</b>	integer value	Maximum number of (green) threads to work concurrently.
<b>rpc_use_acks = False</b>	boolean value	Wait for message acknowledgements from receivers. This mechanism works only via proxy without PUB/SUB.
<b>rpc_zmq_bind_address = *</b>	string value	ZeroMQ bind address. Should be a wildcard (*), an ethernet interface, or IP. The "host" option should point or resolve to this address.
<b>rpc_zmq_bind_port_retries = 100</b>	integer value	Number of retries to find free port number before fail with ZMQBindError.

Configuration option = Default value	Type	Description
<b>rpc_zmq_contexts = 1</b>	integer value	Number of ZeroMQ contexts, defaults to 1.
<b>rpc_zmq_host = ccosp-undercloud.usersys.redhat.com</b>	string value	Name of this node. Must be a valid hostname, FQDN, or IP address. Must match "host" option, if running Nova.
<b>rpc_zmq_ipc_dir = /var/run/openstack</b>	string value	Directory for holding IPC sockets.
<b>rpc_zmq_matchmaker = redis</b>	string value	MatchMaker driver.
<b>rpc_zmq_max_port = 65536</b>	integer value	Maximal port number for random ports range.
<b>rpc_zmq_min_port = 49153</b>	port value	Minimal port number for random ports range.
<b>rpc_zmq_serialization = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>rpc_zmq_topic_backlog = None</b>	integer value	Maximum number of ingress messages to locally buffer per topic. Default is unlimited.
<b>subscribe_on = []</b>	list value	List of publisher hosts SubConsumer can subscribe on. This option has higher priority then the default publishers list taken from the matchmaker.
<b>use_dynamic_connections = False</b>	boolean value	This option makes direct connections dynamic or static. It makes sense only with use_router_proxy=False which means to use direct connections for direct message types (ignored otherwise).
<b>use_pub_sub = False</b>	boolean value	Use PUB/SUB pattern for fanout methods. PUB/SUB always uses proxy.
<b>use_router_proxy = False</b>	boolean value	Use ROUTER remote proxy.
<b>zmq_failover_connections = 2</b>	integer value	How many additional connections to a host will be made for failover reasons. This option is actual only in dynamic connections mode.

Configuration option = Default value	Type	Description
<b>zmq_immediate = True</b>	boolean value	This option configures round-robin mode in zmq socket. True means not keeping a queue when server side disconnects. False means to keep queue and messages even if server is disconnected, when the server appears we send all accumulated messages to it.
<b>zmq_linger = -1</b>	integer value	Number of seconds to wait before all pending messages will be sent after closing a socket. The default value of -1 specifies an infinite linger period. The value of 0 specifies no linger period. Pending messages shall be discarded immediately when the socket is closed. Positive values specify an upper bound for the linger period.
<b>zmq_target_expire = 300</b>	integer value	Expiration timeout in seconds of a name service record about existing target (< 0 means no timeout).
<b>zmq_target_update = 180</b>	integer value	Update period in seconds of a name service record about existing target.
<b>zmq_tcp_keepalive = -1</b>	integer value	Enable/disable TCP keepalive (KA) mechanism. The default value of -1 (or any other negative value) means to skip any overrides and leave it to OS default; 0 and 1 (or any other positive value) mean to disable and enable the option respectively.
<b>zmq_tcp_keepalive_cnt = -1</b>	integer value	The number of retransmissions to be carried out before declaring that remote end is not available. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_idle = -1</b>	integer value	The duration between two keepalive transmissions in idle condition. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_intvl = -1</b>	integer value	The duration between two successive keepalive retransmissions, if acknowledgement to the previous keepalive transmission is not received. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.

### 3.1.22. oslo\_middleware

The following table outlines the options available under the **[oslo\_middleware]** group in the `/etc/cinder/cinder.conf` file.

Table 3.22. oslo\_middleware

Configuration option = Default value	Type	Description
<b>enable_proxy_headers_parsing = False</b>	boolean value	Whether the application is behind a proxy or not. This determines if the middleware should parse the headers or not.
<b>max_request_body_size = 114688</b>	integer value	The maximum body size for each request, in bytes.
<b>secure_proxy_ssl_header = X-Forwarded-Proto</b>	string value	The HTTP Header that will be used to determine what the original request protocol scheme was, even if it was hidden by a SSL termination proxy.

### 3.1.23. oslo\_policy

The following table outlines the options available under the **[oslo\_policy]** group in the `/etc/cinder/cinder.conf` file.

Table 3.23. oslo\_policy

Configuration option = Default value	Type	Description
<b>enforce_scope = False</b>	boolean value	This option controls whether or not to enforce scope when evaluating policies. If <b>True</b> , the scope of the token used in the request is compared to the <b>scope_types</b> of the policy being enforced. If the scopes do not match, an <b>InvalidScope</b> exception will be raised. If <b>False</b> , a message will be logged informing operators that policies are being invoked with mismatching scope.
<b>policy_default_rule = default</b>	string value	Default rule. Enforced when a requested rule is not found.
<b>policy_dirs = [u'policy.d']</b>	multi valued	Directories where policy configuration files are stored. They can be relative to any directory in the search path defined by the <code>config_dir</code> option, or absolute paths. The file defined by <code>policy_file</code> must exist for these directories to be searched. Missing or empty directories are ignored.
<b>policy_file = policy.json</b>	string value	The file that defines policies.

Configuration option = Default value	Type	Description
<b>remote_content_type = application/x-www-form-urlencoded</b>	string value	Content Type to send and receive data for REST based policy check
<b>remote_ssl_ca_cert_file = None</b>	string value	Absolute path to ca cert file for REST based policy check
<b>remote_ssl_client_cert_file = None</b>	string value	Absolute path to client cert for REST based policy check
<b>remote_ssl_client_key_file = None</b>	string value	Absolute path client key file REST based policy check
<b>remote_ssl_verify_server_cert = False</b>	boolean value	server identity verification for REST based policy check

### 3.1.24. oslo\_reports

The following table outlines the options available under the **[oslo\_reports]** group in the `/etc/cinder/cinder.conf` file.

Table 3.24. oslo\_reports

Configuration option = Default value	Type	Description
<b>file_event_handler = None</b>	string value	The path to a file to watch for changes to trigger the reports, instead of signals. Setting this option disables the signal trigger for the reports. If application is running as a WSGI application it is recommended to use this instead of signals.
<b>file_event_handler_interval = 1</b>	integer value	How many seconds to wait between polls when <code>file_event_handler</code> is set
<b>log_dir = None</b>	string value	Path to a log directory where to create a file

### 3.1.25. oslo\_versionedobjects

The following table outlines the options available under the **[oslo\_versionedobjects]** group in the `/etc/cinder/cinder.conf` file.

Table 3.25. oslo\_versionedobjects

Configuration option = Default value	Type	Description
<b>fatal_exception_format_errors = False</b>	boolean value	Make exception message format errors fatal

### 3.1.26. profiler

The following table outlines the options available under the **[profiler]** group in the `/etc/cinder/cinder.conf` file.

Table 3.26. profiler

Configuration option = Default value	Type	Description
<b>connection_string = messaging://</b>	string value	<p>Connection string for a notifier backend. Default value is <code>messaging://</code> which sets the notifier to <code>oslo_messaging</code>.</p> <p>Examples of possible values:</p> <ul style="list-style-type: none"> <li>● <code>messaging://</code>: use <code>oslo_messaging</code> driver for sending notifications.</li> <li>● <code>mongodb://127.0.0.1:27017</code> : use <code>mongodb</code> driver for sending notifications.</li> <li>● <code>elasticsearch://127.0.0.1:9200</code> : use <code>elasticsearch</code> driver for sending notifications.</li> </ul>
<b>enabled = False</b>	boolean value	<p>Enables the profiling for all services on this node. Default value is <code>False</code> (fully disable the profiling feature).</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● <code>True</code>: Enables the feature</li> <li>● <code>False</code>: Disables the feature. The profiling cannot be started via this project operations. If the profiling is triggered by another project, this project part will be empty.</li> </ul>
<b>es_doc_type = notification</b>	string value	Document type for notification indexing in elasticsearch.
<b>es_scroll_size = 10000</b>	integer value	Elasticsearch splits large requests in batches. This parameter defines maximum size of each batch (for example: <code>es_scroll_size=10000</code> ).

Configuration option = Default value	Type	Description
<b>es_scroll_time = 2m</b>	string value	This parameter is a time value parameter (for example: es_scroll_time=2m), indicating for how long the nodes that participate in the search will maintain relevant resources in order to continue and support it.
<b>hmac_keys = SECRET_KEY</b>	string value	<p>Secret key(s) to use for encrypting context data for performance profiling. This string value should have the following format: &lt;key1&gt;[,&lt;key2&gt;,...&lt;keyn&gt;], where each key is some random string. A user who triggers the profiling via the REST API has to set one of these keys in the headers of the REST API call to include profiling results of this node for this particular project.</p> <p>Both "enabled" flag and "hmac_keys" config options should be set to enable profiling. Also, to generate correct profiling information across all services at least one key needs to be consistent between OpenStack projects. This ensures it can be used from client side to generate the trace, containing information from all possible resources.</p>
<b>sentinel_service_name = mymaster</b>	string value	Redis sentinel uses a service name to identify a master redis service. This parameter defines the name (for example: sentinel_service_name=mymaster).
<b>socket_timeout = 0.1</b>	floating point value	Redis sentinel provides a timeout option on the connections. This parameter defines that timeout (for example: socket_timeout=0.1).
<b>trace_sqlalchemy = False</b>	boolean value	<p>Enables SQL requests profiling in services. Default value is False (SQL requests won't be traced).</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• True: Enables SQL requests profiling. Each SQL query will be part of the trace and can be analyzed by how much time was spent for that.</li> <li>• False: Disables SQL requests profiling. The spent time is only shown on a higher level of operations. Single SQL queries cannot be analyzed this way.</li> </ul>

### 3.1.27. service\_user

The following table outlines the options available under the **[service\_user]** group in the **/etc/cinder/cinder.conf** file.

Table 3.27. `service_user`

Configuration option = Default value	Type	Description
<b>send_service_user_token = False</b>	boolean value	When True, if sending a user token to an REST API, also send a service token.

### 3.1.28. `ssl`

The following table outlines the options available under the `[ssl]` group in the `/etc/cinder/cinder.conf` file.

Table 3.28. `ssl`

Configuration option = Default value	Type	Description
<b>ca_file = None</b>	string value	CA certificate file to use to verify connecting clients.
<b>cert_file = None</b>	string value	Certificate file to use when starting the server securely.
<b>ciphers = None</b>	string value	Sets the list of available ciphers. value should be a string in the OpenSSL cipher list format.
<b>key_file = None</b>	string value	Private key file to use when starting the server securely.
<b>version = None</b>	string value	SSL version to use (valid only if SSL enabled). Valid values are TLSv1 and SSLv23. SSLv2, SSLv3, TLSv1_1, and TLSv1_2 may be available on some distributions.

### 3.1.29. `vault`

The following table outlines the options available under the `[vault]` group in the `/etc/cinder/cinder.conf` file.

Table 3.29. `vault`

Configuration option = Default value	Type	Description
<b>root_token_id = None</b>	string value	root token for vault
<b>ssl_ca_cert_file = None</b>	string value	Absolute path to ca cert file
<b>use_ssl = False</b>	boolean value	SSL Enabled/Disabled



Configuration option = Default value	Type	Description
<b>vault_url = http://127.0.0.1:8200</b>	string value	Use this endpoint to connect to Vault, for example: "http://127.0.0.1:8200"

## CHAPTER 4. DESIGNATE

The following chapter contains information about the configuration options in the **designate** service.

### 4.1. DESIGNATE.CONF

This section contains options for the `/etc/designate/designate.conf` file.

#### 4.1.1. DEFAULT

The following table outlines the options available under the **[DEFAULT]** group in the `/etc/designate/designate.conf` file.

Table 4.1. DEFAULT

Configuration option = Default value	Type	Description
<b>backdoor_port = None</b>	string value	Enable eventlet backdoor. Acceptable values are 0, <port>, and <start>:<end>, where 0 results in listening on a random tcp port number; <port> results in listening on the specified port number (and not enabling backdoor if that port is in use); and <start>:<end> results in listening on the smallest unused port number within the specified range of port numbers. The chosen port is displayed in the service's log file.
<b>backdoor_socket = None</b>	string value	Enable eventlet backdoor, using the provided path as a unix socket that can receive connections. This option is mutually exclusive with <i>backdoor_port</i> in that only one should be provided. If both are provided then the existence of this option overrides the usage of that option.
<b>backlog = 4096</b>	integer value	Number of backlog requests to configure the socket with
<b>central-topic = central</b>	string value	Central Topic
<b>conn_pool_min_size = 2</b>	integer value	The pool size limit for connections expiration policy
<b>conn_pool_ttl = 1200</b>	integer value	The time-to-live in sec of idle connections in the pool
<b>control_exchange = designate</b>	string value	The default exchange under which topics are scoped. May be overridden by an exchange name specified in the <i>transport_url</i> option.
<b>debug = False</b>	boolean value	If set to true, the logging level will be set to DEBUG instead of the default INFO level.

Configuration option = Default value	Type	Description
<b>default-soa-expire = 86400</b>	integer value	SOA expire
<b>default-soa-minimum = 3600</b>	integer value	SOA minimum value
<b>default-soa-refresh-max = 3600</b>	integer value	SOA max value
<b>default-soa-refresh-min = 3500</b>	integer value	SOA refresh-min value
<b>default-soa-retry = 600</b>	integer value	SOA retry
<b>default-ttl = 3600</b>	integer value	TTL Value
<b>default_log_levels = [u'amqplib=WARN', u'amqp=WARN', u'boto=WARN', u'eventlet.wsgi.server=WARN', u'iso8601=WARN', u'kazoo.client=WARN', u'keystone=INFO', u'keystonemiddleware.auth_token=INFO', u'oslo_messaging=WARN', u'oslo.messaging=INFO', u'oslo_service.loopingcall=WARN', u'sqlalchemy=WARN', u'stevedore=WARN', u'suds=INFO']</b>	list value	List of package logging levels in logger=LEVEL pairs. This option is ignored if log_config_append is set.
<b>executor_thread_pool_size = 64</b>	integer value	Size of executor thread pool when executor is threading or eventlet.
<b>fatal_deprecations = False</b>	boolean value	Enables or disables fatal status of deprecations.
<b>graceful_shutdown_timeout = 60</b>	integer value	Specify a timeout after which a gracefully shutdown server will exit. Zero value means endless wait.
<b>host = ccosp-undercloud.usersys.redhat.com</b>	string value	Name of this node

Configuration option = Default value	Type	Description
<code>`instance_format = [instance: %(uuid)s] `</code>	string value	The format for an instance that is passed with the log message.
<code>`instance_uuid_format = [instance: %(uuid)s] `</code>	string value	The format for an instance UUID that is passed with the log message.
<b>log-config-append = None</b>	string value	The name of a logging configuration file. This file is appended to any existing logging configuration files. For details about logging configuration files, see the Python logging module documentation. Note that when logging configuration files are used then all logging configuration is set in the configuration file and other logging configuration options are ignored (for example, <code>logging_context_format_string</code> ).
<b>log-date-format = %Y-%m-%d %H:%M:%S</b>	string value	Defines the format string for <code>%(asctime)s</code> in log records. Default: <code>%(default)s</code> . This option is ignored if <code>log_config_append</code> is set.
<b>log-dir = None</b>	string value	(Optional) The base directory used for relative <code>log_file</code> paths. This option is ignored if <code>log_config_append</code> is set.
<b>log-file = None</b>	string value	(Optional) Name of log file to send logging output to. If no default is set, logging will go to <code>stderr</code> as defined by <code>use_stderr</code> . This option is ignored if <code>log_config_append</code> is set.
<b>log_options = True</b>	boolean value	Enables or disables logging values of all registered options when starting a service (at <code>DEBUG</code> level).
<b>logging_context_format_string = %(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [%(request_id)s %(user_identity)s] %(instance)s%(message)s</b>	string value	Format string to use for log messages with context.
<b>logging_debug_format_s uffix = %(funcName)s % (pathname)s:%(lineno)d</b>	string value	Additional data to append to log message when logging level for the message is <code>DEBUG</code> .

Configuration option = Default value	Type	Description
<b>logging_default_format_string</b> = <code>%(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [-] %(instance)s%(message)s</code>	string value	Format string to use for log messages when context is undefined.
<b>logging_exception_prefix</b> = <code>%(asctime)s.%(msecs)03d %(process)d ERROR %(name)s %(instance)s</code>	string value	Prefix each line of exception output with this format.
<b>logging_user_identity_format</b> = <code>%(user)s %(tenant)s %(domain)s %(user_domain)s %(project_domain)s</code>	string value	Defines the format string for <code>%(user_identity)s</code> that is used in <code>logging_context_format_string</code> .
<b>mdns-topic</b> = <code>mdns</code>	string value	mDNS Topic
<b>network_api</b> = <code>neutron</code>	string value	Which API to use.
<b>notification-plugin</b> = <code>default</code>	string value	The notification plugin to use
<b>notify_api_faults</b> = <code>False</code>	boolean value	Send notifications if there's a failure in the API.
<b>pool-manager-topic</b> = <code>pool_manager</code>	string value	Pool Manager Topic
<b>publish_errors</b> = <code>False</code>	boolean value	Enables or disables publication of error events.
<b>pybasedir</b> = <code>/usr/lib/python2.7/site-packages</code>	string value	Directory where the designate python module is installed
<b>quota-api-export-size</b> = <code>1000</code>	integer value	Number of recordsets allowed in a zone export
<b>quota-driver</b> = <code>storage</code>	string value	Quota driver to use
<b>quota-recordset-records</b> = <code>20</code>	integer value	Number of records allowed per recordset

Configuration option = Default value	Type	Description
<b>quota-zone-records = 500</b>	integer value	Number of records allowed per zone
<b>quota-zone-recordsets = 500</b>	integer value	Number of recordsets allowed per zone
<b>quota-zones = 10</b>	integer value	Number of zones allowed per tenant
<b>rate_limit_burst = 0</b>	integer value	Maximum number of logged messages per rate_limit_interval.
<b>rate_limit_except_level = CRITICAL</b>	string value	Log level name used by rate limiting: CRITICAL, ERROR, INFO, WARNING, DEBUG or empty string. Logs with level greater or equal to rate_limit_except_level are not filtered. An empty string means that all levels are filtered.
<b>rate_limit_interval = 0</b>	integer value	Interval, number of seconds, of log rate limiting.
<b>root-helper = sudo designate-rootwrap /etc/designate/rootwrap.c onf</b>	string value	designate-rootwrap configuration
<b>rpc_ack_timeout_base = 15</b>	integer value	Number of seconds to wait for an ack from a cast/call. After each retry attempt this timeout is multiplied by some specified multiplier.
<b>rpc_ack_timeout_multiplier = 2</b>	integer value	Number to multiply base ack timeout by after each retry attempt.
<b>rpc_backend = rabbit</b>	string value	The messaging driver to use, defaults to rabbit. Other drivers include amqp and zmq.
<b>rpc_conn_pool_size = 30</b>	integer value	Size of RPC connection pool.
<b>rpc_message_ttl = 300</b>	integer value	Expiration timeout in seconds of a sent/received message after which it is not tracked anymore by a client/server.
<b>rpc_poll_timeout = 1</b>	integer value	The default number of seconds that poll should wait. Poll raises timeout exception when timeout expired.
<b>rpc_response_timeout = 60</b>	integer value	Seconds to wait for a response from a call.

Configuration option = Default value	Type	Description
<b>rpc_retry_attempts = 3</b>	integer value	Default number of message sending attempts in case of any problems occurred: positive value N means at most N retries, 0 means no retries, None or -1 (or any other negative values) mean to retry forever. This option is used only if acknowledgments are enabled.
<b>rpc_thread_pool_size = 100</b>	integer value	Maximum number of (green) threads to work concurrently.
<b>rpc_use_acks = False</b>	boolean value	Wait for message acknowledgements from receivers. This mechanism works only via proxy without PUB/SUB.
<b>rpc_zmq_bind_address = *</b>	string value	ZeroMQ bind address. Should be a wildcard (*), an ethernet interface, or IP. The "host" option should point or resolve to this address.
<b>rpc_zmq_bind_port_retries = 100</b>	integer value	Number of retries to find free port number before fail with ZMQBindError.
<b>rpc_zmq_contexts = 1</b>	integer value	Number of ZeroMQ contexts, defaults to 1.
<b>rpc_zmq_host = ccsosp-undercloud.usersys.redhat.com</b>	string value	Name of this node. Must be a valid hostname, FQDN, or IP address. Must match "host" option, if running Nova.
<b>rpc_zmq_ipc_dir = /var/run/openstack</b>	string value	Directory for holding IPC sockets.
<b>rpc_zmq_matchmaker = redis</b>	string value	MatchMaker driver.
<b>rpc_zmq_max_port = 65536</b>	integer value	Maximal port number for random ports range.
<b>rpc_zmq_min_port = 49153</b>	port value	Minimal port number for random ports range.
<b>rpc_zmq_serialization = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>rpc_zmq_topic_backlog = None</b>	integer value	Maximum number of ingress messages to locally buffer per topic. Default is unlimited.
<b>run_external_periodic_tasks = True</b>	boolean value	Some periodic tasks can be run in a separate process. Should we run them here?

Configuration option = Default value	Type	Description
<b>state-path =</b> <b>/var/lib/designate</b>	string value	Top-level directory for maintaining designate's state
<b>subscribe_on = []</b>	list value	List of publisher hosts SubConsumer can subscribe on. This option has higher priority then the default publishers list taken from the matchmaker.
<b>supported-record-type =</b> <b>[u'A', u'AAAA', u'CNAME',</b> <b>u'MX', u'SRV', u'TXT',</b> <b>u'SPF', u'NS', u'PTR',</b> <b>u'SSHFP', u'SOA']</b>	list value	Supported record types
<b>syslog-log-facility =</b> <b>LOG_USER</b>	string value	Syslog facility to receive log lines. This option is ignored if log_config_append is set.
<b>tcp_keepidle = 600</b>	integer value	Sets the value of TCP_KEEPIDLE in seconds for each server socket. Not supported on OS X.
<b>transport_url = None</b>	string value	The network address and optional user credentials for connecting to the messaging backend, in URL format. The expected format is:  driver://[user:pass@]host:port[, [userN:passN@]hostN:portN]/virtual_host?query  Example: rabbit://rabbitmq:password@127.0.0.1:5672//  For full details on the fields in the URL see the documentation of oslo_messaging.TransportURL at <a href="https://docs.openstack.org/oslo.messaging/latest/reference/transport.html">https://docs.openstack.org/oslo.messaging/latest/reference/transport.html</a>
<b>use-journal = False</b>	boolean value	Enable journald for logging. If running in a systemd environment you may wish to enable journal support. Doing so will use the journal native protocol which includes structured metadata in addition to log messages. This option is ignored if log_config_append is set.
<b>use-json = False</b>	boolean value	Use JSON formatting for logging. This option is ignored if log_config_append is set.
<b>use-syslog = False</b>	boolean value	Use syslog for logging. Existing syslog format is DEPRECATED and will be changed later to honor RFC5424. This option is ignored if log_config_append is set.



Configuration option = Default value	Type	Description
<b>use_dynamic_connections = False</b>	boolean value	This option makes direct connections dynamic or static. It makes sense only with use_router_proxy=False which means to use direct connections for direct message types (ignored otherwise).
<b>use_pub_sub = False</b>	boolean value	Use PUB/SUB pattern for fanout methods. PUB/SUB always uses proxy.
<b>use_router_proxy = False</b>	boolean value	Use ROUTER remote proxy.
<b>use_stderr = False</b>	boolean value	Log output to standard error. This option is ignored if log_config_append is set.
<b>watch-log-file = False</b>	boolean value	Uses logging handler designed to watch file system. When log file is moved or removed this handler will open a new log file with specified path instantaneously. It makes sense only if log_file option is specified and Linux platform is used. This option is ignored if log_config_append is set.
<b>worker-topic = worker</b>	string value	Worker Topic
<b>xfr_timeout = 10</b>	integer value	Timeout in seconds for XFR's.
<b>zmq_failover_connections = 2</b>	integer value	How many additional connections to a host will be made for failover reasons. This option is actual only in dynamic connections mode.
<b>zmq_immediate = True</b>	boolean value	This option configures round-robin mode in zmq socket. True means not keeping a queue when server side disconnects. False means to keep queue and messages even if server is disconnected, when the server appears we send all accumulated messages to it.
<b>zmq_linger = -1</b>	integer value	Number of seconds to wait before all pending messages will be sent after closing a socket. The default value of -1 specifies an infinite linger period. The value of 0 specifies no linger period. Pending messages shall be discarded immediately when the socket is closed. Positive values specify an upper bound for the linger period.
<b>zmq_target_expire = 300</b>	integer value	Expiration timeout in seconds of a name service record about existing target (< 0 means no timeout).

Configuration option = Default value	Type	Description
<b>zmq_target_update = 180</b>	integer value	Update period in seconds of a name service record about existing target.
<b>zmq_tcp_keepalive = -1</b>	integer value	Enable/disable TCP keepalive (KA) mechanism. The default value of -1 (or any other negative value) means to skip any overrides and leave it to OS default; 0 and 1 (or any other positive value) mean to disable and enable the option respectively.
<b>zmq_tcp_keepalive_cnt = -1</b>	integer value	The number of retransmissions to be carried out before declaring that remote end is not available. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_idle = -1</b>	integer value	The duration between two keepalive transmissions in idle condition. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_intvl = -1</b>	integer value	The duration between two successive keepalive retransmissions, if acknowledgement to the previous keepalive transmission is not received. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.

#### 4.1.2. backend:agent:bind9

The following table outlines the options available under the **[backend:agent:bind9]** group in the `/etc/designate/designate.conf` file.

Table 4.2. backend:agent:bind9

Configuration option = Default value	Type	Description
<b>query-destination = 127.0.0.1</b>	string value	Host to query when finding zones
<b>rndc-config-file = None</b>	string value	RNDC Config File
<b>rndc-host = 127.0.0.1</b>	string value	RNDC Host

Configuration option = Default value	Type	Description
<b>rndc-key-file = None</b>	string value	RNDC Key File
<b>rndc-port = 953</b>	integer value	RNDC Port
<b>zone-file-path = \$state_path/zones</b>	string value	Path where zone files are stored

### 4.1.3. backend:agent:denominator

The following table outlines the options available under the **[backend:agent:denominator]** group in the `/etc/designate/designate.conf` file.

Table 4.3. backend:agent:denominator

Configuration option = Default value	Type	Description
<b>config_file = /etc/denominator.conf</b>	string value	Path to Denominator configuration file
<b>name = fake</b>	string value	Name of the affected provider

### 4.1.4. backend:agent:djbdns

The following table outlines the options available under the **[backend:agent:djbdns]** group in the `/etc/designate/designate.conf` file.

Table 4.4. backend:agent:djbdns

Configuration option = Default value	Type	Description
<b>axfr-get-cmd-name = axfr-get</b>	string value	axfr-get executable path or rootwrap command name
<b>query-destination = 127.0.0.1</b>	string value	Host to query when finding zones
<b>tcpclient-cmd-name = tcpclient</b>	string value	tcpclient executable path or rootwrap command name
<b>tinydns-data-cmd-name = tinydns-data</b>	string value	tinydns-data executable path or rootwrap command name
<b>tinydns-datadir = /var/lib/djbdns</b>	string value	TinyDNS data directory

### 4.1.5. backend:agent:gdnsd

The following table outlines the options available under the **[backend:agent:gdnsd]** group in the **/etc/designate/designate.conf** file.

Table 4.5. backend:agent:gdnsd

Configuration option = Default value	Type	Description
<b>confdir-path = /etc/gdnsd</b>	string value	gdnsd configuration directory path
<b>gdnsd-cmd-name = gdnsd</b>	string value	gdnsd executable path or rootwrap command name
<b>query-destination = 127.0.0.1</b>	string value	Host to query when finding zones

### 4.1.6. backend:agent:knot2

The following table outlines the options available under the **[backend:agent:knot2]** group in the **/etc/designate/designate.conf** file.

Table 4.6. backend:agent:knot2

Configuration option = Default value	Type	Description
<b>knotc-cmd-name = knotc</b>	string value	knotc executable path or rootwrap command name
<b>query-destination = 127.0.0.1</b>	string value	Host to query when finding zones

### 4.1.7. coordination

The following table outlines the options available under the **[coordination]** group in the **/etc/designate/designate.conf** file.

Table 4.7. coordination

Configuration option = Default value	Type	Description
<b>backend_url = None</b>	string value	The backend URL to use for distributed coordination. If unset services that need coordination will function as a standalone service. This is a <b>tooz</b> url - see <a href="https://docs.openstack.org/tooz/latest/user/compatibility.html">https://docs.openstack.org/tooz/latest/user/compatibility.html</a>
<b>heartbeat_interval = 1.0</b>	floating point value	Number of seconds between heartbeats for distributed coordination.

Configuration option = Default value	Type	Description
<b>run_watchers_interval = 10.0</b>	floating point value	Number of seconds between checks to see if group membership has changed

#### 4.1.8. cors

The following table outlines the options available under the **[cors]** group in the `/etc/designate/designate.conf` file.

Table 4.8. cors

Configuration option = Default value	Type	Description
<b>allow_credentials = True</b>	boolean value	Indicate that the actual request can include user credentials
<b>allow_headers = [u'X-Auth-Token', u'X-Auth-Sudo-Tenant-ID', u'X-Auth-Sudo-Project-ID', u'X-Auth-All-Projects', u'X-Designate-Edit-Managed-Records', u'OpenStack-DNS-Hide-Counts']</b>	list value	Indicate which header field names may be used during the actual request.
<b>allow_methods = [u'GET', u'PUT', u'POST', u'DELETE', u'PATCH', u'HEAD']</b>	list value	Indicate which methods can be used during the actual request.
<b>allowed_origin = None</b>	list value	Indicate whether this resource may be shared with the domain received in the requests "origin" header. Format: "<protocol>://<host>[:<port>]", no trailing slash. Example: <a href="https://horizon.example.com">https://horizon.example.com</a>
<b>expose_headers = [u'X-OpenStack-Request-ID', u'Host']</b>	list value	Indicate which headers are safe to expose to the API. Defaults to HTTP Simple Headers.
<b>max_age = 3600</b>	integer value	Maximum cache age of CORS preflight requests.

#### 4.1.9. database

The following table outlines the options available under the **[database]** group in the `/etc/designate/designate.conf` file.

Table 4.9. database

Configuration option = Default value	Type	Description
<b>backend = sqlalchemy</b>	string value	The back end to use for the database.
<b>connection = None</b>	string value	The SQLAlchemy connection string to use to connect to the database.
<b>connection_debug = 0</b>	integer value	Verbosity of SQL debugging information: 0=None, 100=Everything.
<b>connection_recycle_time = 3600</b>	integer value	Connections which have been present in the connection pool longer than this number of seconds will be replaced with a new one the next time they are checked out from the pool.
<b>connection_trace = False</b>	boolean value	Add Python stack traces to SQL as comment strings.
<b>db_inc_retry_interval = True</b>	boolean value	If True, increases the interval between retries of a database operation up to <code>db_max_retry_interval</code> .
<b>db_max_retries = 20</b>	integer value	Maximum retries in case of connection error or deadlock error before error is raised. Set to -1 to specify an infinite retry count.
<b>db_max_retry_interval = 10</b>	integer value	If <code>db_inc_retry_interval</code> is set, the maximum seconds between retries of a database operation.
<b>db_retry_interval = 1</b>	integer value	Seconds between retries of a database transaction.
<b>max_overflow = 50</b>	integer value	If set, use this value for <code>max_overflow</code> with SQLAlchemy.
<b>max_pool_size = 5</b>	integer value	Maximum number of SQL connections to keep open in a pool. Setting a value of 0 indicates no limit.
<b>max_retries = 10</b>	integer value	Maximum number of database connection retries during startup. Set to -1 to specify an infinite retry count.
<b>min_pool_size = 1</b>	integer value	Minimum number of SQL connections to keep open in a pool.
<b>mysql_enable_ndb = False</b>	boolean value	If True, transparently enables support for handling MySQL Cluster (NDB).

Configuration option = Default value	Type	Description
<b>mysql_sql_mode = TRADITIONAL</b>	string value	The SQL mode to be used for MySQL sessions. This option, including the default, overrides any server-set SQL mode. To use whatever SQL mode is set by the server configuration, set this to no value. Example: <code>mysql_sql_mode=</code>
<b>pool_timeout = None</b>	integer value	If set, use this value for <code>pool_timeout</code> with SQLAlchemy.
<b>retry_interval = 10</b>	integer value	Interval between retries of opening a SQL connection.
<b>slave_connection = None</b>	string value	The SQLAlchemy connection string to use to connect to the slave database.
<b>sqlite_synchronous = True</b>	boolean value	If True, SQLite uses synchronous mode.
<b>use_db_reconnect = False</b>	boolean value	Enable the experimental use of database reconnect on connection lost.

#### 4.1.10. handler:neutron\_floatingip

The following table outlines the options available under the **[handler:neutron\_floatingip]** group in the `/etc/designate/designate.conf` file.

Table 4.10. handler:neutron\_floatingip

Configuration option = Default value	Type	Description
<b>control-exchange = neutron</b>	string value	control-exchange for neutron notification
<b>format = None</b>	multi valued	format which replaced by <code>formatv4/formatv6</code>
<b>formatv4 = None</b>	multi valued	IPv4 format
<b>formatv6 = None</b>	multi valued	IPv6 format
<b>notification-topics = [u'notifications']</b>	list value	notification any events from neutron
<b>zone-id = None</b>	string value	Zone ID with each notification

### 4.1.11. handler:nova\_fixed

The following table outlines the options available under the **[handler:nova\_fixed]** group in the `/etc/designate/designate.conf` file.

Table 4.11. handler:nova\_fixed

Configuration option = Default value	Type	Description
<b>control-exchange = nova</b>	string value	control-exchange for nova notification
<b>format = None</b>	multi valued	format which replaced by formatv4/formatv6
<b>formatv4 = None</b>	multi valued	IPv4 format
<b>formatv6 = None</b>	multi valued	IPv6 format
<b>notification-topics = [u'notifications']</b>	list value	notification any events from nova
<b>zone-id = None</b>	string value	Zone ID with each notification

### 4.1.12. healthcheck

The following table outlines the options available under the **[healthcheck]** group in the `/etc/designate/designate.conf` file.

Table 4.12. healthcheck

Configuration option = Default value	Type	Description
<b>backends = []</b>	list value	Additional backends that can perform health checks and report that information back as part of a request.
<b>detailed = False</b>	boolean value	Show more detailed information as part of the response
<b>disable_by_file_path = None</b>	string value	Check the presence of a file to determine if an application is running on a port. Used by DisableByFileHealthcheck plugin.
<b>disable_by_file_paths = []</b>	list value	Check the presence of a file based on a port to determine if an application is running on a port. Expects a "port:path" list of strings. Used by DisableByFilesPortsHealthcheck plugin.
<b>path = /healthcheck</b>	string value	The path to respond to healthcheck requests on.



### 4.1.13. heartbeat\_emitter

The following table outlines the options available under the **[heartbeat\_emitter]** group in the `/etc/designate/designate.conf` file.

Table 4.13. heartbeat\_emitter

Configuration option = Default value	Type	Description
<b>emitter_type = rpc</b>	string value	Emitter to use
<b>heartbeat_interval = 5.0</b>	floating point value	Number of seconds between heartbeats for reporting state

### 4.1.14. keystone\_authtoken

The following table outlines the options available under the **[keystone\_authtoken]** group in the `/etc/designate/designate.conf` file.

Table 4.14. keystone\_authtoken

Configuration option = Default value	Type	Description
<b>auth_section = None</b>	string value	Config Section from which to load plugin specific options
<b>auth_type = None</b>	string value	Authentication type to load
<b>auth_uri = None</b>	string value	Complete "public" Identity API endpoint. This endpoint should not be an "admin" endpoint, as it should be accessible by all end users. Unauthenticated clients are redirected to this endpoint to authenticate. Although this endpoint should ideally be unversioned, client support in the wild varies. If you're using a versioned v2 endpoint here, then this should <b>not</b> be the same endpoint the service user utilizes for validating tokens, because normal end users may not be able to reach that endpoint. This option is deprecated in favor of <code>www_authenticate_uri</code> and will be removed in the S release.
<b>auth_version = None</b>	string value	API version of the admin Identity API endpoint.
<b>cache = None</b>	string value	Request environment key where the Swift cache object is stored. When <code>auth_token</code> middleware is deployed with a Swift cache, use this option to have the middleware share a caching backend with swift. Otherwise, use the <b>memcached_servers</b> option instead.

Configuration option = Default value	Type	Description
<b>cafile = None</b>	string value	A PEM encoded Certificate Authority to use when verifying HTTPs connections. Defaults to system CAs.
<b>certfile = None</b>	string value	Required if identity server requires client certificate
<b>check_revocations_for_cached = False</b>	boolean value	If true, the revocation list will be checked for cached tokens. This requires that PKI tokens are configured on the identity server.
<b>delay_auth_decision = False</b>	boolean value	Do not handle authorization requests within the middleware, but delegate the authorization decision to downstream WSGI components.
<b>enforce_token_bind = permissive</b>	string value	Used to control the use and type of token binding. Can be set to: "disabled" to not check token binding. "permissive" (default) to validate binding information if the bind type is of a form known to the server and ignore it if not. "strict" like "permissive" but if the bind type is unknown the token will be rejected. "required" any form of token binding is needed to be allowed. Finally the name of a binding method that must be present in tokens.
<b>hash_algorithms = [u'md5']</b>	list value	Hash algorithms to use for hashing PKI tokens. This may be a single algorithm or multiple. The algorithms are those supported by Python standard hashlib.new(). The hashes will be tried in the order given, so put the preferred one first for performance. The result of the first hash will be stored in the cache. This will typically be set to multiple values only while migrating from a less secure algorithm to a more secure one. Once all the old tokens are expired this option should be set to a single value for better performance.
<b>http_connect_timeout = None</b>	integer value	Request timeout value for communicating with Identity API server.
<b>http_request_max_retries = 3</b>	integer value	How many times are we trying to reconnect when communicating with Identity API Server.
<b>include_service_catalog = True</b>	boolean value	(Optional) Indicate whether to set the X-Service-Catalog header. If False, middleware will not ask for service catalog on token validation and will not set the X-Service-Catalog header.
<b>insecure = False</b>	boolean value	Verify HTTPS connections.

Configuration option = Default value	Type	Description
<b>keyfile = None</b>	string value	Required if identity server requires client certificate
<b>memcache_pool_conn_get_timeout = 10</b>	integer value	(Optional) Number of seconds that an operation will wait to get a memcached client connection from the pool.
<b>memcache_pool_dead_retry = 300</b>	integer value	(Optional) Number of seconds memcached server is considered dead before it is tried again.
<b>memcache_pool_maxsize = 10</b>	integer value	(Optional) Maximum total number of open connections to every memcached server.
<b>memcache_pool_socket_timeout = 3</b>	integer value	(Optional) Socket timeout in seconds for communicating with a memcached server.
<b>memcache_pool_unused_timeout = 60</b>	integer value	(Optional) Number of seconds a connection to memcached is held unused in the pool before it is closed.
<b>memcache_secret_key = None</b>	string value	(Optional, mandatory if memcache_security_strategy is defined) This string is used for key derivation.
<b>memcache_security_strategy = None</b>	string value	(Optional) If defined, indicate whether token data should be authenticated or authenticated and encrypted. If MAC, token data is authenticated (with HMAC) in the cache. If ENCRYPT, token data is encrypted and authenticated in the cache. If the value is not one of these options or empty, auth_token will raise an exception on initialization.
<b>memcache_use_advanced_pool = False</b>	boolean value	(Optional) Use the advanced (eventlet safe) memcached client pool. The advanced pool will only work under python 2.x.
<b>memcached_servers = None</b>	list value	Optionally specify a list of memcached server(s) to use for caching. If left undefined, tokens will instead be cached in-process.
<b>region_name = None</b>	string value	The region in which the identity server can be found.

Configuration option = Default value	Type	Description
<b>revocation_cache_time = 10</b>	integer value	Determines the frequency at which the list of revoked tokens is retrieved from the Identity service (in seconds). A high number of revocation events combined with a low cache duration may significantly reduce performance. Only valid for PKI tokens. This option has been deprecated in the Ocata release and will be removed in the P release.
<b>service_token_roles = [u'service']</b>	list value	A choice of roles that must be present in a service token. Service tokens are allowed to request that an expired token can be used and so this check should tightly control that only actual services should be sending this token. Roles here are applied as an ANY check so any role in this list must be present. For backwards compatibility reasons this currently only affects the allow_expired check.
<b>service_token_roles_required = False</b>	boolean value	For backwards compatibility reasons we must let valid service tokens pass that don't pass the service_token_roles check as valid. Setting this true will become the default in a future release and should be enabled if possible.
<b>signing_dir = None</b>	string value	Directory used to cache files related to PKI tokens. This option has been deprecated in the Ocata release and will be removed in the P release.
<b>token_cache_time = 300</b>	integer value	In order to prevent excessive effort spent validating tokens, the middleware caches previously-seen tokens for a configurable duration (in seconds). Set to -1 to disable caching completely.
<b>www_authenticate_uri = None</b>	string value	Complete "public" Identity API endpoint. This endpoint should not be an "admin" endpoint, as it should be accessible by all end users. Unauthenticated clients are redirected to this endpoint to authenticate. Although this endpoint should ideally be unversioned, client support in the wild varies. If you're using a versioned v2 endpoint here, then this should <b>not</b> be the same endpoint the service user utilizes for validating tokens, because normal end users may not be able to reach that endpoint.

#### 4.1.15. matchmaker\_redis

The following table outlines the options available under the **[matchmaker\_redis]** group in the **/etc/designate/designate.conf** file.

Table 4.15. matchmaker\_redis

Configuration option = Default value	Type	Description
<b>check_timeout = 20000</b>	integer value	Time in ms to wait before the transaction is killed.
<b>host = 127.0.0.1</b>	string value	Host to locate redis.
<code>`password = `</code>	string value	Password for Redis server (optional).
<b>port = 6379</b>	port value	Use this port to connect to redis host.
<b>sentinel_group_name = oslo-messaging-zeromq</b>	string value	Redis replica set name.
<b>sentinel_hosts = []</b>	list value	List of Redis Sentinel hosts (fault tolerance mode), e.g., [host:port, host1:port ... ]
<b>socket_timeout = 10000</b>	integer value	Timeout in ms on blocking socket operations.
<b>wait_timeout = 2000</b>	integer value	Time in ms to wait between connection attempts.

#### 4.1.16. monasca:statsd

The following table outlines the options available under the **[monasca:statsd]** group in the `/etc/designate/designate.conf` file.

Table 4.16. monasca:statsd

Configuration option = Default value	Type	Description
<b>enabled = False</b>	boolean value	enable
<b>hostname = 127.0.0.1</b>	string value	hostname
<b>port = 8125</b>	integer value	UDP port

#### 4.1.17. network\_api:neutron

The following table outlines the options available under the **[network\_api:neutron]** group in the `/etc/designate/designate.conf` file.

Table 4.17. network\_api:neutron

Configuration option = Default value	Type	Description
<b>admin_password = None</b>	string value	password for connecting to neutron in admin context
<b>admin_tenant_name = None</b>	string value	tenant name for connecting to neutron in admin context
<b>admin_username = None</b>	string value	username for connecting to neutron in admin context
<b>auth_strategy = keystone</b>	string value	auth strategy for connecting to neutron in admin context
<b>auth_url = None</b>	string value	auth url for connecting to neutron in admin context
<b>ca_certificates_file = None</b>	string value	Location of ca certificates file to use for neutron client requests.
<b>endpoint_type = publicURL</b>	string value	Endpoint type to use
<b>endpoints = None</b>	list value	URL to use if None in the ServiceCatalog that is passed by the request context. Format: <region> <url>
<b>insecure = False</b>	boolean value	if set, ignore any SSL validation issues
<b>timeout = 30</b>	integer value	timeout value for connecting to neutron in seconds

#### 4.1.18. oslo\_concurrency

The following table outlines the options available under the **[oslo\_concurrency]** group in the `/etc/designate/designate.conf` file.

Table 4.18. oslo\_concurrency

Configuration option = Default value	Type	Description
<b>disable_process_locking = False</b>	boolean value	Enables or disables inter-process locks.

Configuration option = Default value	Type	Description
<b>lock_path = \$state_path</b>	string value	Directory to use for lock files. For security, the specified directory should only be writable by the user running the processes that need locking. Defaults to environment variable OSLO_LOCK_PATH. If external locks are used, a lock path must be set.

#### 4.1.19. oslo\_messaging\_amqp

The following table outlines the options available under the **[oslo\_messaging\_amqp]** group in the `/etc/designate/designate.conf` file.

Table 4.19. oslo\_messaging\_amqp

Configuration option = Default value	Type	Description
<b>addressing_mode = dynamic</b>	string value	Indicates the addressing mode used by the driver. Permitted values: <i>legacy</i> - use legacy non-routable addressing <i>routable</i> - use routable addresses <i>dynamic</i> - use legacy addresses if the message bus does not support routing otherwise use routable addressing
<b>allow_insecure_clients = False</b>	boolean value	Accept clients using either SSL or plain TCP
<b>anycast_address = anycast</b>	string value	Appended to the address prefix when sending to a group of consumers. Used by the message bus to identify messages that should be delivered in a round-robin fashion across consumers.
<b>broadcast_prefix = broadcast</b>	string value	address prefix used when broadcasting to all servers
<b>connection_retry_backoff = 2</b>	integer value	Increase the connection_retry_interval by this many seconds after each unsuccessful failover attempt.
<b>connection_retry_interval = 1</b>	integer value	Seconds to pause before attempting to re-connect.
<b>connection_retry_interval_max = 30</b>	integer value	Maximum limit for connection_retry_interval + connection_retry_backoff
<b>container_name = None</b>	string value	Name for the AMQP container. must be globally unique. Defaults to a generated UUID

Configuration option = Default value	Type	Description
<b>default_notification_exchange = None</b>	string value	Exchange name used in notification addresses. Exchange name resolution precedence: Target.exchange if set else default_notification_exchange if set else control_exchange if set else <i>notify</i>
<b>default_notify_timeout = 30</b>	integer value	The deadline for a sent notification message delivery. Only used when caller does not provide a timeout expiry.
<b>default_reply_retry = 0</b>	integer value	The maximum number of attempts to re-send a reply message which failed due to a recoverable error.
<b>default_reply_timeout = 30</b>	integer value	The deadline for an rpc reply message delivery.
<b>default_rpc_exchange = None</b>	string value	Exchange name used in RPC addresses. Exchange name resolution precedence: Target.exchange if set else default_rpc_exchange if set else control_exchange if set else <i>rpc</i>
<b>default_send_timeout = 30</b>	integer value	The deadline for an rpc cast or call message delivery. Only used when caller does not provide a timeout expiry.
<b>default_sender_link_timeout = 600</b>	integer value	The duration to schedule a purge of idle sender links. Detach link after expiry.
<b>group_request_prefix = unicast</b>	string value	address prefix when sending to any server in group
<b>idle_timeout = 0</b>	integer value	Timeout for inactive connections (in seconds)
<b>link_retry_delay = 10</b>	integer value	Time to pause between re-connecting an AMQP 1.0 link that failed due to a recoverable error.
<b>multicast_address = multicast</b>	string value	Appended to the address prefix when sending a fanout message. Used by the message bus to identify fanout messages.
<b>notify_address_prefix = openstack.org/om/notify</b>	string value	Address prefix for all generated Notification addresses
<b>notify_server_credit = 100</b>	integer value	Window size for incoming Notification messages



Configuration option = Default value	Type	Description
<code>`password = `</code>	string value	Password for message broker authentication
<b>pre_settled = [u'rpc-cast', u'rpc-reply']</b>	multi valued	Send messages of this type pre-settled. Pre-settled messages will not receive acknowledgement from the peer. Note well: pre-settled messages may be silently discarded if the delivery fails. Permitted values: <i>rpc-call</i> - send RPC Calls pre-settled <i>rpc-reply</i> - send RPC Replies pre-settled <i>rpc-cast</i> - Send RPC Casts pre-settled <i>notify</i> - Send Notifications pre-settled
<b>pseudo_vhost = True</b>	boolean value	Enable virtual host support for those message buses that do not natively support virtual hosting (such as <i>qpidd</i> ). When set to true the virtual host name will be added to all message bus addresses, effectively creating a private <i>subnet</i> per virtual host. Set to False if the message bus supports virtual hosting using the <i>hostname</i> field in the AMQP 1.0 Open performative as the name of the virtual host.
<b>reply_link_credit = 200</b>	integer value	Window size for incoming RPC Reply messages.
<b>rpc_address_prefix = openstack.org/om/rpc</b>	string value	Address prefix for all generated RPC addresses
<b>rpc_server_credit = 100</b>	integer value	Window size for incoming RPC Request messages
<code>`sas_config_dir = `</code>	string value	Path to directory that contains the SASL configuration
<code>`sas_config_name = `</code>	string value	Name of configuration file (without <i>.conf</i> suffix)
<code>`sas_default_realm = `</code>	string value	SASL realm to use if no realm present in username
<code>`sas_mechanisms = `</code>	string value	Space separated list of acceptable SASL mechanisms
<b>server_request_prefix = exclusive</b>	string value	address prefix used when sending to a specific server
<b>ssl = False</b>	boolean value	Attempt to connect via SSL. If no other <i>ssl</i> -related parameters are given, it will use the system's CA-bundle to verify the server's certificate.
<code>`ssl_ca_file = `</code>	string value	CA certificate PEM file used to verify the server's certificate

Configuration option = Default value	Type	Description
<code>`ssl_cert_file = `</code>	string value	Self-identifying certificate PEM file for client authentication
<code>`ssl_key_file = `</code>	string value	Private key PEM file used to sign <code>ssl_cert_file</code> certificate (optional)
<b>ssl_key_password = None</b>	string value	Password for decrypting <code>ssl_key_file</code> (if encrypted)
<b>ssl_verify_vhost = False</b>	boolean value	By default SSL checks that the name in the server's certificate matches the hostname in the <code>transport_url</code> . In some configurations it may be preferable to use the virtual hostname instead, for example if the server uses the Server Name Indication TLS extension (rfc6066) to provide a certificate per virtual host. Set <code>ssl_verify_vhost</code> to True if the server's SSL certificate uses the virtual host name instead of the DNS name.
<b>trace = False</b>	boolean value	Debug: dump AMQP frames to stdout
<b>unicast_address = unicast</b>	string value	Appended to the address prefix when sending to a particular RPC/Notification server. Used by the message bus to identify messages sent to a single destination.
<code>`username = `</code>	string value	User name for message broker authentication

#### 4.1.20. oslo\_messaging\_kafka

The following table outlines the options available under the **[oslo\_messaging\_kafka]** group in the `/etc/designate/designate.conf` file.

Table 4.20. oslo\_messaging\_kafka

Configuration option = Default value	Type	Description
<b>conn_pool_min_size = 2</b>	integer value	The pool size limit for connections expiration policy
<b>conn_pool_ttl = 1200</b>	integer value	The time-to-live in sec of idle connections in the pool
<b>consumer_group = oslo_messaging_consumer</b>	string value	Group id for Kafka consumer. Consumers in one group will coordinate message consumption

Configuration option = Default value	Type	Description
<b>kafka_consumer_timeout = 1.0</b>	floating point value	Default timeout(s) for Kafka consumers
<b>kafka_default_host = localhost</b>	string value	Default Kafka broker Host
<b>kafka_default_port = 9092</b>	port value	Default Kafka broker Port
<b>kafka_max_fetch_bytes = 1048576</b>	integer value	Max fetch bytes of Kafka consumer
<b>pool_size = 10</b>	integer value	Pool Size for Kafka Consumers
<b>producer_batch_size = 16384</b>	integer value	Size of batch for the producer async send
<b>producer_batch_timeout = 0.0</b>	floating point value	Upper bound on the delay for KafkaProducer batching in seconds

#### 4.1.21. oslo\_messaging\_notifications

The following table outlines the options available under the **[oslo\_messaging\_notifications]** group in the **/etc/designate/designate.conf** file.

Table 4.21. oslo\_messaging\_notifications

Configuration option = Default value	Type	Description
<b>driver = []</b>	multi valued	The Drivers(s) to handle sending notifications. Possible values are messaging, messagingv2, routing, log, test, noop
<b>retry = -1</b>	integer value	The maximum number of attempts to re-send a notification message which failed to be delivered due to a recoverable error. 0 - No retry, -1 - indefinite
<b>topics = [u'notifications']</b>	list value	AMQP topic used for OpenStack notifications.
<b>transport_url = None</b>	string value	A URL representing the messaging driver to use for notifications. If not set, we fall back to the same configuration used for RPC.

#### 4.1.22. oslo\_messaging\_rabbit

The following table outlines the options available under the `[oslo_messaging_rabbit]` group in the `/etc/designate/designate.conf` file.

Table 4.22. oslo\_messaging\_rabbit

Configuration option = Default value	Type	Description
<b>amqp_auto_delete = False</b>	boolean value	Auto-delete queues in AMQP.
<b>amqp_durable_queues = False</b>	boolean value	Use durable queues in AMQP.
<b>channel_max = None</b>	integer value	Maximum number of channels to allow
<b>connection_factory = single</b>	string value	Connection factory implementation
<b>default_notification_exchange = \${control_exchange}_notification</b>	string value	Exchange name for sending notifications
<b>default_notification_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending notification, -1 means infinite retry.
<b>default_rpc_exchange = \${control_exchange}_rpc</b>	string value	Exchange name for sending RPC messages
<b>default_rpc_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending RPC message, -1 means infinite retry. If actual retry attempts in not 0 the rpc request could be processed more than one time
<b>default_serializer_type = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>fake_rabbit = False</b>	boolean value	Deprecated, use <code>rpc_backend=kombu+memory</code> or <code>rpc_backend=fake</code>
<b>frame_max = None</b>	integer value	The maximum byte size for an AMQP frame
<b>heartbeat_interval = 3</b>	integer value	How often to send heartbeats for consumer's connections
<b>heartbeat_rate = 2</b>	integer value	How often times during the <code>heartbeat_timeout_threshold</code> we check the heartbeat.

Configuration option = Default value	Type	Description
<b>heartbeat_timeout_thresh old = 60</b>	integer value	Number of seconds after which the Rabbit broker is considered down if heartbeat's keep-alive fails (0 disable the heartbeat). EXPERIMENTAL
<b>host_connection_reconn ect_delay = 0.25</b>	floating point value	Set delay for reconnection to some host which has connection error
<b>kombu_compression = None</b>	string value	EXPERIMENTAL: Possible values are: gzip, bz2. If not set compression will not be used. This option may not be available in future versions.
<b>kombu_failover_strategy = round-robin</b>	string value	Determines how the next RabbitMQ node is chosen in case the one we are currently connected to becomes unavailable. Takes effect only if more than one RabbitMQ node is provided in config.
<b>kombu_missing_consum er_retry_timeout = 60</b>	integer value	How long to wait a missing client before abandoning to send it its replies. This value should not be longer than <code>rpc_response_timeout</code> .
<b>kombu_reconnect_delay = 1.0</b>	floating point value	How long to wait before reconnecting in response to an AMQP consumer cancel notification.
<b>notification_listener_pref etch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to notification listener.
<b>notification_persistence = False</b>	boolean value	Persist notification messages.
<b>notification_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending notification message
<b>pool_max_overflow = 0</b>	integer value	Maximum number of connections to create above <b>pool_max_size</b> .
<b>pool_max_size = 30</b>	integer value	Maximum number of connections to keep queued.
<b>pool_recycle = 600</b>	integer value	Lifetime of a connection (since creation) in seconds or None for no recycling. Expired connections are closed on acquire.
<b>pool_stale = 60</b>	integer value	Threshold at which inactive (since release) connections are considered stale in seconds or None for no staleness. Stale connections are closed on acquire.

Configuration option = Default value	Type	Description
<b>pool_timeout = 30</b>	integer value	Default number of seconds to wait for a connections to available
<b>rabbit_ha_queues = False</b>	boolean value	Try to use HA queues in RabbitMQ (x-ha-policy: all). If you change this option, you must wipe the RabbitMQ database. In RabbitMQ 3.0, queue mirroring is no longer controlled by the x-ha-policy argument when declaring a queue. If you just want to make sure that all queues (except those with auto-generated names) are mirrored across all nodes, run: <code>rabbitmqctl set_policy HA ^(?!amq\).* {"ha-mode": "all"}</code>
<b>rabbit_host = localhost</b>	string value	The RabbitMQ broker address where a single node is used.
<b>rabbit_hosts = [u'\$rabbit_host:\$rabbit_port']</b>	list value	RabbitMQ HA cluster host:port pairs.
<b>rabbit_interval_max = 30</b>	integer value	Maximum interval of RabbitMQ connection retries. Default is 30 seconds.
<b>rabbit_login_method = AMQPLAIN</b>	string value	The RabbitMQ login method.
<b>rabbit_max_retries = 0</b>	integer value	Maximum number of RabbitMQ connection retries. Default is 0 (infinite retry count).
<b>rabbit_password = guest</b>	string value	The RabbitMQ password.
<b>rabbit_port = 5672</b>	port value	The RabbitMQ broker port where a single node is used.
<b>rabbit_qos_prefetch_count = 0</b>	integer value	Specifies the number of messages to prefetch. Setting to zero allows unlimited messages.
<b>rabbit_retry_backoff = 2</b>	integer value	How long to backoff for between retries when connecting to RabbitMQ.
<b>rabbit_retry_interval = 1</b>	integer value	How frequently to retry connecting with RabbitMQ.
<b>rabbit_transient_queues_ttl = 1800</b>	integer value	Positive integer representing duration in seconds for queue TTL (x-expires). Queues which are unused for the duration of the TTL are automatically deleted. The parameter affects only reply and fanout queues.

Configuration option = Default value	Type	Description
<b>rabbit_userid = guest</b>	string value	The RabbitMQ userid.
<b>rabbit_virtual_host = /</b>	string value	The RabbitMQ virtual host.
<b>rpc_listener_prefetch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to rpc listener.
<b>rpc_queue_expiration = 60</b>	integer value	Time to live for rpc queues without consumers in seconds.
<b>rpc_reply_exchange = \${control_exchange}_rpc_reply</b>	string value	Exchange name for receiving RPC replies
<b>rpc_reply_listener_prefetch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to rpc reply listener.
<b>rpc_reply_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending reply. -1 means infinite retry during rpc_timeout
<b>rpc_reply_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending reply.
<b>rpc_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending RPC message
<b>socket_timeout = 0.25</b>	floating point value	Set socket timeout in seconds for connection's socket
<b>ssl = None</b>	boolean value	Enable SSL
<code>`ssl_ca_file = `</code>	string value	SSL certification authority file (valid only if SSL enabled).
<code>`ssl_cert_file = `</code>	string value	SSL cert file (valid only if SSL enabled).
<code>`ssl_key_file = `</code>	string value	SSL key file (valid only if SSL enabled).
<b>ssl_options = None</b>	dict value	Arguments passed to ssl.wrap_socket
<code>`ssl_version = `</code>	string value	SSL version to use (valid only if SSL enabled). Valid values are TLSv1 and SSLv23. SSLv2, SSLv3, TLSv1_1, and TLSv1_2 may be available on some distributions.

Configuration option = Default value	Type	Description
<b>tcp_user_timeout = 0.25</b>	floating point value	Set TCP_USER_TIMEOUT in seconds for connection's socket

#### 4.1.23. oslo\_messaging\_zmq

The following table outlines the options available under the **[oslo\_messaging\_zmq]** group in the **/etc/designate/designate.conf** file.

Table 4.23. oslo\_messaging\_zmq

Configuration option = Default value	Type	Description
<b>rpc_ack_timeout_base = 15</b>	integer value	Number of seconds to wait for an ack from a cast/call. After each retry attempt this timeout is multiplied by some specified multiplier.
<b>rpc_ack_timeout_multiplier = 2</b>	integer value	Number to multiply base ack timeout by after each retry attempt.
<b>rpc_message_ttl = 300</b>	integer value	Expiration timeout in seconds of a sent/received message after which it is not tracked anymore by a client/server.
<b>rpc_poll_timeout = 1</b>	integer value	The default number of seconds that poll should wait. Poll raises timeout exception when timeout expired.
<b>rpc_retry_attempts = 3</b>	integer value	Default number of message sending attempts in case of any problems occurred: positive value N means at most N retries, 0 means no retries, None or -1 (or any other negative values) mean to retry forever. This option is used only if acknowledgments are enabled.
<b>rpc_thread_pool_size = 100</b>	integer value	Maximum number of (green) threads to work concurrently.
<b>rpc_use_acks = False</b>	boolean value	Wait for message acknowledgements from receivers. This mechanism works only via proxy without PUB/SUB.
<b>rpc_zmq_bind_address = *</b>	string value	ZeroMQ bind address. Should be a wildcard (*), an ethernet interface, or IP. The "host" option should point or resolve to this address.
<b>rpc_zmq_bind_port_retries = 100</b>	integer value	Number of retries to find free port number before fail with ZMQBindError.



Configuration option = Default value	Type	Description
<b>rpc_zmq_contexts = 1</b>	integer value	Number of ZeroMQ contexts, defaults to 1.
<b>rpc_zmq_host = ccosp-undercloud.usersys.redhat.com</b>	string value	Name of this node. Must be a valid hostname, FQDN, or IP address. Must match "host" option, if running Nova.
<b>rpc_zmq_ipc_dir = /var/run/openstack</b>	string value	Directory for holding IPC sockets.
<b>rpc_zmq_matchmaker = redis</b>	string value	MatchMaker driver.
<b>rpc_zmq_max_port = 65536</b>	integer value	Maximal port number for random ports range.
<b>rpc_zmq_min_port = 49153</b>	port value	Minimal port number for random ports range.
<b>rpc_zmq_serialization = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>rpc_zmq_topic_backlog = None</b>	integer value	Maximum number of ingress messages to locally buffer per topic. Default is unlimited.
<b>subscribe_on = []</b>	list value	List of publisher hosts SubConsumer can subscribe on. This option has higher priority then the default publishers list taken from the matchmaker.
<b>use_dynamic_connections = False</b>	boolean value	This option makes direct connections dynamic or static. It makes sense only with use_router_proxy=False which means to use direct connections for direct message types (ignored otherwise).
<b>use_pub_sub = False</b>	boolean value	Use PUB/SUB pattern for fanout methods. PUB/SUB always uses proxy.
<b>use_router_proxy = False</b>	boolean value	Use ROUTER remote proxy.
<b>zmq_failover_connections = 2</b>	integer value	How many additional connections to a host will be made for failover reasons. This option is actual only in dynamic connections mode.

Configuration option = Default value	Type	Description
<b>zmq_immediate = True</b>	boolean value	This option configures round-robin mode in zmq socket. True means not keeping a queue when server side disconnects. False means to keep queue and messages even if server is disconnected, when the server appears we send all accumulated messages to it.
<b>zmq_linger = -1</b>	integer value	Number of seconds to wait before all pending messages will be sent after closing a socket. The default value of -1 specifies an infinite linger period. The value of 0 specifies no linger period. Pending messages shall be discarded immediately when the socket is closed. Positive values specify an upper bound for the linger period.
<b>zmq_target_expire = 300</b>	integer value	Expiration timeout in seconds of a name service record about existing target (< 0 means no timeout).
<b>zmq_target_update = 180</b>	integer value	Update period in seconds of a name service record about existing target.
<b>zmq_tcp_keepalive = -1</b>	integer value	Enable/disable TCP keepalive (KA) mechanism. The default value of -1 (or any other negative value) means to skip any overrides and leave it to OS default; 0 and 1 (or any other positive value) mean to disable and enable the option respectively.
<b>zmq_tcp_keepalive_cnt = -1</b>	integer value	The number of retransmissions to be carried out before declaring that remote end is not available. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_idle = -1</b>	integer value	The duration between two keepalive transmissions in idle condition. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_intvl = -1</b>	integer value	The duration between two successive keepalive retransmissions, if acknowledgement to the previous keepalive transmission is not received. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.

#### 4.1.24. oslo\_middlewares

The following table outlines the options available under the **[oslo\_middlewares]** group in the `/etc/designate/designate.conf` file.

Table 4.24. oslo\_middlewares

Configuration option = Default value	Type	Description
<b>enable_proxy_headers_parsing = False</b>	boolean value	Whether the application is behind a proxy or not. This determines if the middleware should parse the headers or not.
<b>max_request_body_size = 114688</b>	integer value	The maximum body size for each request, in bytes.
<b>secure_proxy_ssl_header = X-Forwarded-Proto</b>	string value	The HTTP Header that will be used to determine what the original request protocol scheme was, even if it was hidden by a SSL termination proxy.

#### 4.1.25. oslo\_policy

The following table outlines the options available under the **[oslo\_policy]** group in the `/etc/designate/designate.conf` file.

Table 4.25. oslo\_policy

Configuration option = Default value	Type	Description
<b>enforce_scope = False</b>	boolean value	This option controls whether or not to enforce scope when evaluating policies. If <b>True</b> , the scope of the token used in the request is compared to the <b>scope_types</b> of the policy being enforced. If the scopes do not match, an <b>InvalidScope</b> exception will be raised. If <b>False</b> , a message will be logged informing operators that policies are being invoked with mismatching scope.
<b>policy_default_rule = default</b>	string value	Default rule. Enforced when a requested rule is not found.
<b>policy_dirs = [u'policy.d']</b>	multi valued	Directories where policy configuration files are stored. They can be relative to any directory in the search path defined by the <code>config_dir</code> option, or absolute paths. The file defined by <code>policy_file</code> must exist for these directories to be searched. Missing or empty directories are ignored.
<b>policy_file = policy.json</b>	string value	The file that defines policies.

Configuration option = Default value	Type	Description
<b>remote_content_type = application/x-www-form-urlencoded</b>	string value	Content Type to send and receive data for REST based policy check
<b>remote_ssl_ca_cert_file = None</b>	string value	Absolute path to ca cert file for REST based policy check
<b>remote_ssl_client_cert_file = None</b>	string value	Absolute path to client cert for REST based policy check
<b>remote_ssl_client_key_file = None</b>	string value	Absolute path client key file REST based policy check
<b>remote_ssl_verify_server_cert = False</b>	boolean value	server identity verification for REST based policy check

#### 4.1.26. pool\_manager\_cache:memcache

The following table outlines the options available under the **[pool\_manager\_cache:memcache]** group in the `/etc/designate/designate.conf` file.

Table 4.26. pool\_manager\_cache:memcache

Configuration option = Default value	Type	Description
<b>expiration = 3600</b>	integer value	Time in seconds to expire cache.
<b>memcached_servers = None</b>	list value	Memcached servers or None for in process cache.

#### 4.1.27. pool\_manager\_cache:sqlalchemy

The following table outlines the options available under the **[pool\_manager\_cache:sqlalchemy]** group in the `/etc/designate/designate.conf` file.

Table 4.27. pool\_manager\_cache:sqlalchemy

Configuration option = Default value	Type	Description
<b>backend = sqlalchemy</b>	string value	The back end to use for the database.
<b>connection = None</b>	string value	The SQLAlchemy connection string to use to connect to the database.

Configuration option = Default value	Type	Description
<b>connection_debug = 0</b>	integer value	Verbosity of SQL debugging information: 0=None, 100=Everything.
<b>connection_recycle_time = 3600</b>	integer value	Connections which have been present in the connection pool longer than this number of seconds will be replaced with a new one the next time they are checked out from the pool.
<b>connection_trace = False</b>	boolean value	Add Python stack traces to SQL as comment strings.
<b>db_inc_retry_interval = True</b>	boolean value	If True, increases the interval between retries of a database operation up to db_max_retry_interval.
<b>db_max_retries = 20</b>	integer value	Maximum retries in case of connection error or deadlock error before error is raised. Set to -1 to specify an infinite retry count.
<b>db_max_retry_interval = 10</b>	integer value	If db_inc_retry_interval is set, the maximum seconds between retries of a database operation.
<b>db_retry_interval = 1</b>	integer value	Seconds between retries of a database transaction.
<b>max_overflow = 50</b>	integer value	If set, use this value for max_overflow with SQLAlchemy.
<b>max_pool_size = 5</b>	integer value	Maximum number of SQL connections to keep open in a pool. Setting a value of 0 indicates no limit.
<b>max_retries = 10</b>	integer value	Maximum number of database connection retries during startup. Set to -1 to specify an infinite retry count.
<b>min_pool_size = 1</b>	integer value	Minimum number of SQL connections to keep open in a pool.
<b>mysql_enable_ndb = False</b>	boolean value	If True, transparently enables support for handling MySQL Cluster (NDB).
<b>mysql_sql_mode = TRADITIONAL</b>	string value	The SQL mode to be used for MySQL sessions. This option, including the default, overrides any server-set SQL mode. To use whatever SQL mode is set by the server configuration, set this to no value. Example: mysql_sql_mode=
<b>pool_timeout = None</b>	integer value	If set, use this value for pool_timeout with SQLAlchemy.

Configuration option = Default value	Type	Description
<b>retry_interval = 10</b>	integer value	Interval between retries of opening a SQL connection.
<b>slave_connection = None</b>	string value	The SQLAlchemy connection string to use to connect to the slave database.
<b>sqlite_synchronous = True</b>	boolean value	If True, SQLite uses synchronous mode.
<b>use_db_reconnect = False</b>	boolean value	Enable the experimental use of database reconnect on connection lost.

#### 4.1.28. producer\_task:delayed\_notify

The following table outlines the options available under the **[producer\_task:delayed\_notify]** group in the `/etc/designate/designate.conf` file.

Table 4.28. producer\_task:delayed\_notify

Configuration option = Default value	Type	Description
<b>batch_size = 100</b>	integer value	How many zones to receive NOTIFY on each run
<b>interval = 5</b>	integer value	Run interval in seconds
<b>per_page = 100</b>	integer value	Default amount of results returned per page

#### 4.1.29. producer\_task:periodic\_exists

The following table outlines the options available under the **[producer\_task:periodic\_exists]** group in the `/etc/designate/designate.conf` file.

Table 4.29. producer\_task:periodic\_exists

Configuration option = Default value	Type	Description
<b>interval = 3600</b>	integer value	Run interval in seconds
<b>per_page = 100</b>	integer value	Default amount of results returned per page

#### 4.1.30. producer\_task:periodic\_secondary\_refresh

The following table outlines the options available under the **[producer\_task:periodic\_secondary\_refresh]** group in the `/etc/designate/designate.conf` file.

Table 4.30. `producer_task:periodic_secondary_refresh`

Configuration option = Default value	Type	Description
<b>interval = 3600</b>	integer value	Run interval in seconds
<b>per_page = 100</b>	integer value	Default amount of results returned per page

#### 4.1.31. `producer_task:worker_periodic_recovery`

The following table outlines the options available under the **[producer\_task:worker\_periodic\_recovery]** group in the `/etc/designate/designate.conf` file.

Table 4.31. `producer_task:worker_periodic_recovery`

Configuration option = Default value	Type	Description
<b>interval = 120</b>	integer value	Run interval in seconds
<b>per_page = 100</b>	integer value	Default amount of results returned per page

#### 4.1.32. `producer_task:zone_purge`

The following table outlines the options available under the **[producer\_task:zone\_purge]** group in the `/etc/designate/designate.conf` file.

Table 4.32. `producer_task:zone_purge`

Configuration option = Default value	Type	Description
<b>batch_size = 100</b>	integer value	How many zones to be purged on each run
<b>interval = 3600</b>	integer value	Run interval in seconds
<b>per_page = 100</b>	integer value	Default amount of results returned per page
<b>time_threshold = 604800</b>	integer value	How old deleted zones should be ( <code>deleted_at</code> ) to be purged, in seconds

#### 4.1.33. `proxy`

The following table outlines the options available under the **[proxy]** group in the `/etc/designate/designate.conf` file.

Table 4.33. proxy

Configuration option = Default value	Type	Description
<b>http_proxy = None</b>	string value	Proxy HTTP requests via this proxy.
<b>https_proxy = None</b>	string value	Proxy HTTPS requests via this proxy
<b>no_proxy = []</b>	list value	These addresses should not be proxied

#### 4.1.34. service:agent

The following table outlines the options available under the **[service:agent]** group in the `/etc/designate/designate.conf` file.

Table 4.34. service:agent

Configuration option = Default value	Type	Description
<b>allow-notify = []</b>	list value	List of IP addresses allowed to NOTIFY The Agent
<b>backend-driver = bind9</b>	string value	The backend driver to use, e.g. bind9, djbdns, knot2
<b>host = None</b>	IP address value	Agent Bind Host
<b>listen = [u'0.0.0.0:5358']</b>	list value	Agent host:port pairs to listen on
<b>masters = []</b>	list value	List of masters for the Agent, format ip:port
<b>notify-delay = 0.0</b>	floating point value	Delay after a NOTIFY arrives for a zone that the Agent will pause and drop subsequent NOTIFYs for that zone
<b>port = None</b>	port value	Agent Port Number
<b>tcp-backlog = 100</b>	integer value	The Agent TCP Backlog
<b>tcp-recv-timeout = 0.5</b>	floating point value	Agent TCP Receive Timeout
<b>threads = 1000</b>	integer value	Number of agent greenthreads to spawn
<b>transfer-source = None</b>	string value	An IP address to be used to fetch zones transferred in
<b>workers = None</b>	integer value	Number of agent worker processes to spawn



### 4.1.35. service:api

The following table outlines the options available under the **[service:api]** group in the `/etc/designate/designate.conf` file.

Table 4.35. service:api

Configuration option = Default value	Type	Description
<b>api-base-uri = http://127.0.0.1:9001/</b>	string value	the url used as the base for all API responses, This should consist of the scheme (http/https), the hostname, port, and any paths that are added to the base of Designate is URLs, For example <a href="http://dns.openstack.example.com/dns">http://dns.openstack.example.com/dns</a>
<b>api_host = None</b>	IP address value	API Bind Host
<b>api_paste_config = api-paste.ini</b>	string value	File name for the paste.deploy config for designate-api
<b>api_port = None</b>	port value	API Port Number
<b>auth_strategy = keystone</b>	string value	The strategy to use for auth. Supports noauth or keystone
<b>default-limit-admin = 20</b>	integer value	Default per-page limit for the Admin API, a value of None means show all results by default
<b>default-limit-v2 = 20</b>	integer value	Default per-page limit for the V2 API, a value of None means show all results by default
<b>enable-api-admin = False</b>	boolean value	enable-api-admin
<b>enable-api-v2 = True</b>	boolean value	enable-api-v2 which enable in a future
<b>enable-host-header = False</b>	boolean value	Enable host request headers
<b>enabled-extensions-admin = []</b>	list value	Enabled Admin API Extensions
<b>enabled-extensions-v2 = []</b>	list value	Enabled API Extensions for the V2 API
<b>listen = [u'0.0.0.0:9001']</b>	list value	API host:port pairs to listen on
<b>maintenance-mode = False</b>	boolean value	Enable API Maintenance Mode

Configuration option = Default value	Type	Description
<b>maintenance-mode-role = admin</b>	string value	Role allowed to bypass maintenance mode
<b>max-limit-admin = 1000</b>	integer value	Max per-page limit for the Admin API
<b>max-limit-v2 = 1000</b>	integer value	Max per-page limit for the V2 API
<b>max_header_line = 16384</b>	integer value	Maximum line size of message headers to be accepted. <code>max_header_line</code> may need to be increased when using large tokens (typically those generated by the Keystone v3 API with big service catalogs).
<b>override-proto = None</b>	string value	A scheme that will be used to override the request protocol scheme, even if it was set by an SSL terminating proxy.
<b>secure-proxy-ssl-header = X-Forwarded-Proto</b>	string value	The HTTP Header that will be used to determine which the original request protocol scheme was, even if it was removed by an SSL terminating proxy.
<b>threads = 1000</b>	integer value	Number of api greenthreads to spawn
<b>workers = None</b>	integer value	Number of api worker processes to spawn

#### 4.1.36. service:central

The following table outlines the options available under the **[service:central]** group in the `/etc/designate/designate.conf` file.

Table 4.36. service:central

Configuration option = Default value	Type	Description
<b>central_topic = central</b>	string value	RPC topic name of central service.
<b>default_pool_id = 794ccc2c-d751-44fe-b57f-8894c9f5c842</b>	string value	The name of the default pool
<b>enabled-notification-handlers = []</b>	list value	Enabled Notification Handlers

Configuration option = Default value	Type	Description
<b>managed_resource_email</b> = <b>hostmaster@example.com</b>	string value	E-Mail for Managed resources
<b>managed_resource_tenant_id</b> = <b>00000000-0000-0000-0000-000000000000</b>	string value	The Tenant ID that will own any managed resources.
<b>max_recordset_name_len</b> = <b>255</b>	integer value	Maximum recordset name length
<b>max_zone_name_len</b> = <b>255</b>	integer value	Maximum zone name length
<b>min_ttl</b> = <b>None</b>	integer value	Minimum TTL allowed
<b>scheduler_filters</b> = <b>[u'default_pool']</b>	list value	Enabled Pool Scheduling filters
<b>storage-driver</b> = <b>sqlalchemy</b>	string value	The storage driver to use
<b>threads</b> = <b>1000</b>	integer value	Number of central greenthreads to spawn
<b>workers</b> = <b>None</b>	integer value	Number of central worker processes to spawn

#### 4.1.37. service:mdns

The following table outlines the options available under the **[service:mdns]** group in the **/etc/designate/designate.conf** file.

Table 4.37. service:mdns

Configuration option = Default value	Type	Description
<b>all-tcp</b> = <b>False</b>	boolean value	Send all traffic over TCP
<b>host</b> = <b>None</b>	IP address value	mDNS Bind Host
<b>listen</b> = <b>[u'0.0.0.0:5354']</b>	list value	mDNS host:port pairs to listen on

Configuration option = Default value	Type	Description
<b>max-message-size = 65535</b>	integer value	Maximum message size to emit
<b>mdns_topic = mdns</b>	string value	RPC topic name for mini-DNS
<b>port = None</b>	port value	mDNS Port Number
<b>query-enforce-tsig = False</b>	boolean value	Enforce all incoming queries (including AXFR) are TSIG signed
<b>storage-driver = sqlalchemy</b>	string value	The storage driver to use
<b>tcp-backlog = 100</b>	integer value	mDNS TCP Backlog
<b>tcp-recv-timeout = 0.5</b>	floating point value	mDNS TCP Receive Timeout
<b>threads = 1000</b>	integer value	Number of mdns greenthreads to spawn
<b>workers = None</b>	integer value	Number of mdns worker processes to spawn

#### 4.1.38. service:pool\_manager

The following table outlines the options available under the **[service:pool\_manager]** group in the `/etc/designate/designate.conf` file.

Table 4.38. service:pool\_manager

Configuration option = Default value	Type	Description
<b>cache-driver = memcache</b>	string value	The cache driver to use
<b>enable-recovery-timer = True</b>	boolean value	The flag for the recovery timer
<b>enable-sync-timer = True</b>	boolean value	The flag for the sync timer
<b>periodic-recovery-interval = 120</b>	integer value	The time between recovering from failures
<b>periodic-sync-interval = 1800</b>	integer value	The time between synchronizing the servers with storage

Configuration option = Default value	Type	Description
<b>periodic-sync-max-attempts = 3</b>	integer value	Number of attempts to update a zone during sync
<b>periodic-sync-retry-interval = 30</b>	integer value	Interval between zone update attempts during sync
<b>periodic-sync-seconds = 21600</b>	integer value	Zones Updated within last N seconds will be syncd. Use an empty value to sync all zones.
<b>poll-delay = 5</b>	integer value	The time to wait before sending the first request to a server
<b>poll-max-retries = 10</b>	integer value	The maximum number of times to retry sending a request and wait for a response from a server
<b>poll-retry-interval = 15</b>	integer value	The time between retrying to send a request and waiting for a response from a server
<b>poll-timeout = 30</b>	integer value	The time to wait for a response from a server
<b>pool-id = 794ccc2c-d751-44fe-b57f-8894c9f5c842</b>	string value	The ID of the pool managed by this instance of the Pool Manager
<b>pool_manager_topic = pool_manager</b>	string value	RPC topic name for pool-manager
<b>threads = 1000</b>	integer value	Number of Pool Manager greenthreads to spawn
<b>threshold-percentage = 100</b>	integer value	The percentage of servers requiring a successful update for a zone change to be considered active
<b>workers = None</b>	integer value	Number of Pool Manager worker processes to spawn

#### 4.1.39. service:producer

The following table outlines the options available under the **[service:producer]** group in the `/etc/designate/designate.conf` file.

Table 4.39. service:producer

Configuration option = Default value	Type	Description
<b>enabled_tasks = None</b>	list value	Enabled tasks to run

Configuration option = Default value	Type	Description
<b>export-synchronous = True</b>	boolean value	Whether to allow synchronous zone exports
<b>storage-driver = sqlalchemy</b>	string value	The storage driver to use
<b>threads = 1000</b>	integer value	Number of Producer greenthreads to spawn
<b>workers = None</b>	integer value	Number of Producer worker processes to spawn

#### 4.1.40. service:sink

The following table outlines the options available under the **[service:sink]** group in the **/etc/designate/designate.conf** file.

Table 4.40. service:sink

Configuration option = Default value	Type	Description
<b>enabled-notification-handlers = []</b>	list value	Enabled Notification Handlers
<b>threads = 1000</b>	integer value	Number of sink greenthreads to spawn
<b>workers = None</b>	integer value	Number of sink worker processes to spawn

#### 4.1.41. service:worker

The following table outlines the options available under the **[service:worker]** group in the **/etc/designate/designate.conf** file.

Table 4.41. service:worker

Configuration option = Default value	Type	Description
<b>enabled = False</b>	boolean value	Whether to send events to worker instead of Pool Manager
<b>export-synchronous = True</b>	boolean value	Whether to allow synchronous zone exports

Configuration option = Default value	Type	Description
<b>notify = True</b>	boolean value	Whether to allow worker to send NOTIFYs, this will noop NOTIFYs in mdns if true
<b>poll-delay = 5</b>	integer value	The time to wait before sending the first request to a server
<b>poll-max-retries = 10</b>	integer value	The maximum number of times to retry sending a request and wait for a response from a server
<b>poll-retry-interval = 15</b>	integer value	The time between retrying to send a request and waiting for a response from a server
<b>poll-timeout = 30</b>	integer value	The time to wait for a response from a server
<b>storage-driver = sqlalchemy</b>	string value	The storage driver to use
<b>threads = 200</b>	integer value	Number of Worker threads to spawn per process
<b>threshold-percentage = 100</b>	integer value	The percentage of servers requiring a successful update for a domain change to be considered active
<b>worker_topic = worker</b>	string value	RPC topic for worker component
<b>workers = None</b>	integer value	Number of Worker worker processes to spawn

#### 4.1.42. service:zone\_manager

The following table outlines the options available under the **[service:zone\_manager]** group in the `/etc/designate/designate.conf` file.

Table 4.42. service:zone\_manager

Configuration option = Default value	Type	Description
<b>enabled_tasks = None</b>	list value	Enabled tasks to run
<b>export-synchronous = True</b>	boolean value	Whether to allow synchronous zone exports
<b>storage-driver = sqlalchemy</b>	string value	The storage driver to use
<b>threads = 1000</b>	integer value	Number of Zone Manager greenthreads to spawn

Configuration option = Default value	Type	Description
<b>workers = None</b>	integer value	Number of Zone Manager worker processes to spawn

#### 4.1.43. ssl

The following table outlines the options available under the **[ssl]** group in the `/etc/designate/designate.conf` file.

Table 4.43. ssl

Configuration option = Default value	Type	Description
<b>ca_file = None</b>	string value	CA certificate file to use to verify connecting clients.
<b>cert_file = None</b>	string value	Certificate file to use when starting the server securely.
<b>ciphers = None</b>	string value	Sets the list of available ciphers. value should be a string in the OpenSSL cipher list format.
<b>key_file = None</b>	string value	Private key file to use when starting the server securely.
<b>version = None</b>	string value	SSL version to use (valid only if SSL enabled). Valid values are TLSv1 and SSLv23. SSLv2, SSLv3, TLSv1_1, and TLSv1_2 may be available on some distributions.

#### 4.1.44. storage:sqlalchemy

The following table outlines the options available under the **[storage:sqlalchemy]** group in the `/etc/designate/designate.conf` file.

Table 4.44. storage:sqlalchemy

Configuration option = Default value	Type	Description
<b>backend = sqlalchemy</b>	string value	The back end to use for the database.
<b>connection = None</b>	string value	The SQLAlchemy connection string to use to connect to the database.
<b>connection_debug = 0</b>	integer value	Verbosity of SQL debugging information: 0=None, 100=Everything.



Configuration option = Default value	Type	Description
<b>connection_recycle_time = 3600</b>	integer value	Connections which have been present in the connection pool longer than this number of seconds will be replaced with a new one the next time they are checked out from the pool.
<b>connection_trace = False</b>	boolean value	Add Python stack traces to SQL as comment strings.
<b>db_inc_retry_interval = True</b>	boolean value	If True, increases the interval between retries of a database operation up to db_max_retry_interval.
<b>db_max_retries = 20</b>	integer value	Maximum retries in case of connection error or deadlock error before error is raised. Set to -1 to specify an infinite retry count.
<b>db_max_retry_interval = 10</b>	integer value	If db_inc_retry_interval is set, the maximum seconds between retries of a database operation.
<b>db_retry_interval = 1</b>	integer value	Seconds between retries of a database transaction.
<b>max_overflow = 50</b>	integer value	If set, use this value for max_overflow with SQLAlchemy.
<b>max_pool_size = 5</b>	integer value	Maximum number of SQL connections to keep open in a pool. Setting a value of 0 indicates no limit.
<b>max_retries = 10</b>	integer value	Maximum number of database connection retries during startup. Set to -1 to specify an infinite retry count.
<b>min_pool_size = 1</b>	integer value	Minimum number of SQL connections to keep open in a pool.
<b>mysql_enable_ndb = False</b>	boolean value	If True, transparently enables support for handling MySQL Cluster (NDB).
<b>mysql_sql_mode = TRADITIONAL</b>	string value	The SQL mode to be used for MySQL sessions. This option, including the default, overrides any server-set SQL mode. To use whatever SQL mode is set by the server configuration, set this to no value. Example: mysql_sql_mode=
<b>pool_timeout = None</b>	integer value	If set, use this value for pool_timeout with SQLAlchemy.
<b>retry_interval = 10</b>	integer value	Interval between retries of opening a SQL connection.

Configuration option = Default value	Type	Description
<b>slave_connection = None</b>	string value	The SQLAlchemy connection string to use to connect to the slave database.
<b>sqlite_synchronous = True</b>	boolean value	If True, SQLite uses synchronous mode.
<b>use_db_reconnect = False</b>	boolean value	Enable the experimental use of database reconnect on connection lost.

## CHAPTER 5. GLANCE

The following chapter contains information about the configuration options in the **glance** service.

### 5.1. GLANCE-API.CONF

This section contains options for the `/etc/glance/glance-api.conf` file.

#### 5.1.1. DEFAULT


The following table outlines the options available under the **[DEFAULT]** group in the `/etc/glance/glance-api.conf` file.

Table 5.1. DEFAULT

Configuration option = Default value	Type	Description
<b>admin_password = None</b>	string value	The administrators password. If "use_user_token" is not in effect, then admin credentials can be specified.
<b>admin_role = admin</b>	string value	<p>Role used to identify an authenticated user as administrator.</p> <p>Provide a string value representing a Keystone role to identify an administrative user. Users with this role will be granted administrative privileges. The default value for this option is <i>admin</i>.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A string value which is a valid Keystone role</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<b>admin_tenant_name = None</b>	string value	The tenant name of the administrative user. If "use_user_token" is not in effect, then admin tenant name can be specified.
<b>admin_user = None</b>	string value	The administrators user name. If "use_user_token" is not in effect, then admin credentials can be specified.

Configuration option = Default value	Type	Description
<b>allow_additional_image_properties = True</b>	boolean value	<p>Allow users to add additional/custom properties to images.</p> <p>Glance defines a standard set of properties (in its schema) that appear on every image. These properties are also known as <b>base properties</b>. In addition to these properties, Glance allows users to add custom properties to images. These are known as <b>additional properties</b>.</p> <p>By default, this configuration option is set to <b>True</b> and users are allowed to add additional properties. The number of additional properties that can be added to an image can be controlled via <b>image_property_quota</b> configuration option.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• True</li> <li>• False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• image_property_quota</li> </ul>
<b>allow_anonymous_accesses = False</b>	boolean value	<p>Allow limited access to unauthenticated users.</p> <p>Assign a boolean to determine API access for unauthenticated users. When set to False, the API cannot be accessed by unauthenticated users. When set to True, unauthenticated users can access the API with read-only privileges. This however only applies when using ContextMiddleware.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• True</li> <li>• False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>

Configuration option = Default value	Type	Description
<b>allowed_rpc_exception_modules =</b> <b>[u'glance.common.exception', u'builtins', u'exceptions']</b>	list value	<p>List of allowed exception modules to handle RPC exceptions.</p> <p>Provide a comma separated list of modules whose exceptions are permitted to be recreated upon receiving exception data via an RPC call made to Glance. The default list includes <b>glance.common.exception, builtins, and exceptions.</b></p> <p>The RPC protocol permits interaction with Glance via calls across a network or within the same system. Including a list of exception namespaces with this option enables RPC to propagate the exceptions back to the users.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A comma separated list of valid exception modules</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>


Configuration option = Default value	Type	Description
<b>api_limit_max = 1000</b>	integer value	<p>Maximum number of results that could be returned by a request.</p> <p>As described in the help text of <b>limit_param_default</b>, some requests may return multiple results. The number of results to be returned are governed either by the <b>limit</b> parameter in the request or the <b>limit_param_default</b> configuration option. The value in either case, can't be greater than the absolute maximum defined by this configuration option. Anything greater than this value is trimmed down to the maximum value defined here.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p>Setting this to a very large value may slow down database queries and increase response times. Setting this to a very low value may result in poor user experience.</p> </div> </div> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any positive integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● limit_param_default</li> </ul>
<b>auth_region = None</b>	string value	The region for the authentication service. If "use_user_token" is not in effect and using keystone auth, then region name can be specified.
<b>auth_strategy = noauth</b>	string value	The strategy to use for authentication. If "use_user_token" is not in effect, then auth strategy can be specified.
<b>auth_url = None</b>	string value	The URL to the keystone service. If "use_user_token" is not in effect and using keystone auth, then URL of keystone can be specified.

Configuration option = Default value	Type	Description
<b>backlog = 4096</b>	integer value	<p>Set the number of incoming connection requests.</p> <p>Provide a positive integer value to limit the number of requests in the backlog queue. The default queue size is 4096.</p> <p>An incoming connection to a TCP listener socket is queued before a connection can be established with the server. Setting the backlog for a TCP socket ensures a limited queue size for incoming traffic.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Positive integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>bind_host = 0.0.0.0</b>	unknown value	<p>IP address to bind the glance servers to.</p> <p>Provide an IP address to bind the glance server to. The default value is <b>0.0.0.0</b>.</p> <p>Edit this option to enable the server to listen on one particular IP address on the network card. This facilitates selection of a particular network interface for the server.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid IPv4 address</li> <li>● A valid IPv6 address</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>


Configuration option = Default value	Type	Description
<b>bind_port = None</b>	port value	<p>Port number on which the server will listen.</p> <p>Provide a valid port number to bind the server's socket to. This port is then set to identify processes and forward network messages that arrive at the server. The default bind_port value for the API server is 9292 and for the registry server is 9191.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid port number (0 to 65535)</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>ca_file = None</b>	string value	<p>Absolute path to the CA file.</p> <p>Provide a string value representing a valid absolute path to the Certificate Authority file to use for client authentication.</p> <p>A CA file typically contains necessary trusted certificates to use for the client authentication. This is essential to ensure that a secure connection is established to the server via the internet.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Valid absolute path to the CA file</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>




Configuration option = Default value	Type	Description
<b>cert_file = None</b>	string value	<p>Absolute path to the certificate file.</p> <p>Provide a string value representing a valid absolute path to the certificate file which is required to start the API service securely.</p> <p>A certificate file typically is a public key container and includes the server's public key, server name, server information and the signature which was a result of the verification process using the CA certificate. This is required for a secure connection establishment.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Valid absolute path to the certificate file</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<b>client_socket_timeout = 900</b>	integer value	<p>Timeout for client connections' socket operations.</p> <p>Provide a valid integer value representing time in seconds to set the period of wait before an incoming connection can be closed. The default value is 900 seconds.</p> <p>The value zero implies wait forever.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Zero</li> <li>• Positive integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<b>conn_pool_min_size = 2</b>	integer value	The pool size limit for connections expiration policy
<b>conn_pool_ttl = 1200</b>	integer value	The time-to-live in sec of idle connections in the pool
<b>control_exchange = openstack</b>	string value	The default exchange under which topics are scoped. May be overridden by an exchange name specified in the transport_url option.

Configuration option = Default value	Type	Description
<b>data_api = glance.db.sqlalchemy.api</b>	string value	<p>Python module path of data access API.</p> <p>Specifies the path to the API to use for accessing the data model. This option determines how the image catalog data will be accessed.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● <code>glance.db.sqlalchemy.api</code></li> <li>● <code>glance.db.registry.api</code></li> <li>● <code>glance.db.simple.api</code></li> </ul> <p>If this option is set to <b>glance.db.sqlalchemy.api</b> then the image catalog data is stored in and read from the database via the SQLAlchemy Core and ORM APIs.</p> <p>Setting this option to <b>glance.db.registry.api</b> will force all database access requests to be routed through the Registry service. This avoids data access from the Glance API nodes for an added layer of security, scalability and manageability.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p>In v2 OpenStack Images API, the registry service is optional. In order to use the Registry API in v2, the option <b>enable_v2_registry</b> must be set to <b>True</b>.</p> </div> </div> <p>Finally, when this configuration option is set to <b>glance.db.simple.api</b>, image catalog data is stored in and read from an in-memory data structure. This is primarily used for testing.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>enable_v2_api</code></li> <li>● <code>enable_v2_registry</code></li> </ul>
<b>debug = False</b>	boolean value	<p>If set to true, the logging level will be set to DEBUG instead of the default INFO level.</p>

Configuration option = Default value	Type	Description
<b>default_log_levels =</b> <b>[u'amqp=WARN',</b> <b>u'amqplib=WARN',</b> <b>u'boto=WARN',</b> <b>u'qpid=WARN',</b> <b>u'sqlalchemy=WARN',</b> <b>u'suds=INFO',</b> <b>u'oslo.messaging=INFO',</b> <b>u'oslo_messaging=INFO',</b> <b>u'iso8601=WARN',</b> <b>u'requests.packages.urlli</b> <b>b3.connectionpool=WAR</b> <b>N',</b> <b>u'urllib3.connectionpool=</b> <b>WARN',</b> <b>u'websocket=WARN',</b> <b>u'requests.packages.urlli</b> <b>b3.util.retry=WARN',</b> <b>u'urllib3.util.retry=WARN',</b> <b>u'keystonemiddleware=W</b> <b>ARN',</b> <b>u'routes.middleware=WA</b> <b>RN', u'stevedore=WARN',</b> <b>u'taskflow=WARN',</b> <b>u'keystoneauth=WARN',</b> <b>u'oslo.cache=INFO',</b> <b>u'dogpile.core.dogpile=IN</b> <b>FO']</b>	list value	<p>List of package logging levels in logger=LEVEL pairs. This option is ignored if log_config_append is set.</p>
<b>default_publisher_id =</b> <b>image.localhost</b>	string value	<p>Default publisher_id for outgoing Glance notifications.</p> <p>This is the value that the notification driver will use to identify messages for events originating from the Glance service. Typically, this is the hostname of the instance that generated the message.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any reasonable instance identifier, for example: image.host1</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>

Configuration option = Default value	Type	Description
<b>delayed_delete = False</b>	boolean value	<p>Turn on/off delayed delete.</p> <p>Typically when an image is deleted, the <b>glance-api</b> service puts the image into <b>deleted</b> state and deletes its data at the same time. Delayed delete is a feature in Glance that delays the actual deletion of image data until a later point in time (as determined by the configuration option <b>scrub_time</b>). When delayed delete is turned on, the <b>glance-api</b> service puts the image into <b>pending_delete</b> state upon deletion and leaves the image data in the storage backend for the image scrubber to delete at a later time. The image scrubber will move the image into <b>deleted</b> state upon successful deletion of image data.</p> <div data-bbox="815 831 922 1055" style="display: inline-block; vertical-align: top;">  </div> <p><b>NOTE</b></p> <p>When delayed delete is turned on, image scrubber <b>MUST</b> be running as a periodic task to prevent the backend storage from filling up with undesired usage.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>scrub_time</b></li> <li>● <b>wakeup_time</b></li> <li>● <b>scrub_pool_size</b></li> </ul>

Configuration option = Default value	Type	Description
<b>digest_algorithm = sha256</b>	string value	<p>Digest algorithm to use for digital signature.</p> <p>Provide a string value representing the digest algorithm to use for generating digital signatures. By default, <b>sha256</b> is used.</p> <p>To get a list of the available algorithms supported by the version of OpenSSL on your platform, run the command: <b>openssl list-message-digest-algorithms</b>. Examples are <i>sha1</i>, <i>sha256</i>, and <i>sha512</i>.</p> <div data-bbox="815 645 922 902" style="border: 1px solid black; padding: 5px; width: fit-content;">  </div> <p><b>NOTE</b></p> <p><b>digest_algorithm</b> is not related to Glance's image signing and verification. It is only used to sign the universally unique identifier (UUID) as a part of the certificate file and key file validation.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● An OpenSSL message digest algorithm identifier</li> </ul> <p>Relation options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>

Configuration option = Default value	Type	Description
<b>disabled_notifications = []</b>	list value	<p>List of notifications to be disabled.</p> <p>Specify a list of notifications that should not be emitted. A notification can be given either as a notification type to disable a single event notification, or as a notification group prefix to disable all event notifications within a group.</p> <p>Possible values: A comma-separated list of individual notification types or notification groups to be disabled. Currently supported groups: <b>image</b> <b>image.member</b> <b>task</b> <b>metadef_namespace</b> <b>metadef_object</b> <b>metadef_property</b> <b>metadef_resource_type</b> <b>metadef_tag</b> For a complete listing and description of each event refer to:  <a href="http://docs.openstack.org/developer/glance/notifications.html">http://docs.openstack.org/developer/glance/notifications.html</a></p> <p>The values must be specified as:  &lt;group_name&gt;.&lt;event_name&gt;  For example:  image.create,task.success,metadef_tag</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<b>enable_image_import = True</b>	boolean value	<p>Enables the Image Import workflow introduced in Pike</p> <p>As <i>[DEFAULT]/node_staging_uri</i> is required for the Image Import, it's disabled per default in Pike, enabled per default in Queens and removed in Rocky. This allows Glance to operate with previous version configs upon upgrade.</p> <p>Setting this option to False will disable the endpoints related to Image Import Refactoring work.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <i>[DEFAULT]/node_staging_uri</i></li> </ul>

Configuration option = Default value	Type	Description
<b>enable_v1_api = True</b>	boolean value	<p>Deploy the v1 OpenStack Images API.</p> <p>When this option is set to <b>True</b>, Glance service will respond to requests on registered endpoints conforming to the v1 OpenStack Images API.</p> <p>NOTES:</p> <ul style="list-style-type: none"> <li>● If this option is enabled, then <b>enable_v1_registry</b> must also be set to <b>True</b> to enable mandatory usage of Registry service with v1 API.</li> <li>● If this option is disabled, then the <b>enable_v1_registry</b> option, which is enabled by default, is also recommended to be disabled.</li> <li>● This option is separate from <b>enable_v2_api</b>, both v1 and v2 OpenStack Images API can be deployed independent of each other.</li> <li>● If deploying only the v2 Images API, this option, which is enabled by default, should be disabled.</li> </ul> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● enable_v1_registry</li> <li>● enable_v2_api</li> </ul>

Configuration option = Default value	Type	Description
<b>enable_v1_registry = True</b>	boolean value	<p>Deploy the v1 API Registry service.</p> <p>When this option is set to <b>True</b>, the Registry service will be enabled in Glance for v1 API requests.</p> <p>NOTES:</p> <ul style="list-style-type: none"><li>● Use of Registry is mandatory in v1 API, so this option must be set to <b>True</b> if the <b>enable_v1_api</b> option is enabled.</li><li>● If deploying only the v2 OpenStack Images API, this option, which is enabled by default, should be disabled.</li></ul> <p>Possible values:</p> <ul style="list-style-type: none"><li>● True</li><li>● False</li></ul> <p>Related options:</p> <ul style="list-style-type: none"><li>● enable_v1_api</li></ul>




Configuration option = Default value	Type	Description
<b>enable_v2_api = True</b>	boolean value	<p>Deploy the v2 OpenStack Images API.</p> <p>When this option is set to <b>True</b>, Glance service will respond to requests on registered endpoints conforming to the v2 OpenStack Images API.</p> <p>NOTES:</p> <ul style="list-style-type: none"> <li>● If this option is disabled, then the <b>enable_v2_registry</b> option, which is enabled by default, is also recommended to be disabled.</li> <li>● This option is separate from <b>enable_v1_api</b>, both v1 and v2 OpenStack Images API can be deployed independent of each other.</li> <li>● If deploying only the v1 Images API, this option, which is enabled by default, should be disabled.</li> </ul> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● enable_v2_registry</li> <li>● enable_v1_api</li> </ul>

Configuration option = Default value	Type	Description
<b>enable_v2_registry = True</b>	boolean value	<p>Deploy the v2 API Registry service.</p> <p>When this option is set to <b>True</b>, the Registry service will be enabled in Glance for v2 API requests.</p> <p>NOTES:</p> <ul style="list-style-type: none"> <li>● Use of Registry is optional in v2 API, so this option must only be enabled if both <b>enable_v2_api</b> is set to <b>True</b> and the <b>data_api</b> option is set to <b>glance.db.registry.api</b>.</li> <li>● If deploying only the v1 OpenStack Images API, this option, which is enabled by default, should be disabled.</li> </ul> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● enable_v2_api</li> <li>● data_api</li> </ul>
<b>enabled_import_methods = [u'glance-direct', u'web-download']</b>	list value	<p>List of enabled Image Import Methods</p> <p>Both <i>glance-direct</i> and <i>web-download</i> are enabled by default.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● [DEFAULT]/node_staging_uri</li> <li>● [DEFAULT]/enable_image_import</li> </ul>
<b>executor_thread_pool_size = 64</b>	integer value	Size of executor thread pool when executor is threading or eventlet.
<b>fatal_deprecations = False</b>	boolean value	Enables or disables fatal status of deprecations.


Configuration option = Default value	Type	Description
<b>http_keepalive = True</b>	boolean value	<p>Set keep alive option for HTTP over TCP.</p> <p>Provide a boolean value to determine sending of keep alive packets. If set to <b>False</b>, the server returns the header "Connection: close". If set to <b>True</b>, the server returns a "Connection: Keep-Alive" in its responses. This enables retention of the same TCP connection for HTTP conversations instead of opening a new one with each new request.</p> <p>This option must be set to <b>False</b> if the client socket connection needs to be closed explicitly after the response is received and read successfully by the client.</p> <p>Possible values:</p> <ul style="list-style-type: none"><li>• True</li><li>• False</li></ul> <p>Related options:</p> <ul style="list-style-type: none"><li>• None</li></ul>

Configuration option = Default value	Type	Description
<b>image_cache_dir = None</b>	string value	<p>Base directory for image cache.</p> <p>This is the location where image data is cached and served out of. All cached images are stored directly under this directory. This directory also contains three subdirectories, namely, <b>incomplete</b>, <b>invalid</b> and <b>queue</b>.</p> <p>The <b>incomplete</b> subdirectory is the staging area for downloading images. An image is first downloaded to this directory. When the image download is successful it is moved to the base directory. However, if the download fails, the partially downloaded image file is moved to the <b>invalid</b> subdirectory.</p> <p>The <b>queue</b> subdirectory is used for queuing images for download. This is used primarily by the cache-prefetcher, which can be scheduled as a periodic task like cache-pruner and cache-cleaner, to cache images ahead of their usage. Upon receiving the request to cache an image, Glance touches a file in the <b>queue</b> directory with the image id as the file name. The cache-prefetcher, when running, polls for the files in <b>queue</b> directory and starts downloading them in the order they were created. When the download is successful, the zero-sized file is deleted from the <b>queue</b> directory. If the download fails, the zero-sized file remains and it'll be retried the next time cache-prefetcher runs.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A valid path</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <b>image_cache_sqlite_db</b></li> </ul>

Configuration option = Default value	Type	Description
<b>image_cache_driver = sqlite</b>	string value	<p>The driver to use for image cache management.</p> <p>This configuration option provides the flexibility to choose between the different image-cache drivers available. An image-cache driver is responsible for providing the essential functions of image-cache like write images to/read images from cache, track age and usage of cached images, provide a list of cached images, fetch size of the cache, queue images for caching and clean up the cache, etc.</p> <p>The essential functions of a driver are defined in the base class <b>glance.image_cache.drivers.base.Driver</b>. All image-cache drivers (existing and prospective) must implement this interface. Currently available drivers are <b>sqlite</b> and <b>xattr</b>. These drivers primarily differ in the way they store the information about cached images:</p> <ul style="list-style-type: none"> <li>● The <b>sqlite</b> driver uses a sqlite database (which sits on every glance node locally) to track the usage of cached images.</li> <li>● The <b>xattr</b> driver uses the extended attributes of files to store this information. It also requires a filesystem that sets <b>atime</b> on the files when accessed.</li> </ul> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● sqlite</li> <li>● xattr</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>

Configuration option = Default value	Type	Description
<b>image_cache_max_size = 10737418240</b>	integer value	<p>The upper limit on cache size, in bytes, after which the cache-pruner cleans up the image cache.</p> <div data-bbox="815 371 922 779" style="border: 1px solid black; padding: 5px; width: fit-content;">  </div> <p><b>NOTE</b></p> <p>This is just a threshold for cache-pruner to act upon. It is NOT a hard limit beyond which the image cache would never grow. In fact, depending on how often the cache-pruner runs and how quickly the cache fills, the image cache can far exceed the size specified here very easily. Hence, care must be taken to appropriately schedule the cache-pruner and in setting this limit.</p> <p>Glance caches an image when it is downloaded. Consequently, the size of the image cache grows over time as the number of downloads increases. To keep the cache size from becoming unmanageable, it is recommended to run the cache-pruner as a periodic task. When the cache pruner is kicked off, it compares the current size of image cache and triggers a cleanup if the image cache grew beyond the size specified here. After the cleanup, the size of cache is less than or equal to size specified here.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any non-negative integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>

Configuration option = Default value	Type	Description
<b>image_cache_sqlite_db = cache.db</b>	string value	<p>The relative path to sqlite file database that will be used for image cache management.</p> <p>This is a relative path to the sqlite file database that tracks the age and usage statistics of image cache. The path is relative to image cache base directory, specified by the configuration option <b>image_cache_dir</b>.</p> <p>This is a lightweight database with just one table.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A valid relative path to sqlite file database</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <b>image_cache_dir</b></li> </ul>
<b>image_cache_stall_time = 86400</b>	integer value	<p>The amount of time, in seconds, an incomplete image remains in the cache.</p> <p>Incomplete images are images for which download is in progress. Please see the description of configuration option <b>image_cache_dir</b> for more detail. Sometimes, due to various reasons, it is possible the download may hang and the incompletely downloaded image remains in the <b>incomplete</b> directory. This configuration option sets a time limit on how long the incomplete images should remain in the <b>incomplete</b> directory before they are cleaned up. Once an incomplete image spends more time than is specified here, it'll be removed by cache-cleaner on its next run.</p> <p>It is recommended to run cache-cleaner as a periodic task on the Glance API nodes to keep the incomplete images from occupying disk space.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Any non-negative integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>

Configuration option = Default value	Type	Description
<b>image_location_quota = 10</b>	integer value	<p>Maximum number of locations allowed on an image.</p> <p>Any negative value is interpreted as unlimited.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<b>image_member_quota = 128</b>	integer value	<p>Maximum number of image members per image.</p> <p>This limits the maximum of users an image can be shared with. Any negative value is interpreted as unlimited.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<b>image_property_quota = 128</b>	integer value	<p>Maximum number of properties allowed on an image.</p> <p>This enforces an upper limit on the number of additional properties an image can have. Any negative value is interpreted as unlimited.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p>This won't have any impact if additional properties are disabled. Please refer to <b>allow_additional_image_properties</b>.</p> </div> </div> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <b>allow_additional_image_properties</b></li> </ul>





Configuration option = Default value	Type	Description
<b>image_size_cap = 1099511627776</b>	integer value	<p>Maximum size of image a user can upload in bytes.</p> <p>An image upload greater than the size mentioned here would result in an image creation failure. This configuration option defaults to 1099511627776 bytes (1 TiB).</p> <p>NOTES:</p> <ul style="list-style-type: none"> <li>• This value should only be increased after careful consideration and must be set less than or equal to 8 EiB (9223372036854775808).</li> <li>• This value must be set with careful consideration of the backend storage capacity. Setting this to a very low value may result in a large number of image failures. And, setting this to a very large value may result in faster consumption of storage. Hence, this must be set according to the nature of images created and storage capacity available.</li> </ul> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Any positive number less than or equal to 9223372036854775808</li> </ul>
<b>image_tag_quota = 128</b>	integer value	<p>Maximum number of tags allowed on an image.</p> <p>Any negative value is interpreted as unlimited.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<code>`instance_format = [instance: %(uuid)s] `</code>	string value	The format for an instance that is passed with the log message.
<code>`instance_uuid_format = [instance: %(uuid)s] `</code>	string value	The format for an instance UUID that is passed with the log message.




Configuration option = Default value	Type	Description
<b>key_file = None</b>	string value	<p>Absolute path to a private key file.</p> <p>Provide a string value representing a valid absolute path to a private key file which is required to establish the client-server connection.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Absolute path to the private key file</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>limit_param_default = 25</b>	integer value	<p>The default number of results to return for a request.</p> <p>Responses to certain API requests, like list images, may return multiple items. The number of results returned can be explicitly controlled by specifying the <b>limit</b> parameter in the API request. However, if a <b>limit</b> parameter is not specified, this configuration value will be used as the default number of results to be returned for any API request.</p> <p>NOTES:</p> <ul style="list-style-type: none"> <li>● The value of this configuration option may not be greater than the value specified by <b>api_limit_max</b>.</li> <li>● Setting this to a very large value may slow down database queries and increase response times. Setting this to a very low value may result in poor user experience.</li> </ul> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any positive integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>api_limit_max</code></li> </ul>

Configuration option = Default value	Type	Description
<b>location_strategy = location_order</b>	string value	<p>Strategy to determine the preference order of image locations.</p> <p>This configuration option indicates the strategy to determine the order in which an image's locations must be accessed to serve the image's data. Glance then retrieves the image data from the first responsive active location it finds in this list.</p> <p>This option takes one of two possible values <b>location_order</b> and <b>store_type</b>. The default value is <b>location_order</b>, which suggests that image data be served by using locations in the order they are stored in Glance. The <b>store_type</b> value sets the image location preference based on the order in which the storage backends are listed as a comma separated list for the configuration option <b>store_type_preference</b>.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● location_order</li> <li>● store_type</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● store_type_preference</li> </ul>
<b>log-config-append = None</b>	string value	The name of a logging configuration file. This file is appended to any existing logging configuration files. For details about logging configuration files, see the Python logging module documentation. Note that when logging configuration files are used then all logging configuration is set in the configuration file and other logging configuration options are ignored (for example, logging_context_format_string).
<b>log-date-format = %Y-%m-%d %H:%M:%S</b>	string value	Defines the format string for <code>%(asctime)s</code> in log records. Default: <code>%(default)s</code> . This option is ignored if <code>log_config_append</code> is set.
<b>log-dir = None</b>	string value	(Optional) The base directory used for relative <code>log_file</code> paths. This option is ignored if <code>log_config_append</code> is set.
<b>log-file = None</b>	string value	(Optional) Name of log file to send logging output to. If no default is set, logging will go to <code>stderr</code> as defined by <code>use_stderr</code> . This option is ignored if <code>log_config_append</code> is set.

Configuration option = Default value	Type	Description
<b>logging_context_format_string</b> = <code>%(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [%(request_id)s %(user_identity)s] %(instance)s%(message)s</code>	string value	Format string to use for log messages with context.
<b>logging_debug_format_s uffix</b> = <code>%(funcName)s %(pathname)s:%(lineno)d</code>	string value	Additional data to append to log message when logging level for the message is DEBUG.
<b>logging_default_format_s tring</b> = <code>%(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [-] %(instance)s%(message)s</code>	string value	Format string to use for log messages when context is undefined.
<b>logging_exception_prefix</b> = <code>%(asctime)s.%(msecs)03d %(process)d ERROR %(name)s %(instance)s</code>	string value	Prefix each line of exception output with this format.
<b>logging_user_identity_for mat</b> = <code>%(user)s %(tenant)s %(domain)s %(user_domain)s %(project_domain)s</code>	string value	Defines the format string for <code>%(user_identity)s</code> that is used in <code>logging_context_format_string</code> .

Configuration option = Default value	Type	Description
<b>max_header_line = 16384</b>	integer value	<p>Maximum line size of message headers.</p> <p>Provide an integer value representing a length to limit the size of message headers. The default value is 16384.</p> <div data-bbox="815 472 922 790" style="display: inline-block; vertical-align: top;">  </div> <p><b>NOTE</b></p> <p><b>max_header_line</b> may need to be increased when using large tokens (typically those generated by the Keystone v3 API with big service catalogs). However, it is to be kept in mind that larger values for <b>max_header_line</b> would flood the logs.</p> <p>Setting <b>max_header_line</b> to 0 sets no limit for the line size of message headers.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● 0</li> <li>● Positive integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>max_request_id_length = 64</b>	integer value	<p>Limit the request ID length.</p> <p>Provide an integer value to limit the length of the request ID to the specified length. The default value is 64. Users can change this to any integer value between 0 and 16384 however keeping in mind that a larger value may flood the logs.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Integer value between 0 and 16384</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>

Configuration option = Default value	Type	Description
<b>metadata_encryption_key</b> = None	string value	<p>AES key for encrypting store location metadata.</p> <p>Provide a string value representing the AES cipher to use for encrypting Glance store metadata.</p> <p> <b>NOTE</b></p> <p>The AES key to use must be set to a random string of length 16, 24 or 32 bytes.</p> <p>Possible values:</p> <ul style="list-style-type: none"><li>• String value representing a valid AES key</li></ul> <p>Related options:</p> <ul style="list-style-type: none"><li>• None</li></ul>

Configuration option = Default value	Type	Description
<b>node_staging_uri = file:///tmp/staging/</b>	string value	<p>The URL provides location where the temporary data will be stored</p> <p>This option is for Glance internal use only. Glance will save the image data uploaded by the user to <i>staging</i> endpoint during the image import process.</p> <p>This option does not change the <i>staging</i> API endpoint by any means.</p> <p> <b>NOTE</b></p> <p>It is discouraged to use same path as [task]/work_dir</p> <p> <b>NOTE</b></p> <p><i>file://&lt;absolute-directory-path&gt;</i> is the only option api_image_import flow will support for now.</p> <p> <b>NOTE</b></p> <p>The staging path must be on shared filesystem available to all Glance API nodes.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● String starting with <i>file://</i> followed by absolute FS path</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● [task]/work_dir</li> <li>● [DEFAULT]/enable_image_import (<b>deprecated</b>)</li> </ul>

Configuration option = Default value	Type	Description
<b>owner_is_tenant = True</b>	boolean value	<p>Set the image owner to tenant or the authenticated user.</p> <p>Assign a boolean value to determine the owner of an image. When set to True, the owner of the image is the tenant. When set to False, the owner of the image will be the authenticated user issuing the request. Setting it to False makes the image private to the associated user and sharing with other users within the same tenant (or "project") requires explicit image sharing via image membership.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• True</li> <li>• False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<b>property_protection_file = None</b>	string value	<p>The location of the property protection file.</p> <p>Provide a valid path to the property protection file which contains the rules for property protections and the roles/policies associated with them.</p> <p>A property protection file, when set, restricts the Glance image properties to be created, read, updated and/or deleted by a specific set of users that are identified by either roles or policies. If this configuration option is not set, by default, property protections won't be enforced. If a value is specified and the file is not found, the glance-api service will fail to start. More information on property protections can be found at: <a href="https://docs.openstack.org/glance/latest/admin/property-protections.html">https://docs.openstack.org/glance/latest/admin/property-protections.html</a></p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Empty string</li> <li>• Valid path to the property protection configuration file</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <code>property_protection_rule_format</code></li> </ul>




Configuration option = Default value	Type	Description
<b>property_protection_rule_format = roles</b>	string value	<p>Rule format for property protection.</p> <p>Provide the desired way to set property protection on Glance image properties. The two permissible values are <b>roles</b> and <b>policies</b>. The default value is <b>roles</b>.</p> <p>If the value is <b>roles</b>, the property protection file must contain a comma separated list of user roles indicating permissions for each of the CRUD operations on each property being protected. If set to <b>policies</b>, a policy defined in policy.json is used to express property protections for each of the CRUD operations. Examples of how property protections are enforced based on <b>roles</b> or <b>policies</b> can be found at:  <a href="https://docs.openstack.org/glance/latest/admin/property-protections.html#examples">https://docs.openstack.org/glance/latest/admin/property-protections.html#examples</a></p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● roles</li> <li>● policies</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● property_protection_file</li> </ul>
<b>public_endpoint = None</b>	string value	<p>Public url endpoint to use for Glance versions response.</p> <p>This is the public url endpoint that will appear in the Glance "versions" response. If no value is specified, the endpoint that is displayed in the version's response is that of the host running the API service. Change the endpoint to represent the proxy URL if the API service is running behind a proxy. If the service is running behind a load balancer, add the load balancer's URL for this value.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● None</li> <li>● Proxy URL</li> <li>● Load balancer URL</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>

Configuration option = Default value	Type	Description
<b>publish_errors = False</b>	boolean value	Enables or disables publication of error events.
<b>pydev_worker_debug_host = None</b>	unknown value	<p>Host address of the pydev server.</p> <p>Provide a string value representing the hostname or IP of the pydev server to use for debugging. The pydev server listens for debug connections on this address, facilitating remote debugging in Glance.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Valid hostname</li> <li>Valid IP address</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>None</li> </ul>
<b>pydev_worker_debug_port = 5678</b>	port value	<p>Port number that the pydev server will listen on.</p> <p>Provide a port number to bind the pydev server to. The pydev process accepts debug connections on this port and facilitates remote debugging in Glance.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>A valid port number</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>None</li> </ul>
<b>rate_limit_burst = 0</b>	integer value	Maximum number of logged messages per <code>rate_limit_interval</code> .
<b>rate_limit_except_level = CRITICAL</b>	string value	Log level name used by rate limiting: CRITICAL, ERROR, INFO, WARNING, DEBUG or empty string. Logs with level greater or equal to <code>rate_limit_except_level</code> are not filtered. An empty string means that all levels are filtered.
<b>rate_limit_interval = 0</b>	integer value	Interval, number of seconds, of log rate limiting.

Configuration option = Default value	Type	Description
<b>registry_client_ca_file = None</b>	string value	<p>Absolute path to the Certificate Authority file.</p> <p>Provide a string value representing a valid absolute path to the certificate authority file to use for establishing a secure connection to the registry server.</p> <div data-bbox="815 510 922 857" style="background-image: linear-gradient(to right, transparent 49%, #ccc 49% 51%, #ccc 51% 53%, transparent 53%); background-size: 10px 10px; border: 1px solid #ccc; width: 67px; height: 155px; margin-bottom: 10px;"></div> <p><b>NOTE</b></p> <p>This option must be set if <b>registry_client_protocol</b> is set to <b>https</b>. Alternatively, the <code>GLANCE_CLIENT_CA_FILE</code> environment variable may be set to a filepath of the CA file. This option is ignored if the <b>registry_client_insecure</b> option is set to <b>True</b>.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● String value representing a valid absolute path to the CA file.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● registry_client_protocol</li> <li>● registry_client_insecure</li> </ul>
<b>registry_client_cert_file = None</b>	string value	<p>Absolute path to the certificate file.</p> <p>Provide a string value representing a valid absolute path to the certificate file to use for establishing a secure connection to the registry server.</p> <div data-bbox="815 1525 922 1783" style="background-image: linear-gradient(to right, transparent 49%, #ccc 49% 51%, #ccc 51% 53%, transparent 53%); background-size: 10px 10px; border: 1px solid #ccc; width: 67px; height: 115px; margin-bottom: 10px;"></div> <p><b>NOTE</b></p> <p>This option must be set if <b>registry_client_protocol</b> is set to <b>https</b>. Alternatively, the <code>GLANCE_CLIENT_CERT_FILE</code> environment variable may be set to a filepath of the certificate file.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● String value representing a valid absolute path to the certificate file.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● registry_client_protocol</li> </ul>

Configuration option = Default value	Type	Description
<b>registry_client_insecure = False</b>	boolean value	<p>Set verification of the registry server certificate.</p> <p>Provide a boolean value to determine whether or not to validate SSL connections to the registry server. By default, this option is set to <b>False</b> and the SSL connections are validated.</p> <p>If set to <b>True</b>, the connection to the registry server is not validated via a certifying authority and the <b>registry_client_ca_file</b> option is ignored. This is the registry's equivalent of specifying <code>--insecure</code> on the command line using <code>glanceclient</code> for the API.</p> <p>Possible values:</p> <ul style="list-style-type: none"><li>• True</li><li>• False</li></ul> <p>Related options:</p> <ul style="list-style-type: none"><li>• <code>registry_client_protocol</code></li><li>• <code>registry_client_ca_file</code></li></ul>

Configuration option = Default value	Type	Description
<b>registry_client_key_file = None</b>	string value	<p>Absolute path to the private key file.</p> <p>Provide a string value representing a valid absolute path to the private key file to use for establishing a secure connection to the registry server.</p> <div data-bbox="815 472 922 728" style="border: 1px solid black; padding: 5px; width: fit-content;">  </div> <p><b>NOTE</b></p> <p>This option must be set if <b>registry_client_protocol</b> is set to <b>https</b>. Alternatively, the <code>GLANCE_CLIENT_KEY_FILE</code> environment variable may be set to a filepath of the key file.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● String value representing a valid absolute path to the key file.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>registry_client_protocol</code></li> </ul>

Configuration option = Default value	Type	Description
<b>registry_client_protocol = http</b>	string value	<p>Protocol to use for communication with the registry server.</p> <p>Provide a string value representing the protocol to use for communication with the registry server. By default, this option is set to <b>http</b> and the connection is not secure.</p> <p>This option can be set to <b>https</b> to establish a secure connection to the registry server. In this case, provide a key to use for the SSL connection using the <b>registry_client_key_file</b> option. Also include the CA file and cert file using the options <b>registry_client_ca_file</b> and <b>registry_client_cert_file</b> respectively.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• http</li> <li>• https</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• registry_client_key_file</li> <li>• registry_client_cert_file</li> <li>• registry_client_ca_file</li> </ul>
<b>registry_client_timeout = 600</b>	integer value	<p>Timeout value for registry requests.</p> <p>Provide an integer value representing the period of time in seconds that the API server will wait for a registry request to complete. The default value is 600 seconds.</p> <p>A value of 0 implies that a request will never timeout.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Zero</li> <li>• Positive integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>

Configuration option = Default value	Type	Description
<b>registry_host = 0.0.0.0</b>	unknown value	Address the registry server is hosted on.  Possible values: <ul style="list-style-type: none"> <li>• A valid IP or hostname</li> </ul> Related options: <ul style="list-style-type: none"> <li>• None</li> </ul>
<b>registry_port = 9191</b>	port value	Port the registry server is listening on.  Possible values: <ul style="list-style-type: none"> <li>• A valid port number</li> </ul> Related options: <ul style="list-style-type: none"> <li>• None</li> </ul>
<b>rpc_ack_timeout_base = 15</b>	integer value	Number of seconds to wait for an ack from a cast/call. After each retry attempt this timeout is multiplied by some specified multiplier.
<b>rpc_ack_timeout_multiplier = 2</b>	integer value	Number to multiply base ack timeout by after each retry attempt.
<b>rpc_backend = rabbit</b>	string value	The messaging driver to use, defaults to rabbit. Other drivers include amqp and zmq.
<b>rpc_conn_pool_size = 30</b>	integer value	Size of RPC connection pool.
<b>rpc_message_ttl = 300</b>	integer value	Expiration timeout in seconds of a sent/received message after which it is not tracked anymore by a client/server.
<b>rpc_poll_timeout = 1</b>	integer value	The default number of seconds that poll should wait. Poll raises timeout exception when timeout expired.
<b>rpc_response_timeout = 60</b>	integer value	Seconds to wait for a response from a call.
<b>rpc_retry_attempts = 3</b>	integer value	Default number of message sending attempts in case of any problems occurred: positive value N means at most N retries, 0 means no retries, None or -1 (or any other negative values) mean to retry forever. This option is used only if acknowledgments are enabled.

Configuration option = Default value	Type	Description
<b>rpc_thread_pool_size = 100</b>	integer value	Maximum number of (green) threads to work concurrently.
<b>rpc_use_acks = False</b>	boolean value	Wait for message acknowledgements from receivers. This mechanism works only via proxy without PUB/SUB.
<b>rpc_zmq_bind_address = *</b>	string value	ZeroMQ bind address. Should be a wildcard (*), an ethernet interface, or IP. The "host" option should point or resolve to this address.
<b>rpc_zmq_bind_port_retries = 100</b>	integer value	Number of retries to find free port number before fail with ZMQBindError.
<b>rpc_zmq_contexts = 1</b>	integer value	Number of ZeroMQ contexts, defaults to 1.
<b>rpc_zmq_host = ccosp-undercloud.usersys.redhat.com</b>	string value	Name of this node. Must be a valid hostname, FQDN, or IP address. Must match "host" option, if running Nova.
<b>rpc_zmq_ipc_dir = /var/run/openstack</b>	string value	Directory for holding IPC sockets.
<b>rpc_zmq_matchmaker = redis</b>	string value	MatchMaker driver.
<b>rpc_zmq_max_port = 65536</b>	integer value	Maximal port number for random ports range.
<b>rpc_zmq_min_port = 49153</b>	port value	Minimal port number for random ports range.
<b>rpc_zmq_serialization = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>rpc_zmq_topic_backlog = None</b>	integer value	Maximum number of ingress messages to locally buffer per topic. Default is unlimited.



Configuration option = Default value	Type	Description
<b>scrub_pool_size = 1</b>	integer value	<p>The size of thread pool to be used for scrubbing images.</p> <p>When there are a large number of images to scrub, it is beneficial to scrub images in parallel so that the scrub queue stays in control and the backend storage is reclaimed in a timely fashion. This configuration option denotes the maximum number of images to be scrubbed in parallel. The default value is one, which signifies serial scrubbing. Any value above one indicates parallel scrubbing.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Any non-zero positive integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <b>delayed_delete</b></li> </ul>
<b>scrub_time = 0</b>	integer value	<p>The amount of time, in seconds, to delay image scrubbing.</p> <p>When delayed delete is turned on, an image is put into <b>pending_delete</b> state upon deletion until the scrubber deletes its image data. Typically, soon after the image is put into <b>pending_delete</b> state, it is available for scrubbing. However, scrubbing can be delayed until a later point using this configuration option. This option denotes the time period an image spends in <b>pending_delete</b> state before it is available for scrubbing.</p> <p>It is important to realize that this has storage implications. The larger the <b>scrub_time</b>, the longer the time to reclaim backend storage from deleted images.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Any non-negative integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <b>delayed_delete</b></li> </ul>
<b>secure_proxy_ssl_header = None</b>	string value	<p>The HTTP header used to determine the scheme for the original request, even if it was removed by an SSL terminating proxy. Typical value is "HTTP_X_FORWARDED_PROTO".</p>


Configuration option = Default value	Type	Description
<b>send_identity_headers = False</b>	boolean value	<p>Send headers received from identity when making requests to registry.</p> <p>Typically, Glance registry can be deployed in multiple flavors, which may or may not include authentication. For example, <b>trusted-auth</b> is a flavor that does not require the registry service to authenticate the requests it receives. However, the registry service may still need a user context to be populated to serve the requests. This can be achieved by the caller (the Glance API usually) passing through the headers it received from authenticating with identity for the same request. The typical headers sent are <b>X-User-Id, X-Tenant-Id, X-Roles, X-Identity-Status</b> and <b>X-Service-Catalog</b>.</p> <p>Provide a boolean value to determine whether to send the identity headers to provide tenant and user information along with the requests to registry service. By default, this option is set to <b>False</b>, which means that user and tenant information is not available readily. It must be obtained by authenticating. Hence, if this is set to <b>False, flavor</b> must be set to value that either includes authentication or authenticated user context.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● flavor</li> </ul>

Configuration option = Default value	Type	Description
<b>show_image_direct_url = False</b>	boolean value	<p>Show direct image location when returning an image.</p> <p>This configuration option indicates whether to show the direct image location when returning image details to the user. The direct image location is where the image data is stored in backend storage. This image location is shown under the image property <b>direct_url</b>.</p> <p>When multiple image locations exist for an image, the best location is displayed based on the location strategy indicated by the configuration option <b>location_strategy</b>.</p> <p>NOTES:</p> <ul style="list-style-type: none"> <li>● Revealing image locations can present a GRAVE SECURITY RISK as image locations can sometimes include credentials. Hence, this is set to <b>False</b> by default. Set this to <b>True</b> with EXTREME CAUTION and ONLY IF you know what you are doing!</li> <li>● If an operator wishes to avoid showing any image location(s) to the user, then both this option and <b>show_multiple_locations</b> MUST be set to <b>False</b>.</li> </ul> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● show_multiple_locations</li> <li>● location_strategy</li> </ul>

Configuration option = Default value	Type	Description
<b>show_multiple_locations</b> = <b>False</b>	boolean value	<p>Show all image locations when returning an image.</p> <p>This configuration option indicates whether to show all the image locations when returning image details to the user. When multiple image locations exist for an image, the locations are ordered based on the location strategy indicated by the configuration opt <b>location_strategy</b>. The image locations are shown under the image property <b>locations</b>.</p> <p>NOTES:</p> <ul style="list-style-type: none"> <li>● Revealing image locations can present a GRAVE SECURITY RISK as image locations can sometimes include credentials. Hence, this is set to <b>False</b> by default. Set this to <b>True</b> with EXTREME CAUTION and ONLY IF you know what you are doing!</li> <li>● See <a href="https://wiki.openstack.org/wiki/OSSN/OSSN-0065">https://wiki.openstack.org/wiki/OSSN/OSSN-0065</a> for more information.</li> <li>● If an operator wishes to avoid showing any image location(s) to the user, then both this option and <b>show_image_direct_url</b> MUST be set to <b>False</b>.</li> </ul> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● show_image_direct_url</li> <li>● location_strategy</li> </ul>
<b>subscribe_on</b> = []	list value	List of publisher hosts SubConsumer can subscribe on. This option has higher priority then the default publishers list taken from the matchmaker.
<b>syslog-log-facility</b> = <b>LOG_USER</b>	string value	Syslog facility to receive log lines. This option is ignored if log_config_append is set.

Configuration option = Default value	Type	Description
<b>tcp_keepidle = 600</b>	integer value	<p>Set the wait time before a connection recheck.</p> <p>Provide a positive integer value representing time in seconds which is set as the idle wait time before a TCP keep alive packet can be sent to the host. The default value is 600 seconds.</p> <p>Setting <b>tcp_keepidle</b> helps verify at regular intervals that a connection is intact and prevents frequent TCP connection reestablishment.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Positive integer value representing time in seconds</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>transport_url = None</b>	string value	<p>The network address and optional user credentials for connecting to the messaging backend, in URL format. The expected format is:</p> <pre>driver://[user:pass@]host:port[, [userN:passN@]hostN:portN]/virtual_host?query</pre> <p>Example: rabbit://rabbitmq:password@127.0.0.1:5672//</p> <p>For full details on the fields in the URL see the documentation of <code>oslo_messaging.TransportURL</code> at <a href="https://docs.openstack.org/oslo.messaging/latest/reference/transport.html">https://docs.openstack.org/oslo.messaging/latest/reference/transport.html</a></p>
<b>use-journal = False</b>	boolean value	<p>Enable journald for logging. If running in a systemd environment you may wish to enable journal support. Doing so will use the journal native protocol which includes structured metadata in addition to log messages. This option is ignored if <code>log_config_append</code> is set.</p>
<b>use-json = False</b>	boolean value	<p>Use JSON formatting for logging. This option is ignored if <code>log_config_append</code> is set.</p>
<b>use-syslog = False</b>	boolean value	<p>Use syslog for logging. Existing syslog format is DEPRECATED and will be changed later to honor RFC5424. This option is ignored if <code>log_config_append</code> is set.</p>

Configuration option = Default value	Type	Description
<b>use_dynamic_connections = False</b>	boolean value	This option makes direct connections dynamic or static. It makes sense only with <code>use_router_proxy=False</code> which means to use direct connections for direct message types (ignored otherwise).
<b>use_pub_sub = False</b>	boolean value	Use PUB/SUB pattern for fanout methods. PUB/SUB always uses proxy.
<b>use_router_proxy = False</b>	boolean value	Use ROUTER remote proxy.
<b>use_stderr = False</b>	boolean value	Log output to standard error. This option is ignored if <code>log_config_append</code> is set.
<b>use_user_token = True</b>	boolean value	Whether to pass through the user token when making requests to the registry. To prevent failures with token expiration during big files upload, it is recommended to set this parameter to <code>False</code> . If "use_user_token" is not in effect, then admin credentials can be specified.
<b>user_storage_quota = 0</b>	string value	<p>Maximum amount of image storage per tenant.</p> <p>This enforces an upper limit on the cumulative storage consumed by all images of a tenant across all stores. This is a per-tenant limit.</p> <p>The default unit for this configuration option is Bytes. However, storage units can be specified using case-sensitive literals <b>B</b>, <b>KB</b>, <b>MB</b>, <b>GB</b> and <b>TB</b> representing Bytes, KiloBytes, MegaBytes, GigaBytes and TeraBytes respectively. Note that there should not be any space between the value and unit. Value <b>0</b> signifies no quota enforcement. Negative values are invalid and result in errors.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A string that is a valid concatenation of a non-negative integer representing the storage value and an optional string literal representing storage units as mentioned above.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>

Configuration option = Default value	Type	Description
<b>watch-log-file = False</b>	boolean value	Uses logging handler designed to watch file system. When log file is moved or removed this handler will open a new log file with specified path instantaneously. It makes sense only if log_file option is specified and Linux platform is used. This option is ignored if log_config_append is set.
<b>workers = None</b>	integer value	<p>Number of Glance worker processes to start.</p> <p>Provide a non-negative integer value to set the number of child process workers to service requests. By default, the number of CPUs available is set as the value for <b>workers</b> limited to 8. For example if the processor count is 6, 6 workers will be used, if the processor count is 24 only 8 workers will be used. The limit will only apply to the default value, if 24 workers is configured, 24 is used.</p> <p>Each worker process is made to listen on the port set in the configuration file and contains a greenthread pool of size 1000.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p>Setting the number of workers to zero, triggers the creation of a single API process with a greenthread pool of size 1000.</p> </div> </div> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● 0</li> <li>● Positive integer value (typically equal to the number of CPUs)</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>zmq_failover_connections = 2</b>	integer value	How many additional connections to a host will be made for failover reasons. This option is actual only in dynamic connections mode.
<b>zmq_immediate = True</b>	boolean value	This option configures round-robin mode in zmq socket. True means not keeping a queue when server side disconnects. False means to keep queue and messages even if server is disconnected, when the server appears we send all accumulated messages to it.

Configuration option = Default value	Type	Description
<b>zmq_linger = -1</b>	integer value	Number of seconds to wait before all pending messages will be sent after closing a socket. The default value of -1 specifies an infinite linger period. The value of 0 specifies no linger period. Pending messages shall be discarded immediately when the socket is closed. Positive values specify an upper bound for the linger period.
<b>zmq_target_expire = 300</b>	integer value	Expiration timeout in seconds of a name service record about existing target (< 0 means no timeout).
<b>zmq_target_update = 180</b>	integer value	Update period in seconds of a name service record about existing target.
<b>zmq_tcp_keepalive = -1</b>	integer value	Enable/disable TCP keepalive (KA) mechanism. The default value of -1 (or any other negative value) means to skip any overrides and leave it to OS default; 0 and 1 (or any other positive value) mean to disable and enable the option respectively.
<b>zmq_tcp_keepalive_cnt = -1</b>	integer value	The number of retransmissions to be carried out before declaring that remote end is not available. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_idle = -1</b>	integer value	The duration between two keepalive transmissions in idle condition. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_intvl = -1</b>	integer value	The duration between two successive keepalive retransmissions, if acknowledgement to the previous keepalive transmission is not received. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.

### 5.1.2. cors

The following table outlines the options available under the **[cors]** group in the **/etc/glance/glance-api.conf** file.

Table 5.2. cors



Configuration option = Default value	Type	Description
<b>allow_credentials = True</b>	boolean value	Indicate that the actual request can include user credentials
<b>allow_headers =</b> [u'Content-MD5', u'X-Image-Meta-Checksum', u'X-Storage-Token', u'Accept-Encoding', u'X-Auth-Token', u'X-Identity-Status', u'X-Roles', u'X-Service-Catalog', u'X-User-Id', u'X-Tenant-Id', u'X-OpenStack-Request-ID']	list value	Indicate which header field names may be used during the actual request.
<b>allow_methods =</b> [u'GET', u'PUT', u'POST', u'DELETE', u'PATCH']	list value	Indicate which methods can be used during the actual request.
<b>allowed_origin = None</b>	list value	Indicate whether this resource may be shared with the domain received in the requests "origin" header. Format: "<protocol>://<host>[:<port>]", no trailing slash. Example: <a href="https://horizon.example.com">https://horizon.example.com</a>
<b>expose_headers =</b> [u'X-Image-Meta-Checksum', u'X-Auth-Token', u'X-Subject-Token', u'X-Service-Token', u'X-OpenStack-Request-ID']	list value	Indicate which headers are safe to expose to the API. Defaults to HTTP Simple Headers.
<b>max_age = 3600</b>	integer value	Maximum cache age of CORS preflight requests.

### 5.1.3. database

The following table outlines the options available under the **[database]** group in the `/etc/glance/glance-api.conf` file.

Table 5.3. database

Configuration option = Default value	Type	Description
<b>backend = sqlalchemy</b>	string value	The back end to use for the database.
<b>connection = None</b>	string value	The SQLAlchemy connection string to use to connect to the database.

Configuration option = Default value	Type	Description
<b>connection_debug = 0</b>	integer value	Verbosity of SQL debugging information: 0=None, 100=Everything.
<b>connection_recycle_time = 3600</b>	integer value	Connections which have been present in the connection pool longer than this number of seconds will be replaced with a new one the next time they are checked out from the pool.
<b>connection_trace = False</b>	boolean value	Add Python stack traces to SQL as comment strings.
<b>db_inc_retry_interval = True</b>	boolean value	If True, increases the interval between retries of a database operation up to db_max_retry_interval.
<b>db_max_retries = 20</b>	integer value	Maximum retries in case of connection error or deadlock error before error is raised. Set to -1 to specify an infinite retry count.
<b>db_max_retry_interval = 10</b>	integer value	If db_inc_retry_interval is set, the maximum seconds between retries of a database operation.
<b>db_retry_interval = 1</b>	integer value	Seconds between retries of a database transaction.
<b>max_overflow = 50</b>	integer value	If set, use this value for max_overflow with SQLAlchemy.
<b>max_pool_size = 5</b>	integer value	Maximum number of SQL connections to keep open in a pool. Setting a value of 0 indicates no limit.
<b>max_retries = 10</b>	integer value	Maximum number of database connection retries during startup. Set to -1 to specify an infinite retry count.
<b>min_pool_size = 1</b>	integer value	Minimum number of SQL connections to keep open in a pool.
<b>mysql_enable_ndb = False</b>	boolean value	If True, transparently enables support for handling MySQL Cluster (NDB).
<b>mysql_sql_mode = TRADITIONAL</b>	string value	The SQL mode to be used for MySQL sessions. This option, including the default, overrides any server-set SQL mode. To use whatever SQL mode is set by the server configuration, set this to no value. Example: mysql_sql_mode=
<b>pool_timeout = None</b>	integer value	If set, use this value for pool_timeout with SQLAlchemy.

Configuration option = Default value	Type	Description
<b>retry_interval = 10</b>	integer value	Interval between retries of opening a SQL connection.
<b>slave_connection = None</b>	string value	The SQLAlchemy connection string to use to connect to the slave database.
<b>sqlite_synchronous = True</b>	boolean value	If True, SQLite uses synchronous mode.
<b>use_db_reconnect = False</b>	boolean value	Enable the experimental use of database reconnect on connection lost.
<b>use_tpool = False</b>	boolean value	Enable the experimental use of thread pooling for all DB API calls

#### 5.1.4. glance\_store

The following table outlines the options available under the **[glance\_store]** group in the `/etc/glance/glance-api.conf` file.

Table 5.4. glance\_store

Configuration option = Default value	Type	Description
<b>cinder_api_insecure = False</b>	boolean value	<p>Allow to perform insecure SSL requests to cinder.</p> <p>If this option is set to True, HTTPS endpoint connection is verified using the CA certificates file specified by <b>cinder_ca_certificates_file</b> option.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● cinder_ca_certificates_file</li> </ul>

Configuration option = Default value	Type	Description
<b>cinder_ca_certificates_file = None</b>	string value	<p>Location of a CA certificates file used for cinder client requests.</p> <p>The specified CA certificates file, if set, is used to verify cinder connections via HTTPS endpoint. If the endpoint is HTTP, this value is ignored.</p> <p><b>cinder_api_insecure</b> must be set to <b>True</b> to enable the verification.</p> <p>Possible values:</p> <ul style="list-style-type: none"><li>● Path to a ca certificates file</li></ul> <p>Related options:</p> <ul style="list-style-type: none"><li>● cinder_api_insecure</li></ul>

Configuration option = Default value	Type	Description
<b>cinder_catalog_info = volume2::publicURL</b>	string value	<p>Information to match when looking for cinder in the service catalog.</p> <p>When the <b>cinder_endpoint_template</b> is not set and any of <b>cinder_store_auth_address</b>, <b>cinder_store_user_name</b>, <b>cinder_store_project_name</b>, <b>cinder_store_password</b> is not set, cinder store uses this information to lookup cinder endpoint from the service catalog in the current context. <b>cinder_os_region_name</b>, if set, is taken into consideration to fetch the appropriate endpoint.</p> <p>The service catalog can be listed by the <b>openstack catalog list</b> command.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A string of of the following form: <b>&lt;service_type&gt;:&lt;service_name&gt;: &lt;interface&gt;</b> At least <b>service_type</b> and <b>interface</b> should be specified. <b>service_name</b> can be omitted.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● cinder_os_region_name</li> <li>● cinder_endpoint_template</li> <li>● cinder_store_auth_address</li> <li>● cinder_store_user_name</li> <li>● cinder_store_project_name</li> <li>● cinder_store_password</li> </ul>

Configuration option = Default value	Type	Description
<b>cinder_endpoint_template = None</b>	string value	<p>Override service catalog lookup with template for cinder endpoint.</p> <p>When this option is set, this value is used to generate cinder endpoint, instead of looking up from the service catalog. This value is ignored if <b>cinder_store_auth_address</b>, <b>cinder_store_user_name</b>, <b>cinder_store_project_name</b>, and <b>cinder_store_password</b> are specified.</p> <p>If this configuration option is set, <b>cinder_catalog_info</b> will be ignored.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● URL template string for cinder endpoint, where <b>%%(tenant)s</b> is replaced with the current tenant (project) name. For example: <b>http://cinder.openstack.example.org/v2/%%(tenant)s</b></li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● cinder_store_auth_address</li> <li>● cinder_store_user_name</li> <li>● cinder_store_project_name</li> <li>● cinder_store_password</li> <li>● cinder_catalog_info</li> </ul>
<b>cinder_http_retries = 3</b>	integer value	<p>Number of cinderclient retries on failed http calls.</p> <p>When a call failed by any errors, cinderclient will retry the call up to the specified times after sleeping a few seconds.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A positive integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>


Configuration option = Default value	Type	Description
<b>cinder_os_region_name = None</b>	string value	<p>Region name to lookup cinder service from the service catalog.</p> <p>This is used only when <b>cinder_catalog_info</b> is used for determining the endpoint. If set, the lookup for cinder endpoint by this node is filtered to the specified region. It is useful when multiple regions are listed in the catalog. If this is not set, the endpoint is looked up from every region.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A string that is a valid region name.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● cinder_catalog_info</li> </ul>
<b>cinder_state_transition_timeout = 300</b>	integer value	<p>Time period, in seconds, to wait for a cinder volume transition to complete.</p> <p>When the cinder volume is created, deleted, or attached to the glance node to read/write the volume data, the volume's state is changed. For example, the newly created volume status changes from <b>creating</b> to <b>available</b> after the creation process is completed. This specifies the maximum time to wait for the status change. If a timeout occurs while waiting, or the status is changed to an unexpected value (e.g. <b>error</b>), the image creation fails.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A positive integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>

Configuration option = Default value	Type	Description
<b>cinder_store_auth_address = None</b>	string value	<p>The address where the cinder authentication service is listening.</p> <p>When all of <b>cinder_store_auth_address</b>, <b>cinder_store_user_name</b>, <b>cinder_store_project_name</b>, and <b>cinder_store_password</b> options are specified, the specified values are always used for the authentication. This is useful to hide the image volumes from users by storing them in a project/tenant specific to the image service. It also enables users to share the image volume among other projects under the control of glance's ACL.</p> <p>If either of these options are not set, the cinder endpoint is looked up from the service catalog, and current context's user and project are used.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid authentication service address, for example: <b>http://openstack.example.org/identity/v2.0</b></li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● cinder_store_user_name</li> <li>● cinder_store_password</li> <li>● cinder_store_project_name</li> </ul>
<b>cinder_store_password = None</b>	string value	<p>Password for the user authenticating against cinder.</p> <p>This must be used with all the following related options. If any of these are not specified, the user of the current context is used.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid password for the user specified by <b>cinder_store_user_name</b></li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● cinder_store_auth_address</li> <li>● cinder_store_user_name</li> <li>● cinder_store_project_name</li> </ul>



Configuration option = Default value	Type	Description
<b>cinder_store_project_name = None</b>	string value	<p>Project name where the image volume is stored in cinder.</p> <p>If this configuration option is not set, the project in current context is used.</p> <p>This must be used with all the following related options. If any of these are not specified, the project of the current context is used.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid project name</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>cinder_store_auth_address</b></li> <li>● <b>cinder_store_user_name</b></li> <li>● <b>cinder_store_password</b></li> </ul>
<b>cinder_store_user_name = None</b>	string value	<p>User name to authenticate against cinder.</p> <p>This must be used with all the following related options. If any of these are not specified, the user of the current context is used.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid user name</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● cinder_store_auth_address</li> <li>● cinder_store_password</li> <li>● cinder_store_project_name</li> </ul>

Configuration option = Default value	Type	Description
<b>cinder_volume_type = None</b>	string value	<p>Volume type that will be used for volume creation in cinder.</p> <p>Some cinder backends can have several volume types to optimize storage usage. Adding this option allows an operator to choose a specific volume type in cinder that can be optimized for images.</p> <p>If this is not set, then the default volume type specified in the cinder configuration will be used for volume creation.</p> <p>Possible values:</p> <ul style="list-style-type: none"><li>• A valid volume type from cinder</li></ul> <p>Related options:</p> <ul style="list-style-type: none"><li>• None</li></ul>

Configuration option = Default value	Type	Description
<b>default_store = file</b>	string value	<p>The default scheme to use for storing images.</p> <p>Provide a string value representing the default scheme to use for storing images. If not set, Glance uses <b>file</b> as the default scheme to store images with the <b>file</b> store.</p> <div data-bbox="815 510 922 703" style="display: inline-block; vertical-align: top;">  </div> <p><b>NOTE</b></p> <p>The value given for this configuration option must be a valid scheme for a store registered with the <b>stores</b> configuration option.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● file</li> <li>● filesystem</li> <li>● http</li> <li>● https</li> <li>● swift</li> <li>● swift+http</li> <li>● swift+https</li> <li>● swift+config</li> <li>● rbd</li> <li>● sheepdog</li> <li>● cinder</li> <li>● vsphere</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● stores</li> </ul>

Configuration option = Default value	Type	Description
<b>default_swift_reference = ref1</b>	string value	<p>Reference to default Swift account/backing store parameters.</p> <p>Provide a string value representing a reference to the default set of parameters required for using swift account/backing store for image storage. The default reference value for this configuration option is <i>ref1</i>. This configuration option dereferences the parameters and facilitates image storage in Swift storage backend every time a new image is added.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A valid string value</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<b>filesystem_store_datadir = /var/lib/glance/images</b>	string value	<p>Directory to which the filesystem backend store writes images.</p> <p>Upon start up, Glance creates the directory if it doesn't already exist and verifies write access to the user under which <b>glance-api</b> runs. If the write access isn't available, a <b>BadStoreConfiguration</b> exception is raised and the filesystem store may not be available for adding new images.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p>This directory is used only when filesystem store is used as a storage backend. Either <b>filesystem_store_datadir</b> or <b>filesystem_store_datadirs</b> option must be specified in <b>glance-api.conf</b>. If both options are specified, a <b>BadStoreConfiguration</b> will be raised and the filesystem store may not be available for adding new images.</p> </div> </div> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A valid path to a directory</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <b>filesystem_store_datadirs</b></li> <li>• <b>filesystem_store_file_perm</b></li> </ul>

---

Configuration option = Default value	Type	Description
---	------	-------------

Configuration option = Default value	Type	Description
<b>filesystem_store_datadirs</b> = None	multi valued	<p>List of directories and their priorities to which the filesystem backend store writes images.</p> <p>The filesystem store can be configured to store images in multiple directories as opposed to using a single directory specified by the <b>filesystem_store_datadir</b> configuration option. When using multiple directories, each directory can be given an optional priority to specify the preference order in which they should be used. Priority is an integer that is concatenated to the directory path with a colon where a higher value indicates higher priority. When two directories have the same priority, the directory with most free space is used. When no priority is specified, it defaults to zero.</p> <p>More information on configuring filesystem store with multiple store directories can be found at <a href="http://docs.openstack.org/developer/glance/configuring.html">http://docs.openstack.org/developer/glance/configuring.html</a></p> <div style="display: flex; align-items: flex-start;"> <div style="width: 60px; height: 100px; background: repeating-linear-gradient(45deg, transparent, transparent 2px, #ccc 2px, #ccc 4px); margin-right: 10px;"></div> <div> <p><b>NOTE</b></p> <p>This directory is used only when filesystem store is used as a storage backend. Either <b>filesystem_store_datadir</b> or <b>filesystem_store_datadirs</b> option must be specified in <b>glance-api.conf</b>. If both options are specified, a <b>BadStoreConfiguration</b> will be raised and the filesystem store may not be available for adding new images.</p> </div> </div> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● List of strings of the following form:             <ul style="list-style-type: none"> <li>○ <b>&lt;a valid directory path&gt;:&lt;optional integer priority&gt;</b></li> </ul> </li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>filesystem_store_datadir</b></li> <li>● <b>filesystem_store_file_perm</b></li> </ul>

Configuration option = Default value	Type	Description
<b>filesystem_store_file_permissions = 0</b>	integer value	<p>File access permissions for the image files.</p> <p>Set the intended file access permissions for image data. This provides a way to enable other services, e.g. Nova, to consume images directly from the filesystem store. The users running the services that are intended to be given access to could be made a member of the group that owns the files created. Assigning a value less than or equal to zero for this configuration option signifies that no changes be made to the default permissions. This value will be decoded as an octal digit.</p> <p>For more information, please refer the documentation at <a href="http://docs.openstack.org/developer/glance/configuring.html">http://docs.openstack.org/developer/glance/configuring.html</a></p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid file access permission</li> <li>● Zero</li> <li>● Any negative integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>filesystem_store_metadata_file = None</b>	string value	<p>Filesystem store metadata file.</p> <p>The path to a file which contains the metadata to be returned with any location associated with the filesystem store. The file must contain a valid JSON object. The object should contain the keys <b>id</b> and <b>mountpoint</b>. The value for both keys should be a string.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid path to the store metadata file</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>

Configuration option = Default value	Type	Description
<b>http_proxy_information =</b> {}	dict value	<p>The http/https proxy information to be used to connect to the remote server.</p> <p>This configuration option specifies the http/https proxy information that should be used to connect to the remote server. The proxy information should be a key value pair of the scheme and proxy, for example, http:10.0.0.1:3128. You can also specify proxies for multiple schemes by separating the key value pairs with a comma, for example, http:10.0.0.1:3128, https:10.0.0.1:1080.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A comma separated list of scheme:proxy pairs as described above</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<b>https_ca_certificates_file = None</b>	string value	<p>Path to the CA bundle file.</p> <p>This configuration option enables the operator to use a custom Certificate Authority file to verify the remote server certificate. If this option is set, the <b>https_insecure</b> option will be ignored and the CA file specified will be used to authenticate the server certificate and establish a secure connection to the server.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A valid path to a CA file</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• https_insecure</li> </ul>




Configuration option = Default value	Type	Description
<b>https_insecure = True</b>	boolean value	<p>Set verification of the remote server certificate.</p> <p>This configuration option takes in a boolean value to determine whether or not to verify the remote server certificate. If set to True, the remote server certificate is not verified. If the option is set to False, then the default CA truststore is used for verification.</p> <p>This option is ignored if <b>https_ca_certificates_file</b> is set. The remote server certificate will then be verified using the file specified using the <b>https_ca_certificates_file</b> option.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● https_ca_certificates_file</li> </ul>
<b>rados_connect_timeout = 0</b>	integer value	<p>Timeout value for connecting to Ceph cluster.</p> <p>This configuration option takes in the timeout value in seconds used when connecting to the Ceph cluster i.e. it sets the time to wait for glance-api before closing the connection. This prevents glance-api hangups during the connection to RBD. If the value for this option is set to less than or equal to 0, no timeout is set and the default librados value is used.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>● Any integer value</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>

Configuration option = Default value	Type	Description
<b>rbd_store_ceph_conf = /etc/ceph/ceph.conf</b>	string value	<p>Ceph configuration file path.</p> <p>This configuration option takes in the path to the Ceph configuration file to be used. If the value for this option is not set by the user or is set to None, librados will locate the default configuration file which is located at /etc/ceph/ceph.conf. If using Cephx authentication, this file should include a reference to the right keyring in a client.&lt;USER&gt; section</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>● A valid path to a configuration file</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● rbd_store_user</li> </ul>
<b>rbd_store_chunk_size = 8</b>	integer value	<p>Size, in megabytes, to chunk RADOS images into.</p> <p>Provide an integer value representing the size in megabytes to chunk Glance images into. The default chunk size is 8 megabytes. For optimal performance, the value should be a power of two.</p> <p>When Ceph's RBD object storage system is used as the storage backend for storing Glance images, the images are chunked into objects of the size set using this option. These chunked objects are then stored across the distributed block data store to use for Glance.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>● Any positive integer value</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>

Configuration option = Default value	Type	Description
<b>rbd_store_pool = images</b>	string value	<p>RADOS pool in which images are stored.</p> <p>When RBD is used as the storage backend for storing Glance images, the images are stored by means of logical grouping of the objects (chunks of images) into a <b>pool</b>. Each pool is defined with the number of placement groups it can contain. The default pool that is used is <i>images</i>.</p> <p>More information on the RBD storage backend can be found here: <a href="http://ceph.com/planet/how-data-is-stored-in-ceph-cluster/">http://ceph.com/planet/how-data-is-stored-in-ceph-cluster/</a></p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>● A valid pool name</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>rbd_store_user = None</b>	string value	<p>RADOS user to authenticate as.</p> <p>This configuration option takes in the RADOS user to authenticate as. This is only needed when RADOS authentication is enabled and is applicable only if the user is using Cephx authentication. If the value for this option is not set by the user or is set to None, a default value will be chosen, which will be based on the <code>client.</code> section in <code>rbd_store_ceph_conf</code>.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>● A valid RADOS user</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>rbd_store_ceph_conf</code></li> </ul>

Configuration option = Default value	Type	Description
<b>rootwrap_config = /etc/glance/rootwrap.conf</b>	string value	<p>Path to the rootwrap configuration file to use for running commands as root.</p> <p>The cinder store requires root privileges to operate the image volumes (for connecting to iSCSI/FC volumes and reading/writing the volume data, etc.). The configuration file should allow the required commands by cinder store and os-brick library.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Path to the rootwrap config file</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>sheepdog_store_address = 127.0.0.1</b>	unknown value	<p>Address to bind the Sheepdog daemon to.</p> <p>Provide a string value representing the address to bind the Sheepdog daemon to. The default address set for the <i>sheep</i> is 127.0.0.1.</p> <p>The Sheepdog daemon, also called <i>sheep</i>, manages the storage in the distributed cluster by writing objects across the storage network. It identifies and acts on the messages directed to the address set using <b>sheepdog_store_address</b> option to store chunks of Glance images.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid IPv4 address</li> <li>● A valid IPv6 address</li> <li>● A valid hostname</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● sheepdog_store_port</li> </ul>

Configuration option = Default value	Type	Description
<b>sheepdog_store_chunk_size = 64</b>	integer value	<p>Chunk size for images to be stored in Sheepdog data store.</p> <p>Provide an integer value representing the size in mebibyte (1048576 bytes) to chunk Glance images into. The default chunk size is 64 mebibytes.</p> <p>When using Sheepdog distributed storage system, the images are chunked into objects of this size and then stored across the distributed data store to use for Glance.</p> <p>Chunk sizes, if a power of two, help avoid fragmentation and enable improved performance.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Positive integer value representing size in mebibytes.</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>sheepdog_store_port = 7000</b>	port value	<p>Port number on which the sheep daemon will listen.</p> <p>Provide an integer value representing a valid port number on which you want the Sheepdog daemon to listen on. The default port is 7000.</p> <p>The Sheepdog daemon, also called <i>sheep</i>, manages the storage in the distributed cluster by writing objects across the storage network. It identifies and acts on the messages it receives on the port number set using <b>sheepdog_store_port</b> option to store chunks of Glance images.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid port number (0 to 65535)</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● <code>sheepdog_store_address</code></li> </ul>


Configuration option = Default value	Type	Description
<b>store_capabilities_update_min_interval = 0</b>	integer value	<p>Minimum interval in seconds to execute updating dynamic storage capabilities based on current backend status.</p> <p>Provide an integer value representing time in seconds to set the minimum interval before an update of dynamic storage capabilities for a storage backend can be attempted. Setting <b>store_capabilities_update_min_interval</b> does not mean updates occur periodically based on the set interval. Rather, the update is performed at the elapse of this interval set, if an operation of the store is triggered.</p> <p>By default, this option is set to zero and is disabled. Provide an integer value greater than zero to enable this option.</p> <div data-bbox="815 898 922 1155" style="display: inline-block; vertical-align: top;">  </div> <p><b>NOTE</b></p> <p>For more information on store capabilities and their updates, please visit:  <a href="https://specs.openstack.org/openstack/glance-specs/specs/kilo/store-capabilities.html">https://specs.openstack.org/openstack/glance-specs/specs/kilo/store-capabilities.html</a></p> <p>For more information on setting up a particular store in your deployment and help with the usage of this feature, please contact the storage driver maintainers listed here:  <a href="http://docs.openstack.org/developer/glance_store/drivers/index.html">http://docs.openstack.org/developer/glance_store/drivers/index.html</a></p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Zero</li> <li>● Positive integer</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>

Configuration option = Default value	Type	Description
<b>stores = [u'file', u'http']</b>	list value	<p>List of enabled Glance stores.</p> <p>Register the storage backends to use for storing disk images as a comma separated list. The default stores enabled for storing disk images with Glance are <b>file</b> and <b>http</b>.</p> <p>Possible values:</p> <ul style="list-style-type: none"><li>● A comma separated list that could include:<ul style="list-style-type: none"><li>○ file</li><li>○ http</li><li>○ swift</li><li>○ rbd</li><li>○ sheepdog</li><li>○ cinder</li><li>○ vmware</li></ul></li></ul> <p>Related Options:</p> <ul style="list-style-type: none"><li>● default_store</li></ul>


Configuration option = Default value	Type	Description
<b>swift_buffer_on_upload = False</b>	boolean value	<p>Buffer image segments before upload to Swift.</p> <p>Provide a boolean value to indicate whether or not Glance should buffer image data to disk while uploading to swift. This enables Glance to resume uploads on error.</p> <p>NOTES: When enabling this option, one should take great care as this increases disk usage on the API node. Be aware that depending upon how the file system is configured, the disk space used for buffering may decrease the actual disk space available for the glance image cache. Disk utilization will cap according to the following equation: <b>(swift_store_large_object_chunk_size * workers * 1000)</b></p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• True</li> <li>• False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• swift_upload_buffer_dir</li> </ul>
<b>swift_store_admin_tenant s = []</b>	list value	<p>List of tenants that will be granted admin access.</p> <p>This is a list of tenants that will be granted read/write access on all Swift containers created by Glance in multi-tenant mode. The default value is an empty list.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A comma separated list of strings representing UUIDs of Keystone projects/tenants</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<b>swift_store_auth_address = None</b>	string value	<p>The address where the Swift authentication service is listening.</p>



Configuration option = Default value	Type	Description
<b>swift_store_auth_insecure = False</b>	boolean value	<p>Set verification of the server certificate.</p> <p>This boolean determines whether or not to verify the server certificate. If this option is set to True, swiftclient won't check for a valid SSL certificate when authenticating. If the option is set to False, then the default CA truststore is used for verification.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• True</li> <li>• False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• swift_store_cacert</li> </ul>
<b>swift_store_auth_version = 2</b>	string value	<p>Version of the authentication service to use. Valid versions are 2 and 3 for keystone and 1 (deprecated) for swauth and rackspace.</p>
<b>swift_store_cacert = None</b>	string value	<p>Path to the CA bundle file.</p> <p>This configuration option enables the operator to specify the path to a custom Certificate Authority file for SSL verification when connecting to Swift.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A valid path to a CA file</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• swift_store_auth_insecure</li> </ul>


Configuration option = Default value	Type	Description
<b>swift_store_config_file =</b> <b>None</b>	string value	<p>Absolute path to the file containing the swift account(s) configurations.</p> <p>Include a string value representing the path to a configuration file that has references for each of the configured Swift account(s)/backing stores. By default, no file path is specified and customized Swift referencing is disabled. Configuring this option is highly recommended while using Swift storage backend for image storage as it avoids storage of credentials in the database.</p> <div data-bbox="815 689 922 882" style="display: inline-block; vertical-align: top;">  </div> <p><b>NOTE</b></p> <p>Please do not configure this option if you have set <b>swift_store_multi_tenant</b> to <b>True</b>.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● String value representing an absolute path on the glance-api node</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● swift_store_multi_tenant</li> </ul>

Configuration option = Default value	Type	Description
<b>swift_store_container = glance</b>	string value	<p>Name of single container to store images/name prefix for multiple containers</p> <p>When a single container is being used to store images, this configuration option indicates the container within the Glance account to be used for storing all images. When multiple containers are used to store images, this will be the name prefix for all containers. Usage of single/multiple containers can be controlled using the configuration option <b>swift_store_multiple_containers_seed</b>.</p> <p>When using multiple containers, the containers will be named after the value set for this configuration option with the first N chars of the image UUID as the suffix delimited by an underscore (where N is specified by <b>swift_store_multiple_containers_seed</b>).</p> <p>Example: if the seed is set to 3 and <code>swift_store_container = glance</code>, then an image with UUID <b>fdae39a1-bac5-4238-aba4-69bcc726e848</b> would be placed in the container <b>glance_fda</b>. All dashes in the UUID are included when creating the container name but do not count toward the character limit, so when N=10 the container name would be <b>glance_fdae39a1-ba</b>.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● If using single container, this configuration option can be any string that is a valid swift container name in Glance's Swift account</li> <li>● If using multiple containers, this configuration option can be any string as long as it satisfies the container naming rules enforced by Swift. The value of <b>swift_store_multiple_containers_seed</b> should be taken into account as well.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>swift_store_multiple_containers_seed</b></li> <li>● <b>swift_store_multi_tenant</b></li> <li>● <b>swift_store_create_container_on_put</b></li> </ul>

Configuration option = Default value	Type	Description
<b>swift_store_create_container_on_put = False</b>	boolean value	<p>Create container, if it doesn't already exist, when uploading image.</p> <p>At the time of uploading an image, if the corresponding container doesn't exist, it will be created provided this configuration option is set to True. By default, it won't be created. This behavior is applicable for both single and multiple containers mode.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• True</li> <li>• False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<b>swift_store_endpoint = None</b>	string value	<p>The URL endpoint to use for Swift backend storage.</p> <p>Provide a string value representing the URL endpoint to use for storing Glance images in Swift store. By default, an endpoint is not set and the storage URL returned by <b>auth</b> is used. Setting an endpoint with <b>swift_store_endpoint</b> overrides the storage URL and is used for Glance image storage.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p>The URL should include the path up to, but excluding the container. The location of an object is obtained by appending the container and object to the configured URL.</p> </div> </div> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• String value representing a valid URL path up to a Swift container</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>

Configuration option = Default value	Type	Description
<b>swift_store_endpoint_type = publicURL</b>	string value	<p>Endpoint Type of Swift service.</p> <p>This string value indicates the endpoint type to use to fetch the Swift endpoint. The endpoint type determines the actions the user will be allowed to perform, for instance, reading and writing to the Store. This setting is only used if <code>swift_store_auth_version</code> is greater than 1.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● <code>publicURL</code></li> <li>● <code>adminURL</code></li> <li>● <code>internalURL</code></li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>swift_store_endpoint</code></li> </ul>
<b>swift_store_expire_soon_interval = 60</b>	integer value	<p>Time in seconds defining the size of the window in which a new token may be requested before the current token is due to expire.</p> <p>Typically, the Swift storage driver fetches a new token upon the expiration of the current token to ensure continued access to Swift. However, some Swift transactions (like uploading image segments) may not recover well if the token expires on the fly.</p> <p>Hence, by fetching a new token before the current token expiration, we make sure that the token does not expire or is close to expiry before a transaction is attempted. By default, the Swift storage driver requests for a new token 60 seconds or less before the current token expiration.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Zero</li> <li>● Positive integer value</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>swift_store_key = None</b>	string value	Auth key for the user authenticating against the Swift authentication service.


Configuration option = Default value	Type	Description
<b>swift_store_large_object_ chunk_size = 200</b>	integer value	<p>The maximum size, in MB, of the segments when image data is segmented.</p> <p>When image data is segmented to upload images that are larger than the limit enforced by the Swift cluster, image data is broken into segments that are no bigger than the size specified by this configuration option. Refer to <b>swift_store_large_object_size</b> for more detail.</p> <p>For example: if <b>swift_store_large_object_size</b> is 5GB and <b>swift_store_large_object_chunk_size</b> is 1GB, an image of size 6.2GB will be segmented into 7 segments where the first six segments will be 1GB in size and the seventh segment will be 0.2GB.</p> <p>Possible values:</p> <ul style="list-style-type: none"><li>• A positive integer that is less than or equal to the large object limit enforced by Swift cluster in consideration.</li></ul> <p>Related options:</p> <ul style="list-style-type: none"><li>• <b>swift_store_large_object_size</b></li></ul>


Configuration option = Default value	Type	Description
<b>swift_store_large_object_size = 5120</b>	integer value	<p>The size threshold, in MB, after which Glance will start segmenting image data.</p> <p>Swift has an upper limit on the size of a single uploaded object. By default, this is 5GB. To upload objects bigger than this limit, objects are segmented into multiple smaller objects that are tied together with a manifest file. For more detail, refer to <a href="http://docs.openstack.org/developer/swift/overview_large_objects.html">http://docs.openstack.org/developer/swift/overview_large_objects.html</a></p> <p>This configuration option specifies the size threshold over which the Swift driver will start segmenting image data into multiple smaller files. Currently, the Swift driver only supports creating Dynamic Large Objects.</p> <div data-bbox="815 864 922 1055" style="display: inline-block; vertical-align: top;">  </div> <p><b>NOTE</b></p> <p>This should be set by taking into account the large object limit enforced by the Swift cluster in consideration.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A positive integer that is less than or equal to the large object limit enforced by the Swift cluster in consideration.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>swift_store_large_object_chunk_size</b></li> </ul>

Configuration option = Default value	Type	Description
<b>swift_store_multi_tenant</b> <b>= False</b>	boolean value	<p>Store images in tenant's Swift account.</p> <p>This enables multi-tenant storage mode which causes Glance images to be stored in tenant specific Swift accounts. If this is disabled, Glance stores all images in its own account. More details multi-tenant store can be found at <a href="https://wiki.openstack.org/wiki/GlanceSwiftTenantSpecificStorage">https://wiki.openstack.org/wiki/GlanceSwiftTenantSpecificStorage</a></p> <div data-bbox="815 613 922 808" style="float: left; margin-right: 10px;"> </div> <p><b>NOTE</b></p> <p>If using multi-tenant swift store, please make sure that you do not set a swift configuration file with the <i>swift_store_config_file</i> option.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>swift_store_config_file</code></li> </ul>




Configuration option = Default value	Type	Description
<b>swift_store_multiple_containers_seed = 0</b>	integer value	<p>Seed indicating the number of containers to use for storing images.</p> <p>When using a single-tenant store, images can be stored in one or more than one containers. When set to 0, all images will be stored in one single container. When set to an integer value between 1 and 32, multiple containers will be used to store images. This configuration option will determine how many containers are created. The total number of containers that will be used is equal to <math>16^N</math>, so if this config option is set to 2, then <math>16^2=256</math> containers will be used to store images.</p> <p>Please refer to <b>swift_store_container</b> for more detail on the naming convention. More detail about using multiple containers can be found at <a href="https://specs.openstack.org/openstack/glance-specs/specs/kilo/swift-store-multiple-containers.html">https://specs.openstack.org/openstack/glance-specs/specs/kilo/swift-store-multiple-containers.html</a></p> <div data-bbox="815 1003 922 1137"> </div> <p><b>NOTE</b></p> <p>This is used only when <code>swift_store_multi_tenant</code> is disabled.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A non-negative integer less than or equal to 32</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <b>swift_store_container</b></li> <li>• <b>swift_store_multi_tenant</b></li> <li>• <b>swift_store_create_container_on_put</b></li> </ul>

Configuration option = Default value	Type	Description
<b>swift_store_region = None</b>	string value	<p>The region of Swift endpoint to use by Glance.</p> <p>Provide a string value representing a Swift region where Glance can connect to for image storage. By default, there is no region set.</p> <p>When Glance uses Swift as the storage backend to store images for a specific tenant that has multiple endpoints, setting of a Swift region with <b>swift_store_region</b> allows Glance to connect to Swift in the specified region as opposed to a single region connectivity.</p> <p>This option can be configured for both single-tenant and multi-tenant storage.</p> <div data-bbox="815 819 922 1048" style="display: inline-block; vertical-align: top;">  </div> <p><b>NOTE</b></p> <p>Setting the region with <b>swift_store_region</b> is tenant-specific and is necessary <b>only if</b> the tenant has multiple endpoints across different regions.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A string value representing a valid Swift region.</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>

Configuration option = Default value	Type	Description
<b>swift_store_retry_get_count = 0</b>	integer value	<p>The number of times a Swift download will be retried before the request fails.</p> <p>Provide an integer value representing the number of times an image download must be retried before erroring out. The default value is zero (no retry on a failed image download). When set to a positive integer value, <b>swift_store_retry_get_count</b> ensures that the download is attempted this many more times upon a download failure before sending an error message.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Zero</li> <li>• Positive integer value</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<b>swift_store_service_type = object-store</b>	string value	<p>Type of Swift service to use.</p> <p>Provide a string value representing the service type to use for storing images while using Swift backend storage. The default service type is set to <b>object-store</b>.</p> <div data-bbox="815 1294 922 1547" style="display: inline-block; vertical-align: top;">  </div> <p><b>NOTE</b></p> <p>If <b>swift_store_auth_version</b> is set to 2, the value for this configuration option needs to be <b>object-store</b>. If using a higher version of Keystone or a different auth scheme, this option may be modified.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A string representing a valid service type for Swift storage.</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>

Configuration option = Default value	Type	Description
<b>swift_store_ssl_compression = True</b>	boolean value	<p>SSL layer compression for HTTPS Swift requests.</p> <p>Provide a boolean value to determine whether or not to compress HTTPS Swift requests for images at the SSL layer. By default, compression is enabled.</p> <p>When using Swift as the backend store for Glance image storage, SSL layer compression of HTTPS Swift requests can be set using this option. If set to False, SSL layer compression of HTTPS Swift requests is disabled. Disabling this option may improve performance for images which are already in a compressed format, for example, qcow2.</p> <p>Possible values:</p> <ul style="list-style-type: none"><li>• True</li><li>• False</li></ul> <p>Related Options:</p> <ul style="list-style-type: none"><li>• None</li></ul>

Configuration option = Default value	Type	Description
<b>swift_store_use_trusts = True</b>	boolean value	<p>Use trusts for multi-tenant Swift store.</p> <p>This option instructs the Swift store to create a trust for each add/get request when the multi-tenant store is in use. Using trusts allows the Swift store to avoid problems that can be caused by an authentication token expiring during the upload or download of data.</p> <p>By default, <b>swift_store_use_trusts</b> is set to <b>True</b>(use of trusts is enabled). If set to <b>False</b>, a user token is used for the Swift connection instead, eliminating the overhead of trust creation.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p>This option is considered only when <b>swift_store_multi_tenant</b> is set to <b>True</b></p> </div> </div> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● swift_store_multi_tenant</li> </ul>
<b>swift_store_user = None</b>	string value	The user to authenticate against the Swift authentication service.

Configuration option = Default value	Type	Description
<b>swift_upload_buffer_dir = None</b>	string value	<p>Directory to buffer image segments before upload to Swift.</p> <p>Provide a string value representing the absolute path to the directory on the glance node where image segments will be buffered briefly before they are uploaded to swift.</p> <p>NOTES: * This is required only when the configuration option <b>swift_buffer_on_upload</b> is set to True. * This directory should be provisioned keeping in mind the <b>swift_store_large_object_chunk_size</b> and the maximum number of images that could be uploaded simultaneously by a given glance node.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● String value representing an absolute directory path</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>swift_buffer_on_upload</code></li> <li>● <code>swift_store_large_object_chunk_size</code></li> </ul>
<b>vmware_api_retry_count = 10</b>	integer value	<p>The number of VMware API retries.</p> <p>This configuration option specifies the number of times the VMware ESX/VC server API must be retried upon connection related issues or server API call overload. It is not possible to specify <i>retry forever</i>.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>● Any positive integer value</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>

Configuration option = Default value	Type	Description
<b>vmware_ca_file = None</b>	string value	<p>Absolute path to the CA bundle file.</p> <p>This configuration option enables the operator to use a custom Certificate Authority File to verify the ESX/vCenter certificate.</p> <p>If this option is set, the "vmware_insecure" option will be ignored and the CA file specified will be used to authenticate the ESX/vCenter server certificate and establish a secure connection to the server.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>Any string that is a valid absolute path to a CA file</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>vmware_insecure</li> </ul>
<b>vmware_datastores = None</b>	multi valued	<p>The datastores where the image can be stored.</p> <p>This configuration option specifies the datastores where the image can be stored in the VMWare store backend. This option may be specified multiple times for specifying multiple datastores. The datastore name should be specified after its datacenter path, separated by ":". An optional weight may be given after the datastore name, separated again by ":" to specify the priority. Thus, the required format becomes &lt;datacenter_path&gt;:&lt;datastore_name&gt;:&lt;optional_weight&gt;.</p> <p>When adding an image, the datastore with highest weight will be selected, unless there is not enough free space available in cases where the image size is already known. If no weight is given, it is assumed to be zero and the directory will be considered for selection last. If multiple datastores have the same weight, then the one with the most free space available is selected.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>Any string of the format: &lt;datacenter_path&gt;:&lt;datastore_name&gt;:&lt;optional_weight&gt;</li> </ul> <p>Related options: * None</p>

Configuration option = Default value	Type	Description
<b>vmware_insecure = False</b>	boolean value	<p>Set verification of the ESX/vCenter server certificate.</p> <p>This configuration option takes a boolean value to determine whether or not to verify the ESX/vCenter server certificate. If this option is set to True, the ESX/vCenter server certificate is not verified. If this option is set to False, then the default CA truststore is used for verification.</p> <p>This option is ignored if the "vmware_ca_file" option is set. In that case, the ESX/vCenter server certificate will then be verified using the file specified using the "vmware_ca_file" option .</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● vmware_ca_file</li> </ul>
<b>vmware_server_host = None</b>	unknown value	<p>Address of the ESX/ESXi or vCenter Server target system.</p> <p>This configuration option sets the address of the ESX/ESXi or vCenter Server target system. This option is required when using the VMware storage backend. The address can contain an IP address (127.0.0.1) or a DNS name (www.my-domain.com).</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>● A valid IPv4 or IPv6 address</li> <li>● A valid DNS name</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● vmware_server_username</li> <li>● vmware_server_password</li> </ul>



Configuration option = Default value	Type	Description
<b>vmware_server_password = None</b>	string value	<p>Server password.</p> <p>This configuration option takes the password for authenticating with the VMware ESX/ESXi or vCenter Server. This option is required when using the VMware storage backend.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>Any string that is a password corresponding to the username specified using the "vmware_server_username" option</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>vmware_server_host</li> <li>vmware_server_username</li> </ul>
<b>vmware_server_username = None</b>	string value	<p>Server username.</p> <p>This configuration option takes the username for authenticating with the VMware ESX/ESXi or vCenter Server. This option is required when using the VMware storage backend.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>Any string that is the username for a user with appropriate privileges</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>vmware_server_host</li> <li>vmware_server_password</li> </ul>
<b>vmware_store_image_dir = /openstack_glance</b>	string value	<p>The directory where the glance images will be stored in the datastore.</p> <p>This configuration option specifies the path to the directory where the glance images will be stored in the VMware datastore. If this option is not set, the default directory where the glance images are stored is openstack_glance.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>Any string that is a valid path to a directory</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>None</li> </ul>

Configuration option = Default value	Type	Description
<b>vmware_task_poll_interval = 5</b>	integer value	<p>Interval in seconds used for polling remote tasks invoked on VMware ESX/VC server.</p> <p>This configuration option takes in the sleep time in seconds for polling an on-going async task as part of the VMWare ESX/VC server API call.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>Any positive integer value</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>None</li> </ul>

### 5.1.5. image\_format

The following table outlines the options available under the **[image\_format]** group in the `/etc/glance/glance-api.conf` file.

Table 5.5. image\_format

Configuration option = Default value	Type	Description
<b>container_formats = [u'ami', u'ari', u'aki', u'bare', u'ovf', u'ova', u'docker']</b>	list value	Supported values for the <i>container_format</i> image attribute

Configuration option = Default value	Type	Description
<b>disk_formats = [u'ami', u'ari', u'aki', u'vhd', u'vhdx', u'vmdk', u'raw', u'qcow2', u'vdi', u'iso', u'ploop']</b>	list value	Supported values for the <i>disk_format</i> image attribute

### 5.1.6. keystone\_authtoken

The following table outlines the options available under the **[keystone\_authtoken]** group in the `/etc/glance/glance-api.conf` file.

Table 5.6. keystone\_authtoken

Configuration option = Default value	Type	Description
<b>auth_section = None</b>	string value	Config Section from which to load plugin specific options
<b>auth_type = None</b>	string value	Authentication type to load
<b>auth_uri = None</b>	string value	Complete "public" Identity API endpoint. This endpoint should not be an "admin" endpoint, as it should be accessible by all end users. Unauthenticated clients are redirected to this endpoint to authenticate. Although this endpoint should ideally be unversioned, client support in the wild varies. If you're using a versioned v2 endpoint here, then this should <b>not</b> be the same endpoint the service user utilizes for validating tokens, because normal end users may not be able to reach that endpoint. This option is deprecated in favor of <code>www_authenticate_uri</code> and will be removed in the S release.
<b>auth_version = None</b>	string value	API version of the admin Identity API endpoint.
<b>cache = None</b>	string value	Request environment key where the Swift cache object is stored. When <code>auth_token</code> middleware is deployed with a Swift cache, use this option to have the middleware share a caching backend with swift. Otherwise, use the <b>memcached_servers</b> option instead.
<b>cafile = None</b>	string value	A PEM encoded Certificate Authority to use when verifying HTTPs connections. Defaults to system CAs.

Configuration option = Default value	Type	Description
<b>certfile = None</b>	string value	Required if identity server requires client certificate
<b>check_revocations_for_cached = False</b>	boolean value	If true, the revocation list will be checked for cached tokens. This requires that PKI tokens are configured on the identity server.
<b>delay_auth_decision = False</b>	boolean value	Do not handle authorization requests within the middleware, but delegate the authorization decision to downstream WSGI components.
<b>enforce_token_bind = permissive</b>	string value	Used to control the use and type of token binding. Can be set to: "disabled" to not check token binding. "permissive" (default) to validate binding information if the bind type is of a form known to the server and ignore it if not. "strict" like "permissive" but if the bind type is unknown the token will be rejected. "required" any form of token binding is needed to be allowed. Finally the name of a binding method that must be present in tokens.
<b>hash_algorithms = [u'md5']</b>	list value	Hash algorithms to use for hashing PKI tokens. This may be a single algorithm or multiple. The algorithms are those supported by Python standard hashlib.new(). The hashes will be tried in the order given, so put the preferred one first for performance. The result of the first hash will be stored in the cache. This will typically be set to multiple values only while migrating from a less secure algorithm to a more secure one. Once all the old tokens are expired this option should be set to a single value for better performance.
<b>http_connect_timeout = None</b>	integer value	Request timeout value for communicating with Identity API server.
<b>http_request_max_retries = 3</b>	integer value	How many times are we trying to reconnect when communicating with Identity API Server.
<b>include_service_catalog = True</b>	boolean value	(Optional) Indicate whether to set the X-Service-Catalog header. If False, middleware will not ask for service catalog on token validation and will not set the X-Service-Catalog header.
<b>insecure = False</b>	boolean value	Verify HTTPS connections.
<b>keyfile = None</b>	string value	Required if identity server requires client certificate

Configuration option = Default value	Type	Description
<b>memcache_pool_conn_get_timeout = 10</b>	integer value	(Optional) Number of seconds that an operation will wait to get a memcached client connection from the pool.
<b>memcache_pool_dead_retry = 300</b>	integer value	(Optional) Number of seconds memcached server is considered dead before it is tried again.
<b>memcache_pool_maxsize = 10</b>	integer value	(Optional) Maximum total number of open connections to every memcached server.
<b>memcache_pool_socket_timeout = 3</b>	integer value	(Optional) Socket timeout in seconds for communicating with a memcached server.
<b>memcache_pool_unused_timeout = 60</b>	integer value	(Optional) Number of seconds a connection to memcached is held unused in the pool before it is closed.
<b>memcache_secret_key = None</b>	string value	(Optional, mandatory if memcache_security_strategy is defined) This string is used for key derivation.
<b>memcache_security_strategy = None</b>	string value	(Optional) If defined, indicate whether token data should be authenticated or authenticated and encrypted. If MAC, token data is authenticated (with HMAC) in the cache. If ENCRYPT, token data is encrypted and authenticated in the cache. If the value is not one of these options or empty, auth_token will raise an exception on initialization.
<b>memcache_use_advanced_pool = False</b>	boolean value	(Optional) Use the advanced (eventlet safe) memcached client pool. The advanced pool will only work under python 2.x.
<b>memcached_servers = None</b>	list value	Optionally specify a list of memcached server(s) to use for caching. If left undefined, tokens will instead be cached in-process.
<b>region_name = None</b>	string value	The region in which the identity server can be found.
<b>revocation_cache_time = 10</b>	integer value	Determines the frequency at which the list of revoked tokens is retrieved from the Identity service (in seconds). A high number of revocation events combined with a low cache duration may significantly reduce performance. Only valid for PKI tokens. This option has been deprecated in the Ocata release and will be removed in the P release.

Configuration option = Default value	Type	Description
<b>service_token_roles = [u'service']</b>	list value	A choice of roles that must be present in a service token. Service tokens are allowed to request that an expired token can be used and so this check should tightly control that only actual services should be sending this token. Roles here are applied as an ANY check so any role in this list must be present. For backwards compatibility reasons this currently only affects the allow_expired check.
<b>service_token_roles_required = False</b>	boolean value	For backwards compatibility reasons we must let valid service tokens pass that don't pass the service_token_roles check as valid. Setting this true will become the default in a future release and should be enabled if possible.
<b>signing_dir = None</b>	string value	Directory used to cache files related to PKI tokens. This option has been deprecated in the Ocata release and will be removed in the P release.
<b>token_cache_time = 300</b>	integer value	In order to prevent excessive effort spent validating tokens, the middleware caches previously-seen tokens for a configurable duration (in seconds). Set to -1 to disable caching completely.
<b>www_authenticate_uri = None</b>	string value	Complete "public" Identity API endpoint. This endpoint should not be an "admin" endpoint, as it should be accessible by all end users. Unauthenticated clients are redirected to this endpoint to authenticate. Although this endpoint should ideally be unversioned, client support in the wild varies. If you're using a versioned v2 endpoint here, then this should <b>not</b> be the same endpoint the service user utilizes for validating tokens, because normal end users may not be able to reach that endpoint.

### 5.1.7. matchmaker\_redis

The following table outlines the options available under the **[matchmaker\_redis]** group in the `/etc/glance/glance-api.conf` file.

Table 5.7. matchmaker\_redis

Configuration option = Default value	Type	Description
<b>check_timeout = 20000</b>	integer value	Time in ms to wait before the transaction is killed.

Configuration option = Default value	Type	Description
<b>host = 127.0.0.1</b>	string value	Host to locate redis.
<code>`password = `</code>	string value	Password for Redis server (optional).
<b>port = 6379</b>	port value	Use this port to connect to redis host.
<b>sentinel_group_name = oslo-messaging-zeromq</b>	string value	Redis replica set name.
<b>sentinel_hosts = []</b>	list value	List of Redis Sentinel hosts (fault tolerance mode), e.g., [host:port, host1:port ... ]
<b>socket_timeout = 10000</b>	integer value	Timeout in ms on blocking socket operations.
<b>wait_timeout = 2000</b>	integer value	Time in ms to wait between connection attempts.

### 5.1.8. oslo\_concurrency

The following table outlines the options available under the **[oslo\_concurrency]** group in the `/etc/glance/glance-api.conf` file.

Table 5.8. oslo\_concurrency

Configuration option = Default value	Type	Description
<b>disable_process_locking = False</b>	boolean value	Enables or disables inter-process locks.
<b>lock_path = None</b>	string value	Directory to use for lock files. For security, the specified directory should only be writable by the user running the processes that need locking. Defaults to environment variable OSLO_LOCK_PATH. If external locks are used, a lock path must be set.

### 5.1.9. oslo\_messaging\_amqp

The following table outlines the options available under the **[oslo\_messaging\_amqp]** group in the `/etc/glance/glance-api.conf` file.

Table 5.9. oslo\_messaging\_amqp

Configuration option = Default value	Type	Description
<b>addressing_mode = dynamic</b>	string value	Indicates the addressing mode used by the driver. Permitted values: <i>legacy</i> - use legacy non-routable addressing <i>routable</i> - use routable addresses <i>dynamic</i> - use legacy addresses if the message bus does not support routing otherwise use routable addressing
<b>allow_insecure_clients = False</b>	boolean value	Accept clients using either SSL or plain TCP
<b>anycast_address = anycast</b>	string value	Appended to the address prefix when sending to a group of consumers. Used by the message bus to identify messages that should be delivered in a round-robin fashion across consumers.
<b>broadcast_prefix = broadcast</b>	string value	address prefix used when broadcasting to all servers
<b>connection_retry_backoff = 2</b>	integer value	Increase the <code>connection_retry_interval</code> by this many seconds after each unsuccessful failover attempt.
<b>connection_retry_interval = 1</b>	integer value	Seconds to pause before attempting to re-connect.
<b>connection_retry_interval _max = 30</b>	integer value	Maximum limit for <code>connection_retry_interval</code> + <code>connection_retry_backoff</code>
<b>container_name = None</b>	string value	Name for the AMQP container. must be globally unique. Defaults to a generated UUID
<b>default_notification_exch ange = None</b>	string value	Exchange name used in notification addresses. Exchange name resolution precedence: Target.exchange if set else <code>default_notification_exchange</code> if set else <code>control_exchange</code> if set else <i>notify</i>
<b>default_notify_timeout = 30</b>	integer value	The deadline for a sent notification message delivery. Only used when caller does not provide a timeout expiry.
<b>default_reply_retry = 0</b>	integer value	The maximum number of attempts to re-send a reply message which failed due to a recoverable error.
<b>default_reply_timeout = 30</b>	integer value	The deadline for an rpc reply message delivery.



Configuration option = Default value	Type	Description
<b>default_rpc_exchange = None</b>	string value	Exchange name used in RPC addresses. Exchange name resolution precedence: Target.exchange if set else default_rpc_exchange if set else control_exchange if set else <i>rpc</i>
<b>default_send_timeout = 30</b>	integer value	The deadline for an rpc cast or call message delivery. Only used when caller does not provide a timeout expiry.
<b>default_sender_link_timeout = 600</b>	integer value	The duration to schedule a purge of idle sender links. Detach link after expiry.
<b>group_request_prefix = unicast</b>	string value	address prefix when sending to any server in group
<b>idle_timeout = 0</b>	integer value	Timeout for inactive connections (in seconds)
<b>link_retry_delay = 10</b>	integer value	Time to pause between re-connecting an AMQP 1.0 link that failed due to a recoverable error.
<b>multicast_address = multicast</b>	string value	Appended to the address prefix when sending a fanout message. Used by the message bus to identify fanout messages.
<b>notify_address_prefix = openstack.org/om/notify</b>	string value	Address prefix for all generated Notification addresses
<b>notify_server_credit = 100</b>	integer value	Window size for incoming Notification messages
<code>`password = `</code>	string value	Password for message broker authentication
<b>pre_settled = [u'rpc-cast', u'rpc-reply']</b>	multi valued	Send messages of this type pre-settled. Pre-settled messages will not receive acknowledgement from the peer. Note well: pre-settled messages may be silently discarded if the delivery fails. Permitted values: <i>rpc-call</i> - send RPC Calls pre-settled <i>rpc-reply</i> - send RPC Replies pre-settled <i>rpc-cast</i> - Send RPC Casts pre-settled <i>notify</i> - Send Notifications pre-settled

Configuration option = Default value	Type	Description
<b>pseudo_vhost = True</b>	boolean value	Enable virtual host support for those message buses that do not natively support virtual hosting (such as qpid). When set to true the virtual host name will be added to all message bus addresses, effectively creating a private <i>subnet</i> per virtual host. Set to False if the message bus supports virtual hosting using the <i>hostname</i> field in the AMQP 1.0 Open performative as the name of the virtual host.
<b>reply_link_credit = 200</b>	integer value	Window size for incoming RPC Reply messages.
<b>rpc_address_prefix = openstack.org/om/rpc</b>	string value	Address prefix for all generated RPC addresses
<b>rpc_server_credit = 100</b>	integer value	Window size for incoming RPC Request messages
<code>`sasl_config_dir = `</code>	string value	Path to directory that contains the SASL configuration
<code>`sasl_config_name = `</code>	string value	Name of configuration file (without .conf suffix)
<code>`sasl_default_realm = `</code>	string value	SASL realm to use if no realm present in username
<code>`sasl_mechanisms = `</code>	string value	Space separated list of acceptable SASL mechanisms
<b>server_request_prefix = exclusive</b>	string value	address prefix used when sending to a specific server
<b>ssl = False</b>	boolean value	Attempt to connect via SSL. If no other ssl-related parameters are given, it will use the system's CA-bundle to verify the server's certificate.
<code>`ssl_ca_file = `</code>	string value	CA certificate PEM file used to verify the server's certificate
<code>`ssl_cert_file = `</code>	string value	Self-identifying certificate PEM file for client authentication
<code>`ssl_key_file = `</code>	string value	Private key PEM file used to sign <code>ssl_cert_file</code> certificate (optional)
<b>ssl_key_password = None</b>	string value	Password for decrypting <code>ssl_key_file</code> (if encrypted)

Configuration option = Default value	Type	Description
<b>ssl_verify_vhost = False</b>	boolean value	By default SSL checks that the name in the server's certificate matches the hostname in the transport_url. In some configurations it may be preferable to use the virtual hostname instead, for example if the server uses the Server Name Indication TLS extension (rfc6066) to provide a certificate per virtual host. Set ssl_verify_vhost to True if the server's SSL certificate uses the virtual host name instead of the DNS name.
<b>trace = False</b>	boolean value	Debug: dump AMQP frames to stdout
<b>unicast_address = unicast</b>	string value	Appended to the address prefix when sending to a particular RPC/Notification server. Used by the message bus to identify messages sent to a single destination.
<code>`username = `</code>	string value	User name for message broker authentication

### 5.1.10. oslo\_messaging\_kafka

The following table outlines the options available under the **[oslo\_messaging\_kafka]** group in the `/etc/glance/glance-api.conf` file.

Table 5.10. oslo\_messaging\_kafka

Configuration option = Default value	Type	Description
<b>conn_pool_min_size = 2</b>	integer value	The pool size limit for connections expiration policy
<b>conn_pool_ttl = 1200</b>	integer value	The time-to-live in sec of idle connections in the pool
<b>consumer_group = oslo_messaging_consumer</b>	string value	Group id for Kafka consumer. Consumers in one group will coordinate message consumption
<b>kafka_consumer_timeout = 1.0</b>	floating point value	Default timeout(s) for Kafka consumers
<b>kafka_default_host = localhost</b>	string value	Default Kafka broker Host
<b>kafka_default_port = 9092</b>	port value	Default Kafka broker Port
<b>kafka_max_fetch_bytes = 1048576</b>	integer value	Max fetch bytes of Kafka consumer

Configuration option = Default value	Type	Description
<b>pool_size = 10</b>	integer value	Pool Size for Kafka Consumers
<b>producer_batch_size = 16384</b>	integer value	Size of batch for the producer async send
<b>producer_batch_timeout = 0.0</b>	floating point value	Upper bound on the delay for KafkaProducer batching in seconds

### 5.1.11. oslo\_messaging\_notifications

The following table outlines the options available under the **[oslo\_messaging\_notifications]** group in the `/etc/glance/glance-api.conf` file.

Table 5.11. oslo\_messaging\_notifications

Configuration option = Default value	Type	Description
<b>driver = []</b>	multi valued	The Drivers(s) to handle sending notifications. Possible values are messaging, messagingv2, routing, log, test, noop
<b>retry = -1</b>	integer value	The maximum number of attempts to re-send a notification message which failed to be delivered due to a recoverable error. 0 - No retry, -1 - indefinite
<b>topics = [u'notifications']</b>	list value	AMQP topic used for OpenStack notifications.
<b>transport_url = None</b>	string value	A URL representing the messaging driver to use for notifications. If not set, we fall back to the same configuration used for RPC.

### 5.1.12. oslo\_messaging\_rabbit

The following table outlines the options available under the **[oslo\_messaging\_rabbit]** group in the `/etc/glance/glance-api.conf` file.

Table 5.12. oslo\_messaging\_rabbit

Configuration option = Default value	Type	Description
<b>amqp_auto_delete = False</b>	boolean value	Auto-delete queues in AMQP.
<b>amqp_durable_queues = False</b>	boolean value	Use durable queues in AMQP.

Configuration option = Default value	Type	Description
<b>channel_max = None</b>	integer value	Maximum number of channels to allow
<b>connection_factory = single</b>	string value	Connection factory implementation
<b>default_notification_exchange = \${control_exchange}_notification</b>	string value	Exchange name for sending notifications
<b>default_notification_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending notification, -1 means infinite retry.
<b>default_rpc_exchange = \${control_exchange}_rpc</b>	string value	Exchange name for sending RPC messages
<b>default_rpc_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending RPC message, -1 means infinite retry. If actual retry attempts in not 0 the rpc request could be processed more than one time
<b>default_serializer_type = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>fake_rabbit = False</b>	boolean value	Deprecated, use <code>rpc_backend=kombu+memory</code> or <code>rpc_backend=fake</code>
<b>frame_max = None</b>	integer value	The maximum byte size for an AMQP frame
<b>heartbeat_interval = 3</b>	integer value	How often to send heartbeats for consumer's connections
<b>heartbeat_rate = 2</b>	integer value	How often times during the <code>heartbeat_timeout_threshold</code> we check the heartbeat.
<b>heartbeat_timeout_threshold = 60</b>	integer value	Number of seconds after which the Rabbit broker is considered down if heartbeat's keep-alive fails (0 disable the heartbeat). EXPERIMENTAL
<b>host_connection_reconnect_delay = 0.25</b>	floating point value	Set delay for reconnection to some host which has connection error

Configuration option = Default value	Type	Description
<b>kombu_compression = None</b>	string value	EXPERIMENTAL: Possible values are: gzip, bz2. If not set compression will not be used. This option may not be available in future versions.
<b>kombu_failover_strategy = round-robin</b>	string value	Determines how the next RabbitMQ node is chosen in case the one we are currently connected to becomes unavailable. Takes effect only if more than one RabbitMQ node is provided in config.
<b>kombu_missing_consumer_retry_timeout = 60</b>	integer value	How long to wait a missing client before abandoning to send it its replies. This value should not be longer than <code>rpc_response_timeout</code> .
<b>kombu_reconnect_delay = 1.0</b>	floating point value	How long to wait before reconnecting in response to an AMQP consumer cancel notification.
<b>notification_listener_prefetch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to notification listener.
<b>notification_persistence = False</b>	boolean value	Persist notification messages.
<b>notification_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending notification message
<b>pool_max_overflow = 0</b>	integer value	Maximum number of connections to create above <b>pool_max_size</b> .
<b>pool_max_size = 30</b>	integer value	Maximum number of connections to keep queued.
<b>pool_recycle = 600</b>	integer value	Lifetime of a connection (since creation) in seconds or None for no recycling. Expired connections are closed on acquire.
<b>pool_stale = 60</b>	integer value	Threshold at which inactive (since release) connections are considered stale in seconds or None for no staleness. Stale connections are closed on acquire.
<b>pool_timeout = 30</b>	integer value	Default number of seconds to wait for a connections to available

Configuration option = Default value	Type	Description
<b>rabbit_ha_queues = False</b>	boolean value	Try to use HA queues in RabbitMQ (x-ha-policy: all). If you change this option, you must wipe the RabbitMQ database. In RabbitMQ 3.0, queue mirroring is no longer controlled by the x-ha-policy argument when declaring a queue. If you just want to make sure that all queues (except those with auto-generated names) are mirrored across all nodes, run: <code>"rabbitmqctl set_policy HA ^(?!amq\.).* {"ha-mode": "all"}"</code>
<b>rabbit_host = localhost</b>	string value	The RabbitMQ broker address where a single node is used.
<b>rabbit_hosts =</b> <b>[u'\$rabbit_host:\$rabbit_p</b> <b>ort']</b>	list value	RabbitMQ HA cluster host:port pairs.
<b>rabbit_interval_max = 30</b>	integer value	Maximum interval of RabbitMQ connection retries. Default is 30 seconds.
<b>rabbit_login_method =</b> <b>AMQPLAIN</b>	string value	The RabbitMQ login method.
<b>rabbit_max_retries = 0</b>	integer value	Maximum number of RabbitMQ connection retries. Default is 0 (infinite retry count).
<b>rabbit_password = guest</b>	string value	The RabbitMQ password.
<b>rabbit_port = 5672</b>	port value	The RabbitMQ broker port where a single node is used.
<b>rabbit_qos_prefetch_cou</b> <b>nt = 0</b>	integer value	Specifies the number of messages to prefetch. Setting to zero allows unlimited messages.
<b>rabbit_retry_backoff = 2</b>	integer value	How long to backoff for between retries when connecting to RabbitMQ.
<b>rabbit_retry_interval = 1</b>	integer value	How frequently to retry connecting with RabbitMQ.
<b>rabbit_transient_queues_</b> <b>ttl = 1800</b>	integer value	Positive integer representing duration in seconds for queue TTL (x-expires). Queues which are unused for the duration of the TTL are automatically deleted. The parameter affects only reply and fanout queues.
<b>rabbit_userid = guest</b>	string value	The RabbitMQ userid.
<b>rabbit_virtual_host = /</b>	string value	The RabbitMQ virtual host.

Configuration option = Default value	Type	Description
<b>rpc_listener_prefetch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to rpc listener.
<b>rpc_queue_expiration = 60</b>	integer value	Time to live for rpc queues without consumers in seconds.
<b>rpc_reply_exchange = \${control_exchange}_rpc_reply</b>	string value	Exchange name for receiving RPC replies
<b>rpc_reply_listener_prefetch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to rpc reply listener.
<b>rpc_reply_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending reply. -1 means infinite retry during rpc_timeout
<b>rpc_reply_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending reply.
<b>rpc_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending RPC message
<b>socket_timeout = 0.25</b>	floating point value	Set socket timeout in seconds for connection's socket
<b>ssl = None</b>	boolean value	Enable SSL
<code>`ssl_ca_file = `</code>	string value	SSL certification authority file (valid only if SSL enabled).
<code>`ssl_cert_file = `</code>	string value	SSL cert file (valid only if SSL enabled).
<code>`ssl_key_file = `</code>	string value	SSL key file (valid only if SSL enabled).
<b>ssl_options = None</b>	dict value	Arguments passed to ssl.wrap_socket
<code>`ssl_version = `</code>	string value	SSL version to use (valid only if SSL enabled). Valid values are TLSv1 and SSLv23. SSLv2, SSLv3, TLSv1_1, and TLSv1_2 may be available on some distributions.
<b>tcp_user_timeout = 0.25</b>	floating point value	Set TCP_USER_TIMEOUT in seconds for connection's socket

### 5.1.13. oslo\_messaging\_zmq



The following table outlines the options available under the `[oslo_messaging_zmq]` group in the `/etc/glance/glance-api.conf` file.

Table 5.13. `oslo_messaging_zmq`

Configuration option = Default value	Type	Description
<code>rpc_ack_timeout_base = 15</code>	integer value	Number of seconds to wait for an ack from a cast/call. After each retry attempt this timeout is multiplied by some specified multiplier.
<code>rpc_ack_timeout_multiplier = 2</code>	integer value	Number to multiply base ack timeout by after each retry attempt.
<code>rpc_message_ttl = 300</code>	integer value	Expiration timeout in seconds of a sent/received message after which it is not tracked anymore by a client/server.
<code>rpc_poll_timeout = 1</code>	integer value	The default number of seconds that poll should wait. Poll raises timeout exception when timeout expired.
<code>rpc_retry_attempts = 3</code>	integer value	Default number of message sending attempts in case of any problems occurred: positive value N means at most N retries, 0 means no retries, None or -1 (or any other negative values) mean to retry forever. This option is used only if acknowledgments are enabled.
<code>rpc_thread_pool_size = 100</code>	integer value	Maximum number of (green) threads to work concurrently.
<code>rpc_use_acks = False</code>	boolean value	Wait for message acknowledgements from receivers. This mechanism works only via proxy without PUB/SUB.
<code>rpc_zmq_bind_address = *</code>	string value	ZeroMQ bind address. Should be a wildcard (*), an ethernet interface, or IP. The "host" option should point or resolve to this address.
<code>rpc_zmq_bind_port_retries = 100</code>	integer value	Number of retries to find free port number before fail with ZMQBindError.
<code>rpc_zmq_contexts = 1</code>	integer value	Number of ZeroMQ contexts, defaults to 1.
<code>rpc_zmq_host = ccso-sp-undercloud.usersys.redhat.com</code>	string value	Name of this node. Must be a valid hostname, FQDN, or IP address. Must match "host" option, if running Nova.
<code>rpc_zmq_ipc_dir = /var/run/openstack</code>	string value	Directory for holding IPC sockets.

Configuration option = Default value	Type	Description
<b>rpc_zmq_matchmaker = redis</b>	string value	MatchMaker driver.
<b>rpc_zmq_max_port = 65536</b>	integer value	Maximal port number for random ports range.
<b>rpc_zmq_min_port = 49153</b>	port value	Minimal port number for random ports range.
<b>rpc_zmq_serialization = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>rpc_zmq_topic_backlog = None</b>	integer value	Maximum number of ingress messages to locally buffer per topic. Default is unlimited.
<b>subscribe_on = []</b>	list value	List of publisher hosts SubConsumer can subscribe on. This option has higher priority then the default publishers list taken from the matchmaker.
<b>use_dynamic_connections = False</b>	boolean value	This option makes direct connections dynamic or static. It makes sense only with <code>use_router_proxy=False</code> which means to use direct connections for direct message types (ignored otherwise).
<b>use_pub_sub = False</b>	boolean value	Use PUB/SUB pattern for fanout methods. PUB/SUB always uses proxy.
<b>use_router_proxy = False</b>	boolean value	Use ROUTER remote proxy.
<b>zmq_failover_connections = 2</b>	integer value	How many additional connections to a host will be made for failover reasons. This option is actual only in dynamic connections mode.
<b>zmq_immediate = True</b>	boolean value	This option configures round-robin mode in zmq socket. True means not keeping a queue when server side disconnects. False means to keep queue and messages even if server is disconnected, when the server appears we send all accumulated messages to it.

Configuration option = Default value	Type	Description
<b>zmq_linger = -1</b>	integer value	Number of seconds to wait before all pending messages will be sent after closing a socket. The default value of -1 specifies an infinite linger period. The value of 0 specifies no linger period. Pending messages shall be discarded immediately when the socket is closed. Positive values specify an upper bound for the linger period.
<b>zmq_target_expire = 300</b>	integer value	Expiration timeout in seconds of a name service record about existing target (< 0 means no timeout).
<b>zmq_target_update = 180</b>	integer value	Update period in seconds of a name service record about existing target.
<b>zmq_tcp_keepalive = -1</b>	integer value	Enable/disable TCP keepalive (KA) mechanism. The default value of -1 (or any other negative value) means to skip any overrides and leave it to OS default; 0 and 1 (or any other positive value) mean to disable and enable the option respectively.
<b>zmq_tcp_keepalive_cnt = -1</b>	integer value	The number of retransmissions to be carried out before declaring that remote end is not available. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_idle = -1</b>	integer value	The duration between two keepalive transmissions in idle condition. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_intvl = -1</b>	integer value	The duration between two successive keepalive retransmissions, if acknowledgement to the previous keepalive transmission is not received. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.

#### 5.1.14. oslo\_middlewares

The following table outlines the options available under the **[oslo\_middlewares]** group in the **/etc/glance/glance-api.conf** file.

Table 5.14. oslo\_middlewares

Configuration option = Default value	Type	Description
<b>enable_proxy_headers_parsing = False</b>	boolean value	Whether the application is behind a proxy or not. This determines if the middleware should parse the headers or not.

### 5.1.15. oslo\_policy

The following table outlines the options available under the **[oslo\_policy]** group in the `/etc/glance/glance-api.conf` file.

Table 5.15. oslo\_policy

Configuration option = Default value	Type	Description
<b>enforce_scope = False</b>	boolean value	This option controls whether or not to enforce scope when evaluating policies. If <b>True</b> , the scope of the token used in the request is compared to the <b>scope_types</b> of the policy being enforced. If the scopes do not match, an <b>InvalidScope</b> exception will be raised. If <b>False</b> , a message will be logged informing operators that policies are being invoked with mismatching scope.
<b>policy_default_rule = default</b>	string value	Default rule. Enforced when a requested rule is not found.
<b>policy_dirs = [u'policy.d']</b>	multi valued	Directories where policy configuration files are stored. They can be relative to any directory in the search path defined by the <code>config_dir</code> option, or absolute paths. The file defined by <code>policy_file</code> must exist for these directories to be searched. Missing or empty directories are ignored.
<b>policy_file = policy.json</b>	string value	The file that defines policies.
<b>remote_content_type = application/x-www-form-urlencoded</b>	string value	Content Type to send and receive data for REST based policy check
<b>remote_ssl_ca_cert_file = None</b>	string value	Absolute path to ca cert file for REST based policy check
<b>remote_ssl_client_cert_file = None</b>	string value	Absolute path to client cert for REST based policy check
<b>remote_ssl_client_key_file = None</b>	string value	Absolute path client key file REST based policy check

Configuration option = Default value	Type	Description
<b>remote_ssl_verify_server _cert = False</b>	boolean value	server identity verification for REST based policy check

### 5.1.16. paste\_deploy

The following table outlines the options available under the **[paste\_deploy]** group in the **/etc/glance/glance-api.conf** file.

Table 5.16. paste\_deploy

Configuration option = Default value	Type	Description
<b>config_file = None</b>	string value	<p>Name of the paste configuration file.</p> <p>Provide a string value representing the name of the paste configuration file to use for configuring pipelines for server application deployments.</p> <p>NOTES:</p> <ul style="list-style-type: none"> <li>● Provide the name or the path relative to the glance directory for the paste configuration file and not the absolute path.</li> <li>● The sample paste configuration file shipped with Glance need not be edited in most cases as it comes with ready-made pipelines for all common deployment flavors.</li> </ul> <p>If no value is specified for this option, the <b>paste.ini</b> file with the prefix of the corresponding Glance service's configuration file name will be searched for in the known configuration directories. (For example, if this option is missing from or has no value set in <b>glance-api.conf</b>, the service will look for a file named <b>glance-api-paste.ini</b>.) If the paste configuration file is not found, the service will not start.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A string value representing the name of the paste configuration file.</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● flavor</li> </ul>

Configuration option = Default value	Type	Description
<b>flavor = None</b>	string value	<p>Deployment flavor to use in the server application pipeline.</p> <p>Provide a string value representing the appropriate deployment flavor used in the server application pipeline. This is typically the partial name of a pipeline in the paste configuration file with the service name removed.</p> <p>For example, if your paste section name in the paste configuration file is [pipeline:glance-api-keystone], set <b>flavor</b> to <b>keystone</b>.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● String value representing a partial pipeline name.</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● config_file</li> </ul>

### 5.1.17. profiler

The following table outlines the options available under the **[profiler]** group in the **/etc/glance/glance-api.conf** file.

Table 5.17. profiler

Configuration option = Default value	Type	Description
<b>connection_string = messaging://</b>	string value	<p>Connection string for a notifier backend. Default value is messaging:// which sets the notifier to oslo_messaging.</p> <p>Examples of possible values:</p> <ul style="list-style-type: none"> <li>● messaging://: use oslo_messaging driver for sending notifications.</li> <li>● mongodb://127.0.0.1:27017 : use mongodb driver for sending notifications.</li> <li>● elasticsearch://127.0.0.1:9200 : use elasticsearch driver for sending notifications.</li> </ul>

Configuration option = Default value	Type	Description
<b>enabled = False</b>	boolean value	<p>Enables the profiling for all services on this node. Default value is False (fully disable the profiling feature).</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• True: Enables the feature</li> <li>• False: Disables the feature. The profiling cannot be started via this project operations. If the profiling is triggered by another project, this project part will be empty.</li> </ul>
<b>es_doc_type = notification</b>	string value	Document type for notification indexing in elasticsearch.
<b>es_scroll_size = 10000</b>	integer value	Elasticsearch splits large requests in batches. This parameter defines maximum size of each batch (for example: es_scroll_size=10000).
<b>es_scroll_time = 2m</b>	string value	This parameter is a time value parameter (for example: es_scroll_time=2m), indicating for how long the nodes that participate in the search will maintain relevant resources in order to continue and support it.
<b>hmac_keys = SECRET_KEY</b>	string value	<p>Secret key(s) to use for encrypting context data for performance profiling. This string value should have the following format: &lt;key1&gt;[,&lt;key2&gt;,...&lt;keyn&gt;], where each key is some random string. A user who triggers the profiling via the REST API has to set one of these keys in the headers of the REST API call to include profiling results of this node for this particular project.</p> <p>Both "enabled" flag and "hmac_keys" config options should be set to enable profiling. Also, to generate correct profiling information across all services at least one key needs to be consistent between OpenStack projects. This ensures it can be used from client side to generate the trace, containing information from all possible resources.</p>
<b>sentinel_service_name = mymaster</b>	string value	Redis sentinel uses a service name to identify a master redis service. This parameter defines the name (for example: sentinel_service_name=mymaster).

Configuration option = Default value	Type	Description
<b>socket_timeout = 0.1</b>	floating point value	Redis Sentinel provides a timeout option on the connections. This parameter defines that timeout (for example: socket_timeout=0.1).
<b>trace_sqlalchemy = False</b>	boolean value	Enables SQL requests profiling in services. Default value is False (SQL requests won't be traced).  Possible values: <ul style="list-style-type: none"> <li>• True: Enables SQL requests profiling. Each SQL query will be part of the trace and can be analyzed by how much time was spent for that.</li> <li>• False: Disables SQL requests profiling. The spent time is only shown on a higher level of operations. Single SQL queries cannot be analyzed this way.</li> </ul>


### 5.1.18. store\_type\_location\_strategy

The following table outlines the options available under the **[store\_type\_location\_strategy]** group in the **/etc/glance/glance-api.conf** file.

Table 5.18. store\_type\_location\_strategy

Configuration option = Default value	Type	Description
---	------	-------------




Configuration option = Default value	Type	Description
<b>store_type_preference =</b> <b>[]</b>	list value	<p>Preference order of storage backends.</p> <p>Provide a comma separated list of store names in the order in which images should be retrieved from storage backends. These store names must be registered with the <b>stores</b> configuration option.</p> <div data-bbox="815 510 922 763" style="border: 1px solid black; padding: 5px; width: fit-content;">  </div> <p><b>NOTE</b></p> <p>The <b>store_type_preference</b> configuration option is applied only if <b>store_type</b> is chosen as a value for the <b>location_strategy</b> configuration option. An empty list will not change the location order.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Empty list</li> <li>● Comma separated list of registered store names. Legal values are: <ul style="list-style-type: none"> <li>○ file</li> <li>○ http</li> <li>○ rbd</li> <li>○ swift</li> <li>○ sheepdog</li> <li>○ cinder</li> <li>○ vmware</li> </ul> </li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● location_strategy</li> <li>● stores</li> </ul>

### 5.1.19. task

The following table outlines the options available under the **[task]** group in the **/etc/glance/glance-api.conf** file.

Table 5.19. task

Configuration option = Default value	Type	Description
<b>task_executor = taskflow</b>	string value	<p>Task executor to be used to run task scripts.</p> <p>Provide a string value representing the executor to use for task executions. By default, <b>TaskFlow</b> executor is used.</p> <p><b>TaskFlow</b> helps make task executions easy, consistent, scalable and reliable. It also enables creation of lightweight task objects and/or functions that are combined together into flows in a declarative manner.</p> <p>Possible values:</p> <ul style="list-style-type: none"><li>● taskflow</li></ul> <p>Related Options:</p> <ul style="list-style-type: none"><li>● None</li></ul>
<b>task_time_to_live = 48</b>	integer value	Time in hours for which a task lives after, either succeeding or failing

Configuration option = Default value	Type	Description
<b>work_dir = None</b>	string value	<p>Absolute path to the work directory to use for asynchronous task operations.</p> <p>The directory set here will be used to operate over images - normally before they are imported in the destination store.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p>When providing a value for <b>work_dir</b>, please make sure that enough space is provided for concurrent tasks to run efficiently without running out of space.</p> <p>A rough estimation can be done by multiplying the number of <b>max_workers</b> with an average image size (e.g 500MB). The image size estimation should be done based on the average size in your deployment. Note that depending on the tasks running you may need to multiply this number by some factor depending on what the task does. For example, you may want to double the available size if image conversion is enabled. All this being said, remember these are just estimations and you should do them based on the worst case scenario and be prepared to act in case they were wrong.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● String value representing the absolute path to the working directory</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul> </div> </div>

### 5.1.20. taskflow\_executor

The following table outlines the options available under the **[taskflow\_executor]** group in the `/etc/glance/glance-api.conf` file.

Table 5.20. taskflow\_executor

Configuration option = Default value	Type	Description
---	------	-------------

Configuration option = Default value	Type	Description
<b>conversion_format = None</b>	string value	<p>Set the desired image conversion format.</p> <p>Provide a valid image format to which you want images to be converted before they are stored for consumption by Glance. Appropriate image format conversions are desirable for specific storage backends in order to facilitate efficient handling of bandwidth and usage of the storage infrastructure.</p> <p>By default, <b>conversion_format</b> is not set and must be set explicitly in the configuration file.</p> <p>The allowed values for this option are <b>raw</b>, <b>qcow2</b> and <b>vmdk</b>. The <b>raw</b> format is the unstructured disk format and should be chosen when RBD or Ceph storage backends are used for image storage. <b>qcow2</b> is supported by the QEMU emulator that expands dynamically and supports Copy on Write. The <b>vmdk</b> is another common disk format supported by many common virtual machine monitors like VMWare Workstation.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● qcow2</li> <li>● raw</li> <li>● vmdk</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● disk_formats</li> </ul>

Configuration option = Default value	Type	Description
<b>engine_mode = parallel</b>	string value	<p>Set the taskflow engine mode.</p> <p>Provide a string type value to set the mode in which the taskflow engine would schedule tasks to the workers on the hosts. Based on this mode, the engine executes tasks either in single or multiple threads. The possible values for this configuration option are: <b>serial</b> and <b>parallel</b>. When set to <b>serial</b>, the engine runs all the tasks in a single thread which results in serial execution of tasks. Setting this to <b>parallel</b> makes the engine run tasks in multiple threads. This results in parallel execution of tasks.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● serial</li> <li>● parallel</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● max_workers</li> </ul>
<b>max_workers = 10</b>	integer value	<p>Set the number of engine executable tasks.</p> <p>Provide an integer value to limit the number of workers that can be instantiated on the hosts. In other words, this number defines the number of parallel tasks that can be executed at the same time by the taskflow engine. This value can be greater than one when the engine mode is set to parallel.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Integer value greater than or equal to 1</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● engine_mode</li> </ul>

## 5.2. GLANCE-REGISTRY.CONF


This section contains options for the `/etc/glance/glance-registry.conf` file.

### 5.2.1. DEFAULT

The following table outlines the options available under the **[DEFAULT]** group in the `/etc/glance/glance-registry.conf` file.

Table 5.21. DEFAULT

Configuration option = Default value	Type	Description
<b>admin_role = admin</b>	string value	<p>Role used to identify an authenticated user as administrator.</p> <p>Provide a string value representing a Keystone role to identify an administrative user. Users with this role will be granted administrative privileges. The default value for this option is <i>admin</i>.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A string value which is a valid Keystone role</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>allow_additional_image_properties = True</b>	boolean value	<p>Allow users to add additional/custom properties to images.</p> <p>Glance defines a standard set of properties (in its schema) that appear on every image. These properties are also known as <b>base properties</b>. In addition to these properties, Glance allows users to add custom properties to images. These are known as <b>additional properties</b>.</p> <p>By default, this configuration option is set to <b>True</b> and users are allowed to add additional properties. The number of additional properties that can be added to an image can be controlled via <b>image_property_quota</b> configuration option.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● image_property_quota</li> </ul>


Configuration option = Default value	Type	Description
<b>allow_anonymous_access = False</b>	boolean value	<p>Allow limited access to unauthenticated users.</p> <p>Assign a boolean to determine API access for unauthenticated users. When set to False, the API cannot be accessed by unauthenticated users. When set to True, unauthenticated users can access the API with read-only privileges. This however only applies when using ContextMiddleware.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>api_limit_max = 1000</b>	integer value	<p>Maximum number of results that could be returned by a request.</p> <p>As described in the help text of <b>limit_param_default</b>, some requests may return multiple results. The number of results to be returned are governed either by the <b>limit</b> parameter in the request or the <b>limit_param_default</b> configuration option. The value in either case, can't be greater than the absolute maximum defined by this configuration option. Anything greater than this value is trimmed down to the maximum value defined here.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p>Setting this to a very large value may slow down database queries and increase response times. Setting this to a very low value may result in poor user experience.</p> </div> </div> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any positive integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● limit_param_default</li> </ul>

Configuration option = Default value	Type	Description
<b>backlog = 4096</b>	integer value	<p>Set the number of incoming connection requests.</p> <p>Provide a positive integer value to limit the number of requests in the backlog queue. The default queue size is 4096.</p> <p>An incoming connection to a TCP listener socket is queued before a connection can be established with the server. Setting the backlog for a TCP socket ensures a limited queue size for incoming traffic.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Positive integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<b>bind_host = 0.0.0.0</b>	unknown value	<p>IP address to bind the glance servers to.</p> <p>Provide an IP address to bind the glance server to. The default value is <b>0.0.0.0</b>.</p> <p>Edit this option to enable the server to listen on one particular IP address on the network card. This facilitates selection of a particular network interface for the server.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A valid IPv4 address</li> <li>• A valid IPv6 address</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>




Configuration option = Default value	Type	Description
<b>bind_port = None</b>	port value	<p>Port number on which the server will listen.</p> <p>Provide a valid port number to bind the server's socket to. This port is then set to identify processes and forward network messages that arrive at the server. The default bind_port value for the API server is 9292 and for the registry server is 9191.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid port number (0 to 65535)</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>ca_file = None</b>	string value	<p>Absolute path to the CA file.</p> <p>Provide a string value representing a valid absolute path to the Certificate Authority file to use for client authentication.</p> <p>A CA file typically contains necessary trusted certificates to use for the client authentication. This is essential to ensure that a secure connection is established to the server via the internet.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Valid absolute path to the CA file</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>

Configuration option = Default value	Type	Description
<b>cert_file = None</b>	string value	<p>Absolute path to the certificate file.</p> <p>Provide a string value representing a valid absolute path to the certificate file which is required to start the API service securely.</p> <p>A certificate file typically is a public key container and includes the server's public key, server name, server information and the signature which was a result of the verification process using the CA certificate. This is required for a secure connection establishment.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Valid absolute path to the certificate file</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>None</li> </ul>
<b>client_socket_timeout = 900</b>	integer value	<p>Timeout for client connections' socket operations.</p> <p>Provide a valid integer value representing time in seconds to set the period of wait before an incoming connection can be closed. The default value is 900 seconds.</p> <p>The value zero implies wait forever.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Zero</li> <li>Positive integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>None</li> </ul>
<b>conn_pool_min_size = 2</b>	integer value	The pool size limit for connections expiration policy
<b>conn_pool_ttl = 1200</b>	integer value	The time-to-live in sec of idle connections in the pool
<b>control_exchange = openstack</b>	string value	The default exchange under which topics are scoped. May be overridden by an exchange name specified in the transport_url option.

Configuration option = Default value	Type	Description
<b>data_api = glance.db.sqlalchemy.api</b>	string value	<p>Python module path of data access API.</p> <p>Specifies the path to the API to use for accessing the data model. This option determines how the image catalog data will be accessed.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● <code>glance.db.sqlalchemy.api</code></li> <li>● <code>glance.db.registry.api</code></li> <li>● <code>glance.db.simple.api</code></li> </ul> <p>If this option is set to <b>glance.db.sqlalchemy.api</b> then the image catalog data is stored in and read from the database via the SQLAlchemy Core and ORM APIs.</p> <p>Setting this option to <b>glance.db.registry.api</b> will force all database access requests to be routed through the Registry service. This avoids data access from the Glance API nodes for an added layer of security, scalability and manageability.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p>In v2 OpenStack Images API, the registry service is optional. In order to use the Registry API in v2, the option <b>enable_v2_registry</b> must be set to <b>True</b>.</p> </div> </div> <p>Finally, when this configuration option is set to <b>glance.db.simple.api</b>, image catalog data is stored in and read from an in-memory data structure. This is primarily used for testing.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>enable_v2_api</code></li> <li>● <code>enable_v2_registry</code></li> </ul>
<b>debug = False</b>	boolean value	<p>If set to true, the logging level will be set to DEBUG instead of the default INFO level.</p>

Configuration option = Default value	Type	Description
<b>default_log_levels =</b> <b>[u'amqp=WARN',</b> <b>u'amqplib=WARN',</b> <b>u'boto=WARN',</b> <b>u'qpid=WARN',</b> <b>u'sqlalchemy=WARN',</b> <b>u'suds=INFO',</b> <b>u'oslo.messaging=INFO',</b> <b>u'oslo_messaging=INFO',</b> <b>u'iso8601=WARN',</b> <b>u'requests.packages.urlli</b> <b>b3.connectionpool=WAR</b> <b>N',</b> <b>u'urllib3.connectionpool=</b> <b>WARN',</b> <b>u'websocket=WARN',</b> <b>u'requests.packages.urlli</b> <b>b3.util.retry=WARN',</b> <b>u'urllib3.util.retry=WARN',</b> <b>u'keystonemiddleware=W</b> <b>ARN',</b> <b>u'routes.middleware=WA</b> <b>RN', u'stevedore=WARN',</b> <b>u'taskflow=WARN',</b> <b>u'keystoneauth=WARN',</b> <b>u'oslo.cache=INFO',</b> <b>u'dogpile.core.dogpile=IN</b> <b>FO']</b>	list value	List of package logging levels in logger=LEVEL pairs. This option is ignored if log_config_append is set.

Configuration option = Default value	Type	Description
<b>digest_algorithm = sha256</b>	string value	<p>Digest algorithm to use for digital signature.</p> <p>Provide a string value representing the digest algorithm to use for generating digital signatures. By default, <b>sha256</b> is used.</p> <p>To get a list of the available algorithms supported by the version of OpenSSL on your platform, run the command: <b>openssl list-message-digest-algorithms</b>. Examples are <i>sha1</i>, <i>sha256</i>, and <i>sha512</i>.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p><b>digest_algorithm</b> is not related to Glance's image signing and verification. It is only used to sign the universally unique identifier (UUID) as a part of the certificate file and key file validation.</p> </div> </div> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● An OpenSSL message digest algorithm identifier</li> </ul> <p>Relation options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>enable_image_import = True</b>	boolean value	<p>Enables the Image Import workflow introduced in Pike</p> <p>As <i>[DEFAULT]/node_staging_uri</i> is required for the Image Import, it's disabled per default in Pike, enabled per default in Queens and removed in Rocky. This allows Glance to operate with previous version configs upon upgrade.</p> <p>Setting this option to False will disable the endpoints related to Image Import Refactoring work.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <i>[DEFAULT]/node_staging_uri</i></li> </ul>

Configuration option = Default value	Type	Description
<b>enable_v1_api = True</b>	boolean value	<p>Deploy the v1 OpenStack Images API.</p> <p>When this option is set to <b>True</b>, Glance service will respond to requests on registered endpoints conforming to the v1 OpenStack Images API.</p> <p>NOTES:</p> <ul style="list-style-type: none"> <li>● If this option is enabled, then <b>enable_v1_registry</b> must also be set to <b>True</b> to enable mandatory usage of Registry service with v1 API.</li> <li>● If this option is disabled, then the <b>enable_v1_registry</b> option, which is enabled by default, is also recommended to be disabled.</li> <li>● This option is separate from <b>enable_v2_api</b>, both v1 and v2 OpenStack Images API can be deployed independent of each other.</li> <li>● If deploying only the v2 Images API, this option, which is enabled by default, should be disabled.</li> </ul> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● enable_v1_registry</li> <li>● enable_v2_api</li> </ul>


Configuration option = Default value	Type	Description
<b>enable_v1_registry = True</b>	boolean value	<p>Deploy the v1 API Registry service.</p> <p>When this option is set to <b>True</b>, the Registry service will be enabled in Glance for v1 API requests.</p> <p>NOTES:</p> <ul style="list-style-type: none"><li>● Use of Registry is mandatory in v1 API, so this option must be set to <b>True</b> if the <b>enable_v1_api</b> option is enabled.</li><li>● If deploying only the v2 OpenStack Images API, this option, which is enabled by default, should be disabled.</li></ul> <p>Possible values:</p> <ul style="list-style-type: none"><li>● True</li><li>● False</li></ul> <p>Related options:</p> <ul style="list-style-type: none"><li>● enable_v1_api</li></ul>

Configuration option = Default value	Type	Description
<b>enable_v2_api = True</b>	boolean value	<p>Deploy the v2 OpenStack Images API.</p> <p>When this option is set to <b>True</b>, Glance service will respond to requests on registered endpoints conforming to the v2 OpenStack Images API.</p> <p>NOTES:</p> <ul style="list-style-type: none"><li>● If this option is disabled, then the <b>enable_v2_registry</b> option, which is enabled by default, is also recommended to be disabled.</li><li>● This option is separate from <b>enable_v1_api</b>, both v1 and v2 OpenStack Images API can be deployed independent of each other.</li><li>● If deploying only the v1 Images API, this option, which is enabled by default, should be disabled.</li></ul> <p>Possible values:</p> <ul style="list-style-type: none"><li>● True</li><li>● False</li></ul> <p>Related options:</p> <ul style="list-style-type: none"><li>● enable_v2_registry</li><li>● enable_v1_api</li></ul>



Configuration option = Default value	Type	Description
<b>enable_v2_registry = True</b>	boolean value	<p>Deploy the v2 API Registry service.</p> <p>When this option is set to <b>True</b>, the Registry service will be enabled in Glance for v2 API requests.</p> <p>NOTES:</p> <ul style="list-style-type: none"> <li>● Use of Registry is optional in v2 API, so this option must only be enabled if both <b>enable_v2_api</b> is set to <b>True</b> and the <b>data_api</b> option is set to <b>glance.db.registry.api</b>.</li> <li>● If deploying only the v1 OpenStack Images API, this option, which is enabled by default, should be disabled.</li> </ul> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● enable_v2_api</li> <li>● data_api</li> </ul>
<b>enabled_import_methods = [u'glance-direct', u'web-download']</b>	list value	<p>List of enabled Image Import Methods</p> <p>Both <i>glance-direct</i> and <i>web-download</i> are enabled by default.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● [DEFAULT]/node_staging_uri</li> <li>● [DEFAULT]/enable_image_import</li> </ul>
<b>executor_thread_pool_size = 64</b>	integer value	Size of executor thread pool when executor is threading or eventlet.
<b>fatal_deprecations = False</b>	boolean value	Enables or disables fatal status of deprecations.

Configuration option = Default value	Type	Description
<b>http_keepalive = True</b>	boolean value	<p>Set keep alive option for HTTP over TCP.</p> <p>Provide a boolean value to determine sending of keep alive packets. If set to <b>False</b>, the server returns the header "Connection: close". If set to <b>True</b>, the server returns a "Connection: Keep-Alive" in its responses. This enables retention of the same TCP connection for HTTP conversations instead of opening a new one with each new request.</p> <p>This option must be set to <b>False</b> if the client socket connection needs to be closed explicitly after the response is received and read successfully by the client.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• True</li> <li>• False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<b>image_location_quota = 10</b>	integer value	<p>Maximum number of locations allowed on an image.</p> <p>Any negative value is interpreted as unlimited.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<b>image_member_quota = 128</b>	integer value	<p>Maximum number of image members per image.</p> <p>This limits the maximum of users an image can be shared with. Any negative value is interpreted as unlimited.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>


Configuration option = Default value	Type	Description
<b>image_property_quota = 128</b>	integer value	<p>Maximum number of properties allowed on an image.</p> <p>This enforces an upper limit on the number of additional properties an image can have. Any negative value is interpreted as unlimited.</p> <div data-bbox="815 472 922 696" style="display: inline-block; vertical-align: top;">  </div> <p><b>NOTE</b></p> <p>This won't have any impact if additional properties are disabled. Please refer to <b>allow_additional_image_properties</b>.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>allow_additional_image_properties</b></li> </ul>
<b>image_size_cap = 1099511627776</b>	integer value	<p>Maximum size of image a user can upload in bytes.</p> <p>An image upload greater than the size mentioned here would result in an image creation failure. This configuration option defaults to 1099511627776 bytes (1 TiB).</p> <p>NOTES:</p> <ul style="list-style-type: none"> <li>● This value should only be increased after careful consideration and must be set less than or equal to 8 EiB (9223372036854775808).</li> <li>● This value must be set with careful consideration of the backend storage capacity. Setting this to a very low value may result in a large number of image failures. And, setting this to a very large value may result in faster consumption of storage. Hence, this must be set according to the nature of images created and storage capacity available.</li> </ul> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any positive number less than or equal to 9223372036854775808</li> </ul>

Configuration option = Default value	Type	Description
<b>image_tag_quota = 128</b>	integer value	<p>Maximum number of tags allowed on an image.</p> <p>Any negative value is interpreted as unlimited.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<code>`instance_format = [instance: %(uuid)s] `</code>	string value	The format for an instance that is passed with the log message.
<code>`instance_uuid_format = [instance:%(uuid)s] `</code>	string value	The format for an instance UUID that is passed with the log message.
<b>key_file = None</b>	string value	<p>Absolute path to a private key file.</p> <p>Provide a string value representing a valid absolute path to a private key file which is required to establish the client-server connection.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Absolute path to the private key file</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>




Configuration option = Default value	Type	Description
<b>limit_param_default = 25</b>	integer value	<p>The default number of results to return for a request.</p> <p>Responses to certain API requests, like list images, may return multiple items. The number of results returned can be explicitly controlled by specifying the <b>limit</b> parameter in the API request. However, if a <b>limit</b> parameter is not specified, this configuration value will be used as the default number of results to be returned for any API request.</p> <p>NOTES:</p> <ul style="list-style-type: none"> <li>• The value of this configuration option may not be greater than the value specified by <b>api_limit_max</b>.</li> <li>• Setting this to a very large value may slow down database queries and increase response times. Setting this to a very low value may result in poor user experience.</li> </ul> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Any positive integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <code>api_limit_max</code></li> </ul>
<b>log-config-append = None</b>	string value	<p>The name of a logging configuration file. This file is appended to any existing logging configuration files. For details about logging configuration files, see the Python logging module documentation. Note that when logging configuration files are used then all logging configuration is set in the configuration file and other logging configuration options are ignored (for example, <code>logging_context_format_string</code>).</p>
<b>log-date-format = %Y-%m-%d %H:%M:%S</b>	string value	<p>Defines the format string for <code>%(asctime)s</code> in log records. Default: <code>%(default)s</code>. This option is ignored if <code>log_config_append</code> is set.</p>
<b>log-dir = None</b>	string value	<p>(Optional) The base directory used for relative <code>log_file</code> paths. This option is ignored if <code>log_config_append</code> is set.</p>
<b>log-file = None</b>	string value	<p>(Optional) Name of log file to send logging output to. If no default is set, logging will go to <code>stderr</code> as defined by <code>use_stderr</code>. This option is ignored if <code>log_config_append</code> is set.</p>

Configuration option = Default value	Type	Description
<b>logging_context_format_string</b> = <code>%(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [%(request_id)s %(user_identity)s] %(instance)s%(message)s</code>	string value	Format string to use for log messages with context.
<b>logging_debug_format_s uffix</b> = <code>%(funcName)s %(pathname)s:%(lineno)d</code>	string value	Additional data to append to log message when logging level for the message is DEBUG.
<b>logging_default_format_s tring</b> = <code>%(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [-] %(instance)s%(message)s</code>	string value	Format string to use for log messages when context is undefined.
<b>logging_exception_prefix</b> = <code>%(asctime)s.%(msecs)03d %(process)d ERROR %(name)s %(instance)s</code>	string value	Prefix each line of exception output with this format.
<b>logging_user_identity_for mat</b> = <code>%(user)s %(tenant)s %(domain)s %(user_domain)s %(project_domain)s</code>	string value	Defines the format string for <code>%(user_identity)s</code> that is used in <code>logging_context_format_string</code> .

Configuration option = Default value	Type	Description
<b>max_header_line = 16384</b>	integer value	<p>Maximum line size of message headers.</p> <p>Provide an integer value representing a length to limit the size of message headers. The default value is 16384.</p> <div data-bbox="815 472 922 792" style="float: left; margin-right: 10px;"> </div> <p><b>NOTE</b></p> <p><b>max_header_line</b> may need to be increased when using large tokens (typically those generated by the Keystone v3 API with big service catalogs). However, it is to be kept in mind that larger values for <b>max_header_line</b> would flood the logs.</p> <p>Setting <b>max_header_line</b> to 0 sets no limit for the line size of message headers.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● 0</li> <li>● Positive integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>max_request_id_length = 64</b>	integer value	<p>Limit the request ID length.</p> <p>Provide an integer value to limit the length of the request ID to the specified length. The default value is 64. Users can change this to any ineteger value between 0 and 16384 however keeping in mind that a larger value may flood the logs.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Integer value between 0 and 16384</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>

Configuration option = Default value	Type	Description
<b>metadata_encryption_key</b> = None	string value	<p>AES key for encrypting store location metadata.</p> <p>Provide a string value representing the AES cipher to use for encrypting Glance store metadata.</p> <p> <b>NOTE</b></p> <p>The AES key to use must be set to a random string of length 16, 24 or 32 bytes.</p> <p>Possible values:</p> <ul style="list-style-type: none"><li>• String value representing a valid AES key</li></ul> <p>Related options:</p> <ul style="list-style-type: none"><li>• None</li></ul>



Configuration option = Default value	Type	Description
<b>node_staging_uri = file:///tmp/staging/</b>	string value	<p>The URL provides location where the temporary data will be stored</p> <p>This option is for Glance internal use only. Glance will save the image data uploaded by the user to <i>staging</i> endpoint during the image import process.</p> <p>This option does not change the <i>staging</i> API endpoint by any means.</p> <p> <b>NOTE</b></p> <p>It is discouraged to use same path as [task]/work_dir</p> <p> <b>NOTE</b></p> <p><i>file://&lt;absolute-directory-path&gt;</i> is the only option api_image_import flow will support for now.</p> <p> <b>NOTE</b></p> <p>The staging path must be on shared filesystem available to all Glance API nodes.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● String starting with <i>file://</i> followed by absolute FS path</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● [task]/work_dir</li> <li>● [DEFAULT]/enable_image_import (<b>deprecated</b>)</li> </ul>

Configuration option = Default value	Type	Description
<b>owner_is_tenant = True</b>	boolean value	<p>Set the image owner to tenant or the authenticated user.</p> <p>Assign a boolean value to determine the owner of an image. When set to True, the owner of the image is the tenant. When set to False, the owner of the image will be the authenticated user issuing the request. Setting it to False makes the image private to the associated user and sharing with other users within the same tenant (or "project") requires explicit image sharing via image membership.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• True</li> <li>• False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<b>publish_errors = False</b>	boolean value	<p>Enables or disables publication of error events.</p>
<b>pydev_worker_debug_host = None</b>	unknown value	<p>Host address of the pydev server.</p> <p>Provide a string value representing the hostname or IP of the pydev server to use for debugging. The pydev server listens for debug connections on this address, facilitating remote debugging in Glance.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Valid hostname</li> <li>• Valid IP address</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>

Configuration option = Default value	Type	Description
<b>pydev_worker_debug_port = 5678</b>	port value	<p>Port number that the pydev server will listen on.</p> <p>Provide a port number to bind the pydev server to. The pydev process accepts debug connections on this port and facilitates remote debugging in Glance.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A valid port number</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<b>rate_limit_burst = 0</b>	integer value	Maximum number of logged messages per <code>rate_limit_interval</code> .
<b>rate_limit_except_level = CRITICAL</b>	string value	Log level name used by rate limiting: CRITICAL, ERROR, INFO, WARNING, DEBUG or empty string. Logs with level greater or equal to <code>rate_limit_except_level</code> are not filtered. An empty string means that all levels are filtered.
<b>rate_limit_interval = 0</b>	integer value	Interval, number of seconds, of log rate limiting.
<b>rpc_ack_timeout_base = 15</b>	integer value	Number of seconds to wait for an ack from a cast/call. After each retry attempt this timeout is multiplied by some specified multiplier.
<b>rpc_ack_timeout_multiplier = 2</b>	integer value	Number to multiply base ack timeout by after each retry attempt.
<b>rpc_backend = rabbit</b>	string value	The messaging driver to use, defaults to rabbit. Other drivers include amqp and zmq.
<b>rpc_conn_pool_size = 30</b>	integer value	Size of RPC connection pool.
<b>rpc_message_ttl = 300</b>	integer value	Expiration timeout in seconds of a sent/received message after which it is not tracked anymore by a client/server.
<b>rpc_poll_timeout = 1</b>	integer value	The default number of seconds that poll should wait. Poll raises timeout exception when timeout expired.
<b>rpc_response_timeout = 60</b>	integer value	Seconds to wait for a response from a call.

Configuration option = Default value	Type	Description
<b>rpc_retry_attempts = 3</b>	integer value	Default number of message sending attempts in case of any problems occurred: positive value N means at most N retries, 0 means no retries, None or -1 (or any other negative values) mean to retry forever. This option is used only if acknowledgments are enabled.
<b>rpc_thread_pool_size = 100</b>	integer value	Maximum number of (green) threads to work concurrently.
<b>rpc_use_acks = False</b>	boolean value	Wait for message acknowledgements from receivers. This mechanism works only via proxy without PUB/SUB.
<b>rpc_zmq_bind_address = *</b>	string value	ZeroMQ bind address. Should be a wildcard (*), an ethernet interface, or IP. The "host" option should point or resolve to this address.
<b>rpc_zmq_bind_port_retries = 100</b>	integer value	Number of retries to find free port number before fail with ZMQBindError.
<b>rpc_zmq_contexts = 1</b>	integer value	Number of ZeroMQ contexts, defaults to 1.
<b>rpc_zmq_host = ccosp-undercloud.usersys.redhat.com</b>	string value	Name of this node. Must be a valid hostname, FQDN, or IP address. Must match "host" option, if running Nova.
<b>rpc_zmq_ipc_dir = /var/run/openstack</b>	string value	Directory for holding IPC sockets.
<b>rpc_zmq_matchmaker = redis</b>	string value	MatchMaker driver.
<b>rpc_zmq_max_port = 65536</b>	integer value	Maximal port number for random ports range.
<b>rpc_zmq_min_port = 49153</b>	port value	Minimal port number for random ports range.
<b>rpc_zmq_serialization = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>rpc_zmq_topic_backlog = None</b>	integer value	Maximum number of ingress messages to locally buffer per topic. Default is unlimited.


Configuration option = Default value	Type	Description
<b>secure_proxy_ssl_header = None</b>	string value	The HTTP header used to determine the scheme for the original request, even if it was removed by an SSL terminating proxy. Typical value is "HTTP_X_FORWARDED_PROTO".
<b>show_image_direct_url = False</b>	boolean value	<p>Show direct image location when returning an image.</p> <p>This configuration option indicates whether to show the direct image location when returning image details to the user. The direct image location is where the image data is stored in backend storage. This image location is shown under the image property <b>direct_url</b>.</p> <p>When multiple image locations exist for an image, the best location is displayed based on the location strategy indicated by the configuration option <b>location_strategy</b>.</p> <p>NOTES:</p> <ul style="list-style-type: none"> <li>● Revealing image locations can present a GRAVE SECURITY RISK as image locations can sometimes include credentials. Hence, this is set to <b>False</b> by default. Set this to <b>True</b> with EXTREME CAUTION and ONLY IF you know what you are doing!</li> <li>● If an operator wishes to avoid showing any image location(s) to the user, then both this option and <b>show_multiple_locations</b> MUST be set to <b>False</b>.</li> </ul> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● show_multiple_locations</li> <li>● location_strategy</li> </ul>

Configuration option = Default value	Type	Description
<b>show_multiple_locations</b> = <b>False</b>	boolean value	<p>Show all image locations when returning an image.</p> <p>This configuration option indicates whether to show all the image locations when returning image details to the user. When multiple image locations exist for an image, the locations are ordered based on the location strategy indicated by the configuration opt <b>location_strategy</b>. The image locations are shown under the image property <b>locations</b>.</p> <p>NOTES:</p> <ul style="list-style-type: none"> <li>● Revealing image locations can present a GRAVE SECURITY RISK as image locations can sometimes include credentials. Hence, this is set to <b>False</b> by default. Set this to <b>True</b> with EXTREME CAUTION and ONLY IF you know what you are doing!</li> <li>● See <a href="https://wiki.openstack.org/wiki/OSSN/OSSN-0065">https://wiki.openstack.org/wiki/OSSN/OSSN-0065</a> for more information.</li> <li>● If an operator wishes to avoid showing any image location(s) to the user, then both this option and <b>show_image_direct_url</b> MUST be set to <b>False</b>.</li> </ul> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● show_image_direct_url</li> <li>● location_strategy</li> </ul>
<b>subscribe_on</b> = []	list value	List of publisher hosts SubConsumer can subscribe on. This option has higher priority then the default publishers list taken from the matchmaker.
<b>syslog-log-facility</b> = <b>LOG_USER</b>	string value	Syslog facility to receive log lines. This option is ignored if log_config_append is set.

Configuration option = Default value	Type	Description
<b>tcp_keepidle = 600</b>	integer value	<p>Set the wait time before a connection recheck.</p> <p>Provide a positive integer value representing time in seconds which is set as the idle wait time before a TCP keep alive packet can be sent to the host. The default value is 600 seconds.</p> <p>Setting <b>tcp_keepidle</b> helps verify at regular intervals that a connection is intact and prevents frequent TCP connection reestablishment.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Positive integer value representing time in seconds</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>transport_url = None</b>	string value	<p>The network address and optional user credentials for connecting to the messaging backend, in URL format. The expected format is:</p> <pre>driver://[user:pass@]host:port[, [userN:passN@]hostN:portN]/virtual_host?query</pre> <p>Example: rabbit://rabbitmq:password@127.0.0.1:5672//</p> <p>For full details on the fields in the URL see the documentation of <code>oslo_messaging.TransportURL</code> at <a href="https://docs.openstack.org/oslo.messaging/latest/reference/transport.html">https://docs.openstack.org/oslo.messaging/latest/reference/transport.html</a></p>
<b>use-journal = False</b>	boolean value	<p>Enable journald for logging. If running in a systemd environment you may wish to enable journal support. Doing so will use the journal native protocol which includes structured metadata in addition to log messages. This option is ignored if <code>log_config_append</code> is set.</p>
<b>use-json = False</b>	boolean value	<p>Use JSON formatting for logging. This option is ignored if <code>log_config_append</code> is set.</p>
<b>use-syslog = False</b>	boolean value	<p>Use syslog for logging. Existing syslog format is DEPRECATED and will be changed later to honor RFC5424. This option is ignored if <code>log_config_append</code> is set.</p>

Configuration option = Default value	Type	Description
<b>use_dynamic_connections = False</b>	boolean value	This option makes direct connections dynamic or static. It makes sense only with <code>use_router_proxy=False</code> which means to use direct connections for direct message types (ignored otherwise).
<b>use_pub_sub = False</b>	boolean value	Use PUB/SUB pattern for fanout methods. PUB/SUB always uses proxy.
<b>use_router_proxy = False</b>	boolean value	Use ROUTER remote proxy.
<b>use_stderr = False</b>	boolean value	Log output to standard error. This option is ignored if <code>log_config_append</code> is set.
<b>user_storage_quota = 0</b>	string value	<p>Maximum amount of image storage per tenant.</p> <p>This enforces an upper limit on the cumulative storage consumed by all images of a tenant across all stores. This is a per-tenant limit.</p> <p>The default unit for this configuration option is Bytes. However, storage units can be specified using case-sensitive literals <b>B</b>, <b>KB</b>, <b>MB</b>, <b>GB</b> and <b>TB</b> representing Bytes, KiloBytes, MegaBytes, GigaBytes and TeraBytes respectively. Note that there should not be any space between the value and unit. Value <b>0</b> signifies no quota enforcement. Negative values are invalid and result in errors.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A string that is a valid concatenation of a non-negative integer representing the storage value and an optional string literal representing storage units as mentioned above.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<b>watch-log-file = False</b>	boolean value	Uses logging handler designed to watch file system. When log file is moved or removed this handler will open a new log file with specified path instantaneously. It makes sense only if <code>log_file</code> option is specified and Linux platform is used. This option is ignored if <code>log_config_append</code> is set.



Configuration option = Default value	Type	Description
<b>workers = None</b>	integer value	<p>Number of Glance worker processes to start.</p> <p>Provide a non-negative integer value to set the number of child process workers to service requests. By default, the number of CPUs available is set as the value for <b>workers</b> limited to 8. For example if the processor count is 6, 6 workers will be used, if the processor count is 24 only 8 workers will be used. The limit will only apply to the default value, if 24 workers is configured, 24 is used.</p> <p>Each worker process is made to listen on the port set in the configuration file and contains a greenthread pool of size 1000.</p> <div style="display: flex; align-items: flex-start;">  <div> <p><b>NOTE</b></p> <p>Setting the number of workers to zero, triggers the creation of a single API process with a greenthread pool of size 1000.</p> </div> </div> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● 0</li> <li>● Positive integer value (typically equal to the number of CPUs)</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>zmq_failover_connections = 2</b>	integer value	How many additional connections to a host will be made for failover reasons. This option is actual only in dynamic connections mode.
<b>zmq_immediate = True</b>	boolean value	This option configures round-robin mode in zmq socket. True means not keeping a queue when server side disconnects. False means to keep queue and messages even if server is disconnected, when the server appears we send all accumulated messages to it.

Configuration option = Default value	Type	Description
<b>zmq_linger = -1</b>	integer value	Number of seconds to wait before all pending messages will be sent after closing a socket. The default value of -1 specifies an infinite linger period. The value of 0 specifies no linger period. Pending messages shall be discarded immediately when the socket is closed. Positive values specify an upper bound for the linger period.
<b>zmq_target_expire = 300</b>	integer value	Expiration timeout in seconds of a name service record about existing target (< 0 means no timeout).
<b>zmq_target_update = 180</b>	integer value	Update period in seconds of a name service record about existing target.
<b>zmq_tcp_keepalive = -1</b>	integer value	Enable/disable TCP keepalive (KA) mechanism. The default value of -1 (or any other negative value) means to skip any overrides and leave it to OS default; 0 and 1 (or any other positive value) mean to disable and enable the option respectively.
<b>zmq_tcp_keepalive_cnt = -1</b>	integer value	The number of retransmissions to be carried out before declaring that remote end is not available. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_idle = -1</b>	integer value	The duration between two keepalive transmissions in idle condition. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_intvl = -1</b>	integer value	The duration between two successive keepalive retransmissions, if acknowledgement to the previous keepalive transmission is not received. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.

### 5.2.2. database

The following table outlines the options available under the **[database]** group in the **/etc/glance/glance-registry.conf** file.

Table 5.22. database

Configuration option = Default value	Type	Description
<b>backend = sqlalchemy</b>	string value	The back end to use for the database.
<b>connection = None</b>	string value	The SQLAlchemy connection string to use to connect to the database.
<b>connection_debug = 0</b>	integer value	Verbosity of SQL debugging information: 0=None, 100=Everything.
<b>connection_recycle_time = 3600</b>	integer value	Connections which have been present in the connection pool longer than this number of seconds will be replaced with a new one the next time they are checked out from the pool.
<b>connection_trace = False</b>	boolean value	Add Python stack traces to SQL as comment strings.
<b>db_inc_retry_interval = True</b>	boolean value	If True, increases the interval between retries of a database operation up to db_max_retry_interval.
<b>db_max_retries = 20</b>	integer value	Maximum retries in case of connection error or deadlock error before error is raised. Set to -1 to specify an infinite retry count.
<b>db_max_retry_interval = 10</b>	integer value	If db_inc_retry_interval is set, the maximum seconds between retries of a database operation.
<b>db_retry_interval = 1</b>	integer value	Seconds between retries of a database transaction.
<b>max_overflow = 50</b>	integer value	If set, use this value for max_overflow with SQLAlchemy.
<b>max_pool_size = 5</b>	integer value	Maximum number of SQL connections to keep open in a pool. Setting a value of 0 indicates no limit.
<b>max_retries = 10</b>	integer value	Maximum number of database connection retries during startup. Set to -1 to specify an infinite retry count.
<b>min_pool_size = 1</b>	integer value	Minimum number of SQL connections to keep open in a pool.
<b>mysql_enable_ndb = False</b>	boolean value	If True, transparently enables support for handling MySQL Cluster (NDB).

Configuration option = Default value	Type	Description
<b>mysql_sql_mode = TRADITIONAL</b>	string value	The SQL mode to be used for MySQL sessions. This option, including the default, overrides any server-set SQL mode. To use whatever SQL mode is set by the server configuration, set this to no value. Example: <code>mysql_sql_mode=</code>
<b>pool_timeout = None</b>	integer value	If set, use this value for <code>pool_timeout</code> with SQLAlchemy.
<b>retry_interval = 10</b>	integer value	Interval between retries of opening a SQL connection.
<b>slave_connection = None</b>	string value	The SQLAlchemy connection string to use to connect to the slave database.
<b>sqlite_synchronous = True</b>	boolean value	If True, SQLite uses synchronous mode.
<b>use_db_reconnect = False</b>	boolean value	Enable the experimental use of database reconnect on connection lost.
<b>use_tpool = False</b>	boolean value	Enable the experimental use of thread pooling for all DB API calls

### 5.2.3. keystone\_authtoken

The following table outlines the options available under the **[keystone\_authtoken]** group in the `/etc/glance/glance-registry.conf` file.

Table 5.23. keystone\_authtoken

Configuration option = Default value	Type	Description
<b>auth_section = None</b>	string value	Config Section from which to load plugin specific options
<b>auth_type = None</b>	string value	Authentication type to load

Configuration option = Default value	Type	Description
<b>auth_uri = None</b>	string value	Complete "public" Identity API endpoint. This endpoint should not be an "admin" endpoint, as it should be accessible by all end users. Unauthenticated clients are redirected to this endpoint to authenticate. Although this endpoint should ideally be unversioned, client support in the wild varies. If you're using a versioned v2 endpoint here, then this should <b>not</b> be the same endpoint the service user utilizes for validating tokens, because normal end users may not be able to reach that endpoint. This option is deprecated in favor of <code>www_authenticate_uri</code> and will be removed in the S release.
<b>auth_version = None</b>	string value	API version of the admin Identity API endpoint.
<b>cache = None</b>	string value	Request environment key where the Swift cache object is stored. When <code>auth_token</code> middleware is deployed with a Swift cache, use this option to have the middleware share a caching backend with swift. Otherwise, use the <b>memcached_servers</b> option instead.
<b>cafile = None</b>	string value	A PEM encoded Certificate Authority to use when verifying HTTPs connections. Defaults to system CAs.
<b>certfile = None</b>	string value	Required if identity server requires client certificate
<b>check_revocations_for_cached = False</b>	boolean value	If true, the revocation list will be checked for cached tokens. This requires that PKI tokens are configured on the identity server.
<b>delay_auth_decision = False</b>	boolean value	Do not handle authorization requests within the middleware, but delegate the authorization decision to downstream WSGI components.
<b>enforce_token_bind = permissive</b>	string value	Used to control the use and type of token binding. Can be set to: "disabled" to not check token binding. "permissive" (default) to validate binding information if the bind type is of a form known to the server and ignore it if not. "strict" like "permissive" but if the bind type is unknown the token will be rejected. "required" any form of token binding is needed to be allowed. Finally the name of a binding method that must be present in tokens.

Configuration option = Default value	Type	Description
<b>hash_algorithms = [u'md5']</b>	list value	Hash algorithms to use for hashing PKI tokens. This may be a single algorithm or multiple. The algorithms are those supported by Python standard <code>hashlib.new()</code> . The hashes will be tried in the order given, so put the preferred one first for performance. The result of the first hash will be stored in the cache. This will typically be set to multiple values only while migrating from a less secure algorithm to a more secure one. Once all the old tokens are expired this option should be set to a single value for better performance.
<b>http_connect_timeout = None</b>	integer value	Request timeout value for communicating with Identity API server.
<b>http_request_max_retries = 3</b>	integer value	How many times are we trying to reconnect when communicating with Identity API Server.
<b>include_service_catalog = True</b>	boolean value	(Optional) Indicate whether to set the X-Service-Catalog header. If False, middleware will not ask for service catalog on token validation and will not set the X-Service-Catalog header.
<b>insecure = False</b>	boolean value	Verify HTTPS connections.
<b>keyfile = None</b>	string value	Required if identity server requires client certificate
<b>memcache_pool_conn_get_timeout = 10</b>	integer value	(Optional) Number of seconds that an operation will wait to get a memcached client connection from the pool.
<b>memcache_pool_dead_retry = 300</b>	integer value	(Optional) Number of seconds memcached server is considered dead before it is tried again.
<b>memcache_pool_maxsize = 10</b>	integer value	(Optional) Maximum total number of open connections to every memcached server.
<b>memcache_pool_socket_timeout = 3</b>	integer value	(Optional) Socket timeout in seconds for communicating with a memcached server.
<b>memcache_pool_unused_timeout = 60</b>	integer value	(Optional) Number of seconds a connection to memcached is held unused in the pool before it is closed.
<b>memcache_secret_key = None</b>	string value	(Optional, mandatory if <code>memcache_security_strategy</code> is defined) This string is used for key derivation.

Configuration option = Default value	Type	Description
<b>memcache_security_strategy = None</b>	string value	(Optional) If defined, indicate whether token data should be authenticated or authenticated and encrypted. If MAC, token data is authenticated (with HMAC) in the cache. If ENCRYPT, token data is encrypted and authenticated in the cache. If the value is not one of these options or empty, auth_token will raise an exception on initialization.
<b>memcache_use_advanced_pool = False</b>	boolean value	(Optional) Use the advanced (eventlet safe) memcached client pool. The advanced pool will only work under python 2.x.
<b>memcached_servers = None</b>	list value	Optionally specify a list of memcached server(s) to use for caching. If left undefined, tokens will instead be cached in-process.
<b>region_name = None</b>	string value	The region in which the identity server can be found.
<b>revocation_cache_time = 10</b>	integer value	Determines the frequency at which the list of revoked tokens is retrieved from the Identity service (in seconds). A high number of revocation events combined with a low cache duration may significantly reduce performance. Only valid for PKI tokens. This option has been deprecated in the Ocata release and will be removed in the P release.
<b>service_token_roles = [u'service']</b>	list value	A choice of roles that must be present in a service token. Service tokens are allowed to request that an expired token can be used and so this check should tightly control that only actual services should be sending this token. Roles here are applied as an ANY check so any role in this list must be present. For backwards compatibility reasons this currently only affects the allow_expired check.
<b>service_token_roles_required = False</b>	boolean value	For backwards compatibility reasons we must let valid service tokens pass that don't pass the service_token_roles check as valid. Setting this true will become the default in a future release and should be enabled if possible.
<b>signing_dir = None</b>	string value	Directory used to cache files related to PKI tokens. This option has been deprecated in the Ocata release and will be removed in the P release.

Configuration option = Default value	Type	Description
<b>token_cache_time = 300</b>	integer value	In order to prevent excessive effort spent validating tokens, the middleware caches previously-seen tokens for a configurable duration (in seconds). Set to -1 to disable caching completely.
<b>www_authenticate_uri = None</b>	string value	Complete "public" Identity API endpoint. This endpoint should not be an "admin" endpoint, as it should be accessible by all end users. Unauthenticated clients are redirected to this endpoint to authenticate. Although this endpoint should ideally be unversioned, client support in the wild varies. If you're using a versioned v2 endpoint here, then this should <b>not</b> be the same endpoint the service user utilizes for validating tokens, because normal end users may not be able to reach that endpoint.

#### 5.2.4. matchmaker\_redis

The following table outlines the options available under the **[matchmaker\_redis]** group in the **/etc/glance/glance-registry.conf** file.

Table 5.24. matchmaker\_redis

Configuration option = Default value	Type	Description
<b>check_timeout = 20000</b>	integer value	Time in ms to wait before the transaction is killed.
<b>host = 127.0.0.1</b>	string value	Host to locate redis.
<code>`password = `</code>	string value	Password for Redis server (optional).
<b>port = 6379</b>	port value	Use this port to connect to redis host.
<b>sentinel_group_name = oslo-messaging-zeromq</b>	string value	Redis replica set name.
<b>sentinel_hosts = []</b>	list value	List of Redis Sentinel hosts (fault tolerance mode), e.g., [host:port, host1:port ... ]
<b>socket_timeout = 10000</b>	integer value	Timeout in ms on blocking socket operations.
<b>wait_timeout = 2000</b>	integer value	Time in ms to wait between connection attempts.

#### 5.2.5. oslo\_messaging\_amqp



The following table outlines the options available under the `[oslo_messaging_amqp]` group in the `/etc/glance/glance-registry.conf` file.

Table 5.25. `oslo_messaging_amqp`

Configuration option = Default value	Type	Description
<b>addressing_mode = dynamic</b>	string value	Indicates the addressing mode used by the driver. Permitted values: <i>legacy</i> - use legacy non-routable addressing <i>routable</i> - use routable addresses <i>dynamic</i> - use legacy addresses if the message bus does not support routing otherwise use routable addressing
<b>allow_insecure_clients = False</b>	boolean value	Accept clients using either SSL or plain TCP
<b>anycast_address = anycast</b>	string value	Appended to the address prefix when sending to a group of consumers. Used by the message bus to identify messages that should be delivered in a round-robin fashion across consumers.
<b>broadcast_prefix = broadcast</b>	string value	address prefix used when broadcasting to all servers
<b>connection_retry_backoff = 2</b>	integer value	Increase the <code>connection_retry_interval</code> by this many seconds after each unsuccessful failover attempt.
<b>connection_retry_interval = 1</b>	integer value	Seconds to pause before attempting to re-connect.
<b>connection_retry_interval _max = 30</b>	integer value	Maximum limit for <code>connection_retry_interval</code> + <code>connection_retry_backoff</code>
<b>container_name = None</b>	string value	Name for the AMQP container. must be globally unique. Defaults to a generated UUID
<b>default_notification_exch ange = None</b>	string value	Exchange name used in notification addresses. Exchange name resolution precedence: Target.exchange if set else <code>default_notification_exchange</code> if set else <code>control_exchange</code> if set else <i>notify</i>
<b>default_notify_timeout = 30</b>	integer value	The deadline for a sent notification message delivery. Only used when caller does not provide a timeout expiry.
<b>default_reply_retry = 0</b>	integer value	The maximum number of attempts to re-send a reply message which failed due to a recoverable error.

Configuration option = Default value	Type	Description
<b>default_reply_timeout = 30</b>	integer value	The deadline for an rpc reply message delivery.
<b>default_rpc_exchange = None</b>	string value	Exchange name used in RPC addresses. Exchange name resolution precedence: Target.exchange if set else default_rpc_exchange if set else control_exchange if set else <i>rpc</i>
<b>default_send_timeout = 30</b>	integer value	The deadline for an rpc cast or call message delivery. Only used when caller does not provide a timeout expiry.
<b>default_sender_link_timeout = 600</b>	integer value	The duration to schedule a purge of idle sender links. Detach link after expiry.
<b>group_request_prefix = unicast</b>	string value	address prefix when sending to any server in group
<b>idle_timeout = 0</b>	integer value	Timeout for inactive connections (in seconds)
<b>link_retry_delay = 10</b>	integer value	Time to pause between re-connecting an AMQP 1.0 link that failed due to a recoverable error.
<b>multicast_address = multicast</b>	string value	Appended to the address prefix when sending a fanout message. Used by the message bus to identify fanout messages.
<b>notify_address_prefix = openstack.org/om/notify</b>	string value	Address prefix for all generated Notification addresses
<b>notify_server_credit = 100</b>	integer value	Window size for incoming Notification messages
<b>password = `</b>	string value	Password for message broker authentication
<b>pre_settled = [u'rpc-cast', u'rpc-reply']</b>	multi valued	Send messages of this type pre-settled. Pre-settled messages will not receive acknowledgement from the peer. Note well: pre-settled messages may be silently discarded if the delivery fails. Permitted values: <i>rpc-call</i> - send RPC Calls pre-settled <i>rpc-reply</i> - send RPC Replies pre-settled <i>rpc-cast</i> - Send RPC Casts pre-settled <i>notify</i> - Send Notifications pre-settled

Configuration option = Default value	Type	Description
<b>pseudo_vhost = True</b>	boolean value	Enable virtual host support for those message buses that do not natively support virtual hosting (such as qpid). When set to true the virtual host name will be added to all message bus addresses, effectively creating a private <i>subnet</i> per virtual host. Set to False if the message bus supports virtual hosting using the <i>hostname</i> field in the AMQP 1.0 Open performative as the name of the virtual host.
<b>reply_link_credit = 200</b>	integer value	Window size for incoming RPC Reply messages.
<b>rpc_address_prefix = openstack.org/om/rpc</b>	string value	Address prefix for all generated RPC addresses
<b>rpc_server_credit = 100</b>	integer value	Window size for incoming RPC Request messages
<code>`sasldb_dir = `</code>	string value	Path to directory that contains the SASL configuration
<code>`sasldb_name = `</code>	string value	Name of configuration file (without .conf suffix)
<code>`sasldb_realm = `</code>	string value	SASL realm to use if no realm present in username
<code>`sasldb_mechanisms = `</code>	string value	Space separated list of acceptable SASL mechanisms
<b>server_request_prefix = exclusive</b>	string value	address prefix used when sending to a specific server
<b>ssl = False</b>	boolean value	Attempt to connect via SSL. If no other ssl-related parameters are given, it will use the system's CA-bundle to verify the server's certificate.
<code>`ssl_ca_file = `</code>	string value	CA certificate PEM file used to verify the server's certificate
<code>`ssl_cert_file = `</code>	string value	Self-identifying certificate PEM file for client authentication
<code>`ssl_key_file = `</code>	string value	Private key PEM file used to sign <code>ssl_cert_file</code> certificate (optional)
<b>ssl_key_password = None</b>	string value	Password for decrypting <code>ssl_key_file</code> (if encrypted)

Configuration option = Default value	Type	Description
<b>ssl_verify_vhost = False</b>	boolean value	By default SSL checks that the name in the server's certificate matches the hostname in the transport_url. In some configurations it may be preferable to use the virtual hostname instead, for example if the server uses the Server Name Indication TLS extension (rfc6066) to provide a certificate per virtual host. Set ssl_verify_vhost to True if the server's SSL certificate uses the virtual host name instead of the DNS name.
<b>trace = False</b>	boolean value	Debug: dump AMQP frames to stdout
<b>unicast_address = unicast</b>	string value	Appended to the address prefix when sending to a particular RPC/Notification server. Used by the message bus to identify messages sent to a single destination.
<code>`username = `</code>	string value	User name for message broker authentication

### 5.2.6. oslo\_messaging\_kafka

The following table outlines the options available under the **[oslo\_messaging\_kafka]** group in the `/etc/glance/glance-registry.conf` file.

Table 5.26. oslo\_messaging\_kafka

Configuration option = Default value	Type	Description
<b>conn_pool_min_size = 2</b>	integer value	The pool size limit for connections expiration policy
<b>conn_pool_ttl = 1200</b>	integer value	The time-to-live in sec of idle connections in the pool
<b>consumer_group = oslo_messaging_consumer</b>	string value	Group id for Kafka consumer. Consumers in one group will coordinate message consumption
<b>kafka_consumer_timeout = 1.0</b>	floating point value	Default timeout(s) for Kafka consumers
<b>kafka_default_host = localhost</b>	string value	Default Kafka broker Host
<b>kafka_default_port = 9092</b>	port value	Default Kafka broker Port

Configuration option = Default value	Type	Description
<b>kafka_max_fetch_bytes = 1048576</b>	integer value	Max fetch bytes of Kafka consumer
<b>pool_size = 10</b>	integer value	Pool Size for Kafka Consumers
<b>producer_batch_size = 16384</b>	integer value	Size of batch for the producer async send
<b>producer_batch_timeout = 0.0</b>	floating point value	Upper bound on the delay for KafkaProducer batching in seconds

### 5.2.7. oslo\_messaging\_notifications

The following table outlines the options available under the **[oslo\_messaging\_notifications]** group in the **/etc/glance/glance-registry.conf** file.

Table 5.27. oslo\_messaging\_notifications

Configuration option = Default value	Type	Description
<b>driver = []</b>	multi valued	The Drivers(s) to handle sending notifications. Possible values are messaging, messagingv2, routing, log, test, noop
<b>retry = -1</b>	integer value	The maximum number of attempts to re-send a notification message which failed to be delivered due to a recoverable error. 0 - No retry, -1 - indefinite
<b>topics = [u'notifications']</b>	list value	AMQP topic used for OpenStack notifications.
<b>transport_url = None</b>	string value	A URL representing the messaging driver to use for notifications. If not set, we fall back to the same configuration used for RPC.

### 5.2.8. oslo\_messaging\_rabbit

The following table outlines the options available under the **[oslo\_messaging\_rabbit]** group in the **/etc/glance/glance-registry.conf** file.

Table 5.28. oslo\_messaging\_rabbit

Configuration option = Default value	Type	Description
<b>amqp_auto_delete = False</b>	boolean value	Auto-delete queues in AMQP.

Configuration option = Default value	Type	Description
<b>amqp_durable_queues = False</b>	boolean value	Use durable queues in AMQP.
<b>channel_max = None</b>	integer value	Maximum number of channels to allow
<b>connection_factory = single</b>	string value	Connection factory implementation
<b>default_notification_exchange = \${control_exchange}_notification</b>	string value	Exchange name for sending notifications
<b>default_notification_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending notification, -1 means infinite retry.
<b>default_rpc_exchange = \${control_exchange}_rpc</b>	string value	Exchange name for sending RPC messages
<b>default_rpc_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending RPC message, -1 means infinite retry. If actual retry attempts is not 0 the rpc request could be processed more than one time
<b>default_serializer_type = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>fake_rabbit = False</b>	boolean value	Deprecated, use <code>rpc_backend=kombu+memory</code> or <code>rpc_backend=fake</code>
<b>frame_max = None</b>	integer value	The maximum byte size for an AMQP frame
<b>heartbeat_interval = 3</b>	integer value	How often to send heartbeats for consumer's connections
<b>heartbeat_rate = 2</b>	integer value	How often times during the <code>heartbeat_timeout_threshold</code> we check the heartbeat.
<b>heartbeat_timeout_threshold = 60</b>	integer value	Number of seconds after which the Rabbit broker is considered down if heartbeat's keep-alive fails (0 disable the heartbeat). EXPERIMENTAL
<b>host_connection_reconnect_delay = 0.25</b>	floating point value	Set delay for reconnection to some host which has connection error

Configuration option = Default value	Type	Description
<b>kombu_compression = None</b>	string value	EXPERIMENTAL: Possible values are: gzip, bz2. If not set compression will not be used. This option may not be available in future versions.
<b>kombu_failover_strategy = round-robin</b>	string value	Determines how the next RabbitMQ node is chosen in case the one we are currently connected to becomes unavailable. Takes effect only if more than one RabbitMQ node is provided in config.
<b>kombu_missing_consumer_retry_timeout = 60</b>	integer value	How long to wait a missing client before abandoning to send it its replies. This value should not be longer than <code>rpc_response_timeout</code> .
<b>kombu_reconnect_delay = 1.0</b>	floating point value	How long to wait before reconnecting in response to an AMQP consumer cancel notification.
<b>notification_listener_prefetch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to notification listener.
<b>notification_persistence = False</b>	boolean value	Persist notification messages.
<b>notification_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending notification message
<b>pool_max_overflow = 0</b>	integer value	Maximum number of connections to create above <b>pool_max_size</b> .
<b>pool_max_size = 30</b>	integer value	Maximum number of connections to keep queued.
<b>pool_recycle = 600</b>	integer value	Lifetime of a connection (since creation) in seconds or None for no recycling. Expired connections are closed on acquire.
<b>pool_stale = 60</b>	integer value	Threshold at which inactive (since release) connections are considered stale in seconds or None for no staleness. Stale connections are closed on acquire.
<b>pool_timeout = 30</b>	integer value	Default number of seconds to wait for a connections to available

Configuration option = Default value	Type	Description
<b>rabbit_ha_queues = False</b>	boolean value	Try to use HA queues in RabbitMQ (x-ha-policy: all). If you change this option, you must wipe the RabbitMQ database. In RabbitMQ 3.0, queue mirroring is no longer controlled by the x-ha-policy argument when declaring a queue. If you just want to make sure that all queues (except those with auto-generated names) are mirrored across all nodes, run: "rabbitmqctl set_policy HA ^(?!amq\.).* {"ha-mode": "all"}"
<b>rabbit_host = localhost</b>	string value	The RabbitMQ broker address where a single node is used.
<b>rabbit_hosts =</b> [u'\$rabbit_host:\$rabbit_port']	list value	RabbitMQ HA cluster host:port pairs.
<b>rabbit_interval_max = 30</b>	integer value	Maximum interval of RabbitMQ connection retries. Default is 30 seconds.
<b>rabbit_login_method = AMQPLAIN</b>	string value	The RabbitMQ login method.
<b>rabbit_max_retries = 0</b>	integer value	Maximum number of RabbitMQ connection retries. Default is 0 (infinite retry count).
<b>rabbit_password = guest</b>	string value	The RabbitMQ password.
<b>rabbit_port = 5672</b>	port value	The RabbitMQ broker port where a single node is used.
<b>rabbit_qos_prefetch_count = 0</b>	integer value	Specifies the number of messages to prefetch. Setting to zero allows unlimited messages.
<b>rabbit_retry_backoff = 2</b>	integer value	How long to backoff for between retries when connecting to RabbitMQ.
<b>rabbit_retry_interval = 1</b>	integer value	How frequently to retry connecting with RabbitMQ.
<b>rabbit_transient_queues_ttl = 1800</b>	integer value	Positive integer representing duration in seconds for queue TTL (x-expires). Queues which are unused for the duration of the TTL are automatically deleted. The parameter affects only reply and fanout queues.
<b>rabbit_userid = guest</b>	string value	The RabbitMQ userid.



Configuration option = Default value	Type	Description
<b>rabbit_virtual_host = /</b>	string value	The RabbitMQ virtual host.
<b>rpc_listener_prefetch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to rpc listener.
<b>rpc_queue_expiration = 60</b>	integer value	Time to live for rpc queues without consumers in seconds.
<b>rpc_reply_exchange = \${control_exchange}_rpc_reply</b>	string value	Exchange name for receiving RPC replies
<b>rpc_reply_listener_prefetch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to rpc reply listener.
<b>rpc_reply_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending reply. -1 means infinite retry during rpc_timeout
<b>rpc_reply_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending reply.
<b>rpc_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending RPC message
<b>socket_timeout = 0.25</b>	floating point value	Set socket timeout in seconds for connection's socket
<b>ssl = None</b>	boolean value	Enable SSL
<code>`ssl_ca_file = `</code>	string value	SSL certification authority file (valid only if SSL enabled).
<code>`ssl_cert_file = `</code>	string value	SSL cert file (valid only if SSL enabled).
<code>`ssl_key_file = `</code>	string value	SSL key file (valid only if SSL enabled).
<b>ssl_options = None</b>	dict value	Arguments passed to ssl.wrap_socket
<code>`ssl_version = `</code>	string value	SSL version to use (valid only if SSL enabled). Valid values are TLSv1 and SSLv23. SSLv2, SSLv3, TLSv1_1, and TLSv1_2 may be available on some distributions.
<b>tcp_user_timeout = 0.25</b>	floating point value	Set TCP_USER_TIMEOUT in seconds for connection's socket

## 5.2.9. oslo\_messaging\_zmq

The following table outlines the options available under the **[oslo\_messaging\_zmq]** group in the **/etc/glance/glance-registry.conf** file.

Table 5.29. oslo\_messaging\_zmq

Configuration option = Default value	Type	Description
<b>rpc_ack_timeout_base = 15</b>	integer value	Number of seconds to wait for an ack from a cast/call. After each retry attempt this timeout is multiplied by some specified multiplier.
<b>rpc_ack_timeout_multiplier = 2</b>	integer value	Number to multiply base ack timeout by after each retry attempt.
<b>rpc_message_ttl = 300</b>	integer value	Expiration timeout in seconds of a sent/received message after which it is not tracked anymore by a client/server.
<b>rpc_poll_timeout = 1</b>	integer value	The default number of seconds that poll should wait. Poll raises timeout exception when timeout expired.
<b>rpc_retry_attempts = 3</b>	integer value	Default number of message sending attempts in case of any problems occurred: positive value N means at most N retries, 0 means no retries, None or -1 (or any other negative values) mean to retry forever. This option is used only if acknowledgments are enabled.
<b>rpc_thread_pool_size = 100</b>	integer value	Maximum number of (green) threads to work concurrently.
<b>rpc_use_acks = False</b>	boolean value	Wait for message acknowledgements from receivers. This mechanism works only via proxy without PUB/SUB.
<b>rpc_zmq_bind_address = *</b>	string value	ZeroMQ bind address. Should be a wildcard (*), an ethernet interface, or IP. The "host" option should point or resolve to this address.
<b>rpc_zmq_bind_port_retries = 100</b>	integer value	Number of retries to find free port number before fail with ZMQBindError.
<b>rpc_zmq_contexts = 1</b>	integer value	Number of ZeroMQ contexts, defaults to 1.
<b>rpc_zmq_host = ccosp-undercloud.usersys.redhat.com</b>	string value	Name of this node. Must be a valid hostname, FQDN, or IP address. Must match "host" option, if running Nova.

Configuration option = Default value	Type	Description
<b>rpc_zmq_ipc_dir = /var/run/openstack</b>	string value	Directory for holding IPC sockets.
<b>rpc_zmq_matchmaker = redis</b>	string value	MatchMaker driver.
<b>rpc_zmq_max_port = 65536</b>	integer value	Maximal port number for random ports range.
<b>rpc_zmq_min_port = 49153</b>	port value	Minimal port number for random ports range.
<b>rpc_zmq_serialization = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>rpc_zmq_topic_backlog = None</b>	integer value	Maximum number of ingress messages to locally buffer per topic. Default is unlimited.
<b>subscribe_on = []</b>	list value	List of publisher hosts SubConsumer can subscribe on. This option has higher priority then the default publishers list taken from the matchmaker.
<b>use_dynamic_connections = False</b>	boolean value	This option makes direct connections dynamic or static. It makes sense only with <code>use_router_proxy=False</code> which means to use direct connections for direct message types (ignored otherwise).
<b>use_pub_sub = False</b>	boolean value	Use PUB/SUB pattern for fanout methods. PUB/SUB always uses proxy.
<b>use_router_proxy = False</b>	boolean value	Use ROUTER remote proxy.
<b>zmq_failover_connections = 2</b>	integer value	How many additional connections to a host will be made for failover reasons. This option is actual only in dynamic connections mode.
<b>zmq_immediate = True</b>	boolean value	This option configures round-robin mode in zmq socket. True means not keeping a queue when server side disconnects. False means to keep queue and messages even if server is disconnected, when the server appears we send all accumulated messages to it.

Configuration option = Default value	Type	Description
<b>zmq_linger = -1</b>	integer value	Number of seconds to wait before all pending messages will be sent after closing a socket. The default value of -1 specifies an infinite linger period. The value of 0 specifies no linger period. Pending messages shall be discarded immediately when the socket is closed. Positive values specify an upper bound for the linger period.
<b>zmq_target_expire = 300</b>	integer value	Expiration timeout in seconds of a name service record about existing target (< 0 means no timeout).
<b>zmq_target_update = 180</b>	integer value	Update period in seconds of a name service record about existing target.
<b>zmq_tcp_keepalive = -1</b>	integer value	Enable/disable TCP keepalive (KA) mechanism. The default value of -1 (or any other negative value) means to skip any overrides and leave it to OS default; 0 and 1 (or any other positive value) mean to disable and enable the option respectively.
<b>zmq_tcp_keepalive_cnt = -1</b>	integer value	The number of retransmissions to be carried out before declaring that remote end is not available. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_idle = -1</b>	integer value	The duration between two keepalive transmissions in idle condition. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_intvl = -1</b>	integer value	The duration between two successive keepalive retransmissions, if acknowledgement to the previous keepalive transmission is not received. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.

### 5.2.10. oslo\_policy

The following table outlines the options available under the **[oslo\_policy]** group in the **/etc/glance/glance-registry.conf** file.

Table 5.30. oslo\_policy

Configuration option = Default value	Type	Description
<b>enforce_scope = False</b>	boolean value	This option controls whether or not to enforce scope when evaluating policies. If <b>True</b> , the scope of the token used in the request is compared to the <b>scope_types</b> of the policy being enforced. If the scopes do not match, an <b>InvalidScope</b> exception will be raised. If <b>False</b> , a message will be logged informing operators that policies are being invoked with mismatching scope.
<b>policy_default_rule = default</b>	string value	Default rule. Enforced when a requested rule is not found.
<b>policy_dirs = [u'policy.d']</b>	multi valued	Directories where policy configuration files are stored. They can be relative to any directory in the search path defined by the <code>config_dir</code> option, or absolute paths. The file defined by <code>policy_file</code> must exist for these directories to be searched. Missing or empty directories are ignored.
<b>policy_file = policy.json</b>	string value	The file that defines policies.
<b>remote_content_type = application/x-www-form-urlencoded</b>	string value	Content Type to send and receive data for REST based policy check
<b>remote_ssl_ca_cert_file = None</b>	string value	Absolute path to ca cert file for REST based policy check
<b>remote_ssl_client_cert_file = None</b>	string value	Absolute path to client cert for REST based policy check
<b>remote_ssl_client_key_file = None</b>	string value	Absolute path client key file REST based policy check
<b>remote_ssl_verify_server_cert = False</b>	boolean value	server identity verification for REST based policy check

### 5.2.11. paste\_deploy

The following table outlines the options available under the **[paste\_deploy]** group in the `/etc/glance/glance-registry.conf` file.

Table 5.31. paste\_deploy

Configuration option = Default value	Type	Description
<b>config_file = None</b>	string value	<p>Name of the paste configuration file.</p> <p>Provide a string value representing the name of the paste configuration file to use for configuring pipelines for server application deployments.</p> <p>NOTES:</p> <ul style="list-style-type: none"> <li>● Provide the name or the path relative to the glance directory for the paste configuration file and not the absolute path.</li> <li>● The sample paste configuration file shipped with Glance need not be edited in most cases as it comes with ready-made pipelines for all common deployment flavors.</li> </ul> <p>If no value is specified for this option, the <b>paste.ini</b> file with the prefix of the corresponding Glance service's configuration file name will be searched for in the known configuration directories. (For example, if this option is missing from or has no value set in <b>glance-api.conf</b>, the service will look for a file named <b>glance-api-paste.ini</b>.) If the paste configuration file is not found, the service will not start.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A string value representing the name of the paste configuration file.</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● flavor</li> </ul>

Configuration option = Default value	Type	Description
<b>flavor = None</b>	string value	<p>Deployment flavor to use in the server application pipeline.</p> <p>Provide a string value representing the appropriate deployment flavor used in the server application pipeline. This is typically the partial name of a pipeline in the paste configuration file with the service name removed.</p> <p>For example, if your paste section name in the paste configuration file is [pipeline:glance-api-keystone], set <b>flavor</b> to <b>keystone</b>.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● String value representing a partial pipeline name.</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● config_file</li> </ul>

### 5.2.12. profiler

The following table outlines the options available under the **[profiler]** group in the **/etc/glance/glance-registry.conf** file.

Table 5.32. profiler

Configuration option = Default value	Type	Description
<b>connection_string = messaging://</b>	string value	<p>Connection string for a notifier backend. Default value is messaging:// which sets the notifier to oslo_messaging.</p> <p>Examples of possible values:</p> <ul style="list-style-type: none"> <li>● messaging://: use oslo_messaging driver for sending notifications.</li> <li>● mongodb://127.0.0.1:27017 : use mongoddb driver for sending notifications.</li> <li>● elasticsearch://127.0.0.1:9200 : use elasticsearch driver for sending notifications.</li> </ul>

Configuration option = Default value	Type	Description
<b>enabled = False</b>	boolean value	<p>Enables the profiling for all services on this node. Default value is False (fully disable the profiling feature).</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• True: Enables the feature</li> <li>• False: Disables the feature. The profiling cannot be started via this project operations. If the profiling is triggered by another project, this project part will be empty.</li> </ul>
<b>es_doc_type = notification</b>	string value	Document type for notification indexing in elasticsearch.
<b>es_scroll_size = 10000</b>	integer value	Elasticsearch splits large requests in batches. This parameter defines maximum size of each batch (for example: es_scroll_size=10000).
<b>es_scroll_time = 2m</b>	string value	This parameter is a time value parameter (for example: es_scroll_time=2m), indicating for how long the nodes that participate in the search will maintain relevant resources in order to continue and support it.
<b>hmac_keys = SECRET_KEY</b>	string value	<p>Secret key(s) to use for encrypting context data for performance profiling. This string value should have the following format: &lt;key1&gt;[,&lt;key2&gt;,...&lt;keyn&gt;], where each key is some random string. A user who triggers the profiling via the REST API has to set one of these keys in the headers of the REST API call to include profiling results of this node for this particular project.</p> <p>Both "enabled" flag and "hmac_keys" config options should be set to enable profiling. Also, to generate correct profiling information across all services at least one key needs to be consistent between OpenStack projects. This ensures it can be used from client side to generate the trace, containing information from all possible resources.</p>
<b>sentinel_service_name = mymaster</b>	string value	Redis sentinel uses a service name to identify a master redis service. This parameter defines the name (for example: sentinel_service_name=mymaster).



Configuration option = Default value	Type	Description
<b>socket_timeout = 0.1</b>	floating point value	Redis Sentinel provides a timeout option on the connections. This parameter defines that timeout (for example: socket_timeout=0.1).
<b>trace_sqlalchemy = False</b>	boolean value	Enables SQL requests profiling in services. Default value is False (SQL requests won't be traced).  Possible values: <ul style="list-style-type: none"> <li>• True: Enables SQL requests profiling. Each SQL query will be part of the trace and can be analyzed by how much time was spent for that.</li> <li>• False: Disables SQL requests profiling. The spent time is only shown on a higher level of operations. Single SQL queries cannot be analyzed this way.</li> </ul>

## 5.3. GLANCE-SCRUBBER.CONF

This section contains options for the `/etc/glance/glance-scrubber.conf` file.


### 5.3.1. DEFAULT

The following table outlines the options available under the **[DEFAULT]** group in the `/etc/glance/glance-scrubber.conf` file.


Table 5.33. DEFAULT

Configuration option = Default value	Type	Description
---	------	-------------


Configuration option = Default value	Type	Description
<b>allow_additional_image_properties = True</b>	boolean value	<p>Allow users to add additional/custom properties to images.</p> <p>Glance defines a standard set of properties (in its schema) that appear on every image. These properties are also known as <b>base properties</b>. In addition to these properties, Glance allows users to add custom properties to images. These are known as <b>additional properties</b>.</p> <p>By default, this configuration option is set to <b>True</b> and users are allowed to add additional properties. The number of additional properties that can be added to an image can be controlled via <b>image_property_quota</b> configuration option.</p> <p>Possible values:</p> <ul style="list-style-type: none"><li>• True</li><li>• False</li></ul> <p>Related options:</p> <ul style="list-style-type: none"><li>• image_property_quota</li></ul>


Configuration option = Default value	Type	Description
<b>api_limit_max = 1000</b>	integer value	<p>Maximum number of results that could be returned by a request.</p> <p>As described in the help text of <b>limit_param_default</b>, some requests may return multiple results. The number of results to be returned are governed either by the <b>limit</b> parameter in the request or the <b>limit_param_default</b> configuration option. The value in either case, can't be greater than the absolute maximum defined by this configuration option. Anything greater than this value is trimmed down to the maximum value defined here.</p> <div data-bbox="815 757 922 981" style="display: inline-block; vertical-align: top;">  </div> <p><b>NOTE</b></p> <p>Setting this to a very large value may slow down database queries and increase response times. Setting this to a very low value may result in poor user experience.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any positive integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● limit_param_default</li> </ul>

Configuration option = Default value	Type	Description
<b>daemon = False</b>	boolean value	<p>Run scrubber as a daemon.</p> <p>This boolean configuration option indicates whether scrubber should run as a long-running process that wakes up at regular intervals to scrub images. The wake up interval can be specified using the configuration option <b>wakeup_time</b>.</p> <p>If this configuration option is set to <b>False</b>, which is the default value, scrubber runs once to scrub images and exits. In this case, if the operator wishes to implement continuous scrubbing of images, scrubber needs to be scheduled as a cron job.</p> <p>Possible values:</p> <ul style="list-style-type: none"><li>• True</li><li>• False</li></ul> <p>Related options:</p> <ul style="list-style-type: none"><li>• <b>wakeup_time</b></li></ul>

Configuration option = Default value	Type	Description
<b>data_api = glance.db.sqlalchemy.api</b>	string value	<p>Python module path of data access API.</p> <p>Specifies the path to the API to use for accessing the data model. This option determines how the image catalog data will be accessed.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● <code>glance.db.sqlalchemy.api</code></li> <li>● <code>glance.db.registry.api</code></li> <li>● <code>glance.db.simple.api</code></li> </ul> <p>If this option is set to <b>glance.db.sqlalchemy.api</b> then the image catalog data is stored in and read from the database via the SQLAlchemy Core and ORM APIs.</p> <p>Setting this option to <b>glance.db.registry.api</b> will force all database access requests to be routed through the Registry service. This avoids data access from the Glance API nodes for an added layer of security, scalability and manageability.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p>In v2 OpenStack Images API, the registry service is optional. In order to use the Registry API in v2, the option <b>enable_v2_registry</b> must be set to <b>True</b>.</p> </div> </div> <p>Finally, when this configuration option is set to <b>glance.db.simple.api</b>, image catalog data is stored in and read from an in-memory data structure. This is primarily used for testing.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>enable_v2_api</code></li> <li>● <code>enable_v2_registry</code></li> </ul>
<b>debug = False</b>	boolean value	<p>If set to true, the logging level will be set to DEBUG instead of the default INFO level.</p>

Configuration option = Default value	Type	Description
<b>default_log_levels =</b> <b>[u'amqp=WARN',</b> <b>u'amqplib=WARN',</b> <b>u'boto=WARN',</b> <b>u'qpid=WARN',</b> <b>u'sqlalchemy=WARN',</b> <b>u'suds=INFO',</b> <b>u'oslo.messaging=INFO',</b> <b>u'oslo_messaging=INFO',</b> <b>u'iso8601=WARN',</b> <b>u'requests.packages.urlli</b> <b>b3.connectionpool=WAR</b> <b>N',</b> <b>u'urllib3.connectionpool=</b> <b>WARN',</b> <b>u'websocket=WARN',</b> <b>u'requests.packages.urlli</b> <b>b3.util.retry=WARN',</b> <b>u'urllib3.util.retry=WARN',</b> <b>u'keystonemiddleware=W</b> <b>ARN',</b> <b>u'routes.middleware=WA</b> <b>RN', u'stevedore=WARN',</b> <b>u'taskflow=WARN',</b> <b>u'keystoneauth=WARN',</b> <b>u'oslo.cache=INFO',</b> <b>u'dogpile.core.dogpile=IN</b> <b>FO']</b>	list value	List of package logging levels in logger=LEVEL pairs. This option is ignored if log_config_append is set.

Configuration option = Default value	Type	Description
<b>delayed_delete = False</b>	boolean value	<p>Turn on/off delayed delete.</p> <p>Typically when an image is deleted, the <b>glance-api</b> service puts the image into <b>deleted</b> state and deletes its data at the same time. Delayed delete is a feature in Glance that delays the actual deletion of image data until a later point in time (as determined by the configuration option <b>scrub_time</b>). When delayed delete is turned on, the <b>glance-api</b> service puts the image into <b>pending_delete</b> state upon deletion and leaves the image data in the storage backend for the image scrubber to delete at a later time. The image scrubber will move the image into <b>deleted</b> state upon successful deletion of image data.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p>When delayed delete is turned on, image scrubber <b>MUST</b> be running as a periodic task to prevent the backend storage from filling up with undesired usage.</p> </div> </div> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>scrub_time</b></li> <li>● <b>wakeup_time</b></li> <li>● <b>scrub_pool_size</b></li> </ul>

Configuration option = Default value	Type	Description
<b>digest_algorithm = sha256</b>	string value	<p>Digest algorithm to use for digital signature.</p> <p>Provide a string value representing the digest algorithm to use for generating digital signatures. By default, <b>sha256</b> is used.</p> <p>To get a list of the available algorithms supported by the version of OpenSSL on your platform, run the command: <b>openssl list-message-digest-algorithms</b>. Examples are <i>sha1</i>, <i>sha256</i>, and <i>sha512</i>.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p><b>digest_algorithm</b> is not related to Glance's image signing and verification. It is only used to sign the universally unique identifier (UUID) as a part of the certificate file and key file validation.</p> </div> </div> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● An OpenSSL message digest algorithm identifier</li> </ul> <p>Relation options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>enable_image_import = True</b>	boolean value	<p>Enables the Image Import workflow introduced in Pike</p> <p>As <code>[DEFAULT]/node_staging_uri</code> is required for the Image Import, it's disabled per default in Pike, enabled per default in Queens and removed in Rocky. This allows Glance to operate with previous version configs upon upgrade.</p> <p>Setting this option to False will disable the endpoints related to Image Import Refactoring work.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>[DEFAULT]/node_staging_uri</code></li> </ul>




Configuration option = Default value	Type	Description
<b>enable_v1_api = True</b>	boolean value	<p>Deploy the v1 OpenStack Images API.</p> <p>When this option is set to <b>True</b>, Glance service will respond to requests on registered endpoints conforming to the v1 OpenStack Images API.</p> <p>NOTES:</p> <ul style="list-style-type: none"> <li>● If this option is enabled, then <b>enable_v1_registry</b> must also be set to <b>True</b> to enable mandatory usage of Registry service with v1 API.</li> <li>● If this option is disabled, then the <b>enable_v1_registry</b> option, which is enabled by default, is also recommended to be disabled.</li> <li>● This option is separate from <b>enable_v2_api</b>, both v1 and v2 OpenStack Images API can be deployed independent of each other.</li> <li>● If deploying only the v2 Images API, this option, which is enabled by default, should be disabled.</li> </ul> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● enable_v1_registry</li> <li>● enable_v2_api</li> </ul>

Configuration option = Default value	Type	Description
<b>enable_v1_registry = True</b>	boolean value	<p>Deploy the v1 API Registry service.</p> <p>When this option is set to <b>True</b>, the Registry service will be enabled in Glance for v1 API requests.</p> <p>NOTES:</p> <ul style="list-style-type: none"><li>● Use of Registry is mandatory in v1 API, so this option must be set to <b>True</b> if the <b>enable_v1_api</b> option is enabled.</li><li>● If deploying only the v2 OpenStack Images API, this option, which is enabled by default, should be disabled.</li></ul> <p>Possible values:</p> <ul style="list-style-type: none"><li>● True</li><li>● False</li></ul> <p>Related options:</p> <ul style="list-style-type: none"><li>● enable_v1_api</li></ul>

Configuration option = Default value	Type	Description
<b>enable_v2_api = True</b>	boolean value	<p>Deploy the v2 OpenStack Images API.</p> <p>When this option is set to <b>True</b>, Glance service will respond to requests on registered endpoints conforming to the v2 OpenStack Images API.</p> <p>NOTES:</p> <ul style="list-style-type: none"> <li>● If this option is disabled, then the <b>enable_v2_registry</b> option, which is enabled by default, is also recommended to be disabled.</li> <li>● This option is separate from <b>enable_v1_api</b>, both v1 and v2 OpenStack Images API can be deployed independent of each other.</li> <li>● If deploying only the v1 Images API, this option, which is enabled by default, should be disabled.</li> </ul> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● enable_v2_registry</li> <li>● enable_v1_api</li> </ul>

Configuration option = Default value	Type	Description
<b>enable_v2_registry = True</b>	boolean value	<p>Deploy the v2 API Registry service.</p> <p>When this option is set to <b>True</b>, the Registry service will be enabled in Glance for v2 API requests.</p> <p>NOTES:</p> <ul style="list-style-type: none"> <li>● Use of Registry is optional in v2 API, so this option must only be enabled if both <b>enable_v2_api</b> is set to <b>True</b> and the <b>data_api</b> option is set to <b>glance.db.registry.api</b>.</li> <li>● If deploying only the v1 OpenStack Images API, this option, which is enabled by default, should be disabled.</li> </ul> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● enable_v2_api</li> <li>● data_api</li> </ul>
<b>enabled_import_methods = [u'glance-direct', u'web-download']</b>	list value	<p>List of enabled Image Import Methods</p> <p>Both <i>glance-direct</i> and <i>web-download</i> are enabled by default.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● [DEFAULT]/node_staging_uri</li> <li>● [DEFAULT]/enable_image_import</li> </ul>
<b>fatal_deprecations = False</b>	boolean value	Enables or disables fatal status of deprecations.
<b>image_location_quota = 10</b>	integer value	<p>Maximum number of locations allowed on an image.</p> <p>Any negative value is interpreted as unlimited.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>




Configuration option = Default value	Type	Description
<b>image_member_quota = 128</b>	integer value	<p>Maximum number of image members per image.</p> <p>This limits the maximum of users an image can be shared with. Any negative value is interpreted as unlimited.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<b>image_property_quota = 128</b>	integer value	<p>Maximum number of properties allowed on an image.</p> <p>This enforces an upper limit on the number of additional properties an image can have. Any negative value is interpreted as unlimited.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p>This won't have any impact if additional properties are disabled. Please refer to <b>allow_additional_image_properties</b>.</p> </div> </div> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <b>allow_additional_image_properties</b></li> </ul>

Configuration option = Default value	Type	Description
<b>image_size_cap = 1099511627776</b>	integer value	<p>Maximum size of image a user can upload in bytes.</p> <p>An image upload greater than the size mentioned here would result in an image creation failure. This configuration option defaults to 1099511627776 bytes (1 TiB).</p> <p>NOTES:</p> <ul style="list-style-type: none"> <li>• This value should only be increased after careful consideration and must be set less than or equal to 8 EiB (9223372036854775808).</li> <li>• This value must be set with careful consideration of the backend storage capacity. Setting this to a very low value may result in a large number of image failures. And, setting this to a very large value may result in faster consumption of storage. Hence, this must be set according to the nature of images created and storage capacity available.</li> </ul> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Any positive number less than or equal to 9223372036854775808</li> </ul>
<b>image_tag_quota = 128</b>	integer value	<p>Maximum number of tags allowed on an image.</p> <p>Any negative value is interpreted as unlimited.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<code>`instance_format = [instance: %(uuid)s] `</code>	string value	The format for an instance that is passed with the log message.
<code>`instance_uuid_format = [instance:%(uuid)s] `</code>	string value	The format for an instance UUID that is passed with the log message.

Configuration option = Default value	Type	Description
<b>limit_param_default = 25</b>	integer value	<p>The default number of results to return for a request.</p> <p>Responses to certain API requests, like list images, may return multiple items. The number of results returned can be explicitly controlled by specifying the <b>limit</b> parameter in the API request. However, if a <b>limit</b> parameter is not specified, this configuration value will be used as the default number of results to be returned for any API request.</p> <p>NOTES:</p> <ul style="list-style-type: none"> <li>• The value of this configuration option may not be greater than the value specified by <b>api_limit_max</b>.</li> <li>• Setting this to a very large value may slow down database queries and increase response times. Setting this to a very low value may result in poor user experience.</li> </ul> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Any positive integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <code>api_limit_max</code></li> </ul>
<b>log-config-append = None</b>	string value	<p>The name of a logging configuration file. This file is appended to any existing logging configuration files. For details about logging configuration files, see the Python logging module documentation. Note that when logging configuration files are used then all logging configuration is set in the configuration file and other logging configuration options are ignored (for example, <code>logging_context_format_string</code>).</p>
<b>log-date-format = %Y-%m-%d %H:%M:%S</b>	string value	<p>Defines the format string for <code>%(asctime)s</code> in log records. Default: <code>%(default)s</code>. This option is ignored if <code>log_config_append</code> is set.</p>
<b>log-dir = None</b>	string value	<p>(Optional) The base directory used for relative <code>log_file</code> paths. This option is ignored if <code>log_config_append</code> is set.</p>
<b>log-file = None</b>	string value	<p>(Optional) Name of log file to send logging output to. If no default is set, logging will go to <code>stderr</code> as defined by <code>use_stderr</code>. This option is ignored if <code>log_config_append</code> is set.</p>

Configuration option = Default value	Type	Description
<b>logging_context_format_string</b> = <code>%(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [%(request_id)s %(user_identity)s] %(instance)s%(message)s</code>	string value	Format string to use for log messages with context.
<b>logging_debug_format_s uffix</b> = <code>%(funcName)s %(pathname)s:%(lineno)d</code>	string value	Additional data to append to log message when logging level for the message is DEBUG.
<b>logging_default_format_s tring</b> = <code>%(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [-] %(instance)s%(message)s</code>	string value	Format string to use for log messages when context is undefined.
<b>logging_exception_prefix</b> = <code>%(asctime)s.%(msecs)03d %(process)d ERROR %(name)s %(instance)s</code>	string value	Prefix each line of exception output with this format.
<b>logging_user_identity_for mat</b> = <code>%(user)s %(tenant)s %(domain)s %(user_domain)s %(project_domain)s</code>	string value	Defines the format string for <code>%(user_identity)s</code> that is used in <code>logging_context_format_string</code> .
<b>metadata_encryption_key</b> = <code>None</code>	string value	<p>AES key for encrypting store location metadata.</p> <p>Provide a string value representing the AES cipher to use for encrypting Glance store metadata.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p>The AES key to use must be set to a random string of length 16, 24 or 32 bytes.</p> </div> </div> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● String value representing a valid AES key</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>



Configuration option = Default value	Type	Description
<b>node_staging_uri = file:///tmp/staging/</b>	string value	<p>The URL provides location where the temporary data will be stored</p> <p>This option is for Glance internal use only. Glance will save the image data uploaded by the user to <i>staging</i> endpoint during the image import process.</p> <p>This option does not change the <i>staging</i> API endpoint by any means.</p> <p> <b>NOTE</b></p> <p>It is discouraged to use same path as [task]/work_dir</p> <p> <b>NOTE</b></p> <p><i>file://&lt;absolute-directory-path&gt;</i> is the only option api_image_import flow will support for now.</p> <p> <b>NOTE</b></p> <p>The staging path must be on shared filesystem available to all Glance API nodes.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● String starting with <i>file://</i> followed by absolute FS path</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● [task]/work_dir</li> <li>● [DEFAULT]/enable_image_import (<b>deprecated</b>)</li> </ul>
<b>publish_errors = False</b>	boolean value	Enables or disables publication of error events.

Configuration option = Default value	Type	Description
<b>pydev_worker_debug_host = None</b>	unknown value	<p>Host address of the pydev server.</p> <p>Provide a string value representing the hostname or IP of the pydev server to use for debugging. The pydev server listens for debug connections on this address, facilitating remote debugging in Glance.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Valid hostname</li> <li>● Valid IP address</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>pydev_worker_debug_port = 5678</b>	port value	<p>Port number that the pydev server will listen on.</p> <p>Provide a port number to bind the pydev server to. The pydev process accepts debug connections on this port and facilitates remote debugging in Glance.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid port number</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>rate_limit_burst = 0</b>	integer value	Maximum number of logged messages per <code>rate_limit_interval</code> .
<b>rate_limit_except_level = CRITICAL</b>	string value	Log level name used by rate limiting: CRITICAL, ERROR, INFO, WARNING, DEBUG or empty string. Logs with level greater or equal to <code>rate_limit_except_level</code> are not filtered. An empty string means that all levels are filtered.
<b>rate_limit_interval = 0</b>	integer value	Interval, number of seconds, of log rate limiting.

Configuration option = Default value	Type	Description
<b>scrub_pool_size = 1</b>	integer value	<p>The size of thread pool to be used for scrubbing images.</p> <p>When there are a large number of images to scrub, it is beneficial to scrub images in parallel so that the scrub queue stays in control and the backend storage is reclaimed in a timely fashion. This configuration option denotes the maximum number of images to be scrubbed in parallel. The default value is one, which signifies serial scrubbing. Any value above one indicates parallel scrubbing.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any non-zero positive integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>delayed_delete</b></li> </ul>
<b>scrub_time = 0</b>	integer value	<p>The amount of time, in seconds, to delay image scrubbing.</p> <p>When delayed delete is turned on, an image is put into <b>pending_delete</b> state upon deletion until the scrubber deletes its image data. Typically, soon after the image is put into <b>pending_delete</b> state, it is available for scrubbing. However, scrubbing can be delayed until a later point using this configuration option. This option denotes the time period an image spends in <b>pending_delete</b> state before it is available for scrubbing.</p> <p>It is important to realize that this has storage implications. The larger the <b>scrub_time</b>, the longer the time to reclaim backend storage from deleted images.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any non-negative integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>delayed_delete</b></li> </ul>

Configuration option = Default value	Type	Description
<b>show_image_direct_url = False</b>	boolean value	<p>Show direct image location when returning an image.</p> <p>This configuration option indicates whether to show the direct image location when returning image details to the user. The direct image location is where the image data is stored in backend storage. This image location is shown under the image property <b>direct_url</b>.</p> <p>When multiple image locations exist for an image, the best location is displayed based on the location strategy indicated by the configuration option <b>location_strategy</b>.</p> <p>NOTES:</p> <ul style="list-style-type: none"> <li>● Revealing image locations can present a GRAVE SECURITY RISK as image locations can sometimes include credentials. Hence, this is set to <b>False</b> by default. Set this to <b>True</b> with EXTREME CAUTION and ONLY IF you know what you are doing!</li> <li>● If an operator wishes to avoid showing any image location(s) to the user, then both this option and <b>show_multiple_locations</b> MUST be set to <b>False</b>.</li> </ul> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● show_multiple_locations</li> <li>● location_strategy</li> </ul>

Configuration option = Default value	Type	Description
<b>show_multiple_locations = False</b>	boolean value	<p>Show all image locations when returning an image.</p> <p>This configuration option indicates whether to show all the image locations when returning image details to the user. When multiple image locations exist for an image, the locations are ordered based on the location strategy indicated by the configuration opt <b>location_strategy</b>. The image locations are shown under the image property <b>locations</b>.</p> <p>NOTES:</p> <ul style="list-style-type: none"> <li>● Revealing image locations can present a GRAVE SECURITY RISK as image locations can sometimes include credentials. Hence, this is set to <b>False</b> by default. Set this to <b>True</b> with EXTREME CAUTION and ONLY IF you know what you are doing!</li> <li>● See <a href="https://wiki.openstack.org/wiki/OSSN/OSSN-0065">https://wiki.openstack.org/wiki/OSSN/OSSN-0065</a> for more information.</li> <li>● If an operator wishes to avoid showing any image location(s) to the user, then both this option and <b>show_image_direct_url</b> MUST be set to <b>False</b>.</li> </ul> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● show_image_direct_url</li> <li>● location_strategy</li> </ul>
<b>syslog-log-facility = LOG_USER</b>	string value	Syslog facility to receive log lines. This option is ignored if log_config_append is set.
<b>use-journal = False</b>	boolean value	Enable journald for logging. If running in a systemd environment you may wish to enable journal support. Doing so will use the journal native protocol which includes structured metadata in addition to log messages. This option is ignored if log_config_append is set.
<b>use-json = False</b>	boolean value	Use JSON formatting for logging. This option is ignored if log_config_append is set.

Configuration option = Default value	Type	Description
<b>use-syslog = False</b>	boolean value	Use syslog for logging. Existing syslog format is DEPRECATED and will be changed later to honor RFC5424. This option is ignored if log_config_append is set.
<b>use_stderr = False</b>	boolean value	Log output to standard error. This option is ignored if log_config_append is set.
<b>user_storage_quota = 0</b>	string value	<p>Maximum amount of image storage per tenant.</p> <p>This enforces an upper limit on the cumulative storage consumed by all images of a tenant across all stores. This is a per-tenant limit.</p> <p>The default unit for this configuration option is Bytes. However, storage units can be specified using case-sensitive literals <b>B</b>, <b>KB</b>, <b>MB</b>, <b>GB</b> and <b>TB</b> representing Bytes, KiloBytes, MegaBytes, GigaBytes and TeraBytes respectively. Note that there should not be any space between the value and unit. Value <b>0</b> signifies no quota enforcement. Negative values are invalid and result in errors.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A string that is a valid concatenation of a non-negative integer representing the storage value and an optional string literal representing storage units as mentioned above.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>

Configuration option = Default value	Type	Description
<b>wakeup_time = 300</b>	integer value	<p>Time interval, in seconds, between scrubber runs in daemon mode.</p> <p>Scrubber can be run either as a cron job or daemon. When run as a daemon, this configuration time specifies the time period between two runs. When the scrubber wakes up, it fetches and scrubs all <b>pending_delete</b> images that are available for scrubbing after taking <b>scrub_time</b> into consideration.</p> <p>If the wakeup time is set to a large number, there may be a large number of images to be scrubbed for each run. Also, this impacts how quickly the backend storage is reclaimed.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Any non-negative integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <b>daemon</b></li> <li>• <b>delayed_delete</b></li> </ul>
<b>watch-log-file = False</b>	boolean value	<p>Uses logging handler designed to watch file system. When log file is moved or removed this handler will open a new log file with specified path instantaneously. It makes sense only if log_file option is specified and Linux platform is used. This option is ignored if log_config_append is set.</p>

### 5.3.2. database

The following table outlines the options available under the **[database]** group in the **/etc/glance/glance-scrubber.conf** file.

Table 5.34. database

Configuration option = Default value	Type	Description
<b>backend = sqlalchemy</b>	string value	The back end to use for the database.
<b>connection = None</b>	string value	The SQLAlchemy connection string to use to connect to the database.

Configuration option = Default value	Type	Description
<b>connection_debug = 0</b>	integer value	Verbosity of SQL debugging information: 0=None, 100=Everything.
<b>connection_recycle_time = 3600</b>	integer value	Connections which have been present in the connection pool longer than this number of seconds will be replaced with a new one the next time they are checked out from the pool.
<b>connection_trace = False</b>	boolean value	Add Python stack traces to SQL as comment strings.
<b>db_inc_retry_interval = True</b>	boolean value	If True, increases the interval between retries of a database operation up to db_max_retry_interval.
<b>db_max_retries = 20</b>	integer value	Maximum retries in case of connection error or deadlock error before error is raised. Set to -1 to specify an infinite retry count.
<b>db_max_retry_interval = 10</b>	integer value	If db_inc_retry_interval is set, the maximum seconds between retries of a database operation.
<b>db_retry_interval = 1</b>	integer value	Seconds between retries of a database transaction.
<b>max_overflow = 50</b>	integer value	If set, use this value for max_overflow with SQLAlchemy.
<b>max_pool_size = 5</b>	integer value	Maximum number of SQL connections to keep open in a pool. Setting a value of 0 indicates no limit.
<b>max_retries = 10</b>	integer value	Maximum number of database connection retries during startup. Set to -1 to specify an infinite retry count.
<b>min_pool_size = 1</b>	integer value	Minimum number of SQL connections to keep open in a pool.
<b>mysql_enable_ndb = False</b>	boolean value	If True, transparently enables support for handling MySQL Cluster (NDB).
<b>mysql_sql_mode = TRADITIONAL</b>	string value	The SQL mode to be used for MySQL sessions. This option, including the default, overrides any server-set SQL mode. To use whatever SQL mode is set by the server configuration, set this to no value. Example: mysql_sql_mode=
<b>pool_timeout = None</b>	integer value	If set, use this value for pool_timeout with SQLAlchemy.



Configuration option = Default value	Type	Description
<b>retry_interval = 10</b>	integer value	Interval between retries of opening a SQL connection.
<b>slave_connection = None</b>	string value	The SQLAlchemy connection string to use to connect to the slave database.
<b>sqlite_synchronous = True</b>	boolean value	If True, SQLite uses synchronous mode.
<b>use_db_reconnect = False</b>	boolean value	Enable the experimental use of database reconnect on connection lost.
<b>use_tpool = False</b>	boolean value	Enable the experimental use of thread pooling for all DB API calls

### 5.3.3. glance\_store

The following table outlines the options available under the **[glance\_store]** group in the `/etc/glance/glance-scrubber.conf` file.

Table 5.35. glance\_store

Configuration option = Default value	Type	Description
<b>cinder_api_insecure = False</b>	boolean value	<p>Allow to perform insecure SSL requests to cinder.</p> <p>If this option is set to True, HTTPS endpoint connection is verified using the CA certificates file specified by <b>cinder_ca_certificates_file</b> option.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● cinder_ca_certificates_file</li> </ul>

Configuration option = Default value	Type	Description
<b>cinder_ca_certificates_file = None</b>	string value	<p>Location of a CA certificates file used for cinder client requests.</p> <p>The specified CA certificates file, if set, is used to verify cinder connections via HTTPS endpoint. If the endpoint is HTTP, this value is ignored. <b>cinder_api_insecure</b> must be set to <b>True</b> to enable the verification.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Path to a ca certificates file</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>cinder_api_insecure</code></li> </ul>
<b>cinder_catalog_info = volumev2::publicURL</b>	string value	<p>Information to match when looking for cinder in the service catalog.</p> <p>When the <b>cinder_endpoint_template</b> is not set and any of <b>cinder_store_auth_address</b>, <b>cinder_store_user_name</b>, <b>cinder_store_project_name</b>, <b>cinder_store_password</b> is not set, cinder store uses this information to lookup cinder endpoint from the service catalog in the current context. <b>cinder_os_region_name</b>, if set, is taken into consideration to fetch the appropriate endpoint.</p> <p>The service catalog can be listed by the <b>openstack catalog list</b> command.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A string of of the following form: <b>&lt;service_type&gt;:&lt;service_name&gt;: &lt;interface&gt;</b> At least <b>service_type</b> and <b>interface</b> should be specified. <b>service_name</b> can be omitted.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>cinder_os_region_name</code></li> <li>● <code>cinder_endpoint_template</code></li> <li>● <code>cinder_store_auth_address</code></li> <li>● <code>cinder_store_user_name</code></li> <li>● <code>cinder_store_project_name</code></li> <li>● <code>cinder_store_password</code></li> </ul>


Configuration option = Default value	Type	Description
<b>cinder_endpoint_template = None</b>	string value	<p>Override service catalog lookup with template for cinder endpoint.</p> <p>When this option is set, this value is used to generate cinder endpoint, instead of looking up from the service catalog. This value is ignored if <b>cinder_store_auth_address</b>, <b>cinder_store_user_name</b>, <b>cinder_store_project_name</b>, and <b>cinder_store_password</b> are specified.</p> <p>If this configuration option is set, <b>cinder_catalog_info</b> will be ignored.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● URL template string for cinder endpoint, where <b>%%(tenant)s</b> is replaced with the current tenant (project) name. For example: <b>http://cinder.openstack.example.org/v2/%%(tenant)s</b></li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● cinder_store_auth_address</li> <li>● cinder_store_user_name</li> <li>● cinder_store_project_name</li> <li>● cinder_store_password</li> <li>● cinder_catalog_info</li> </ul>
<b>cinder_http_retries = 3</b>	integer value	<p>Number of cinderclient retries on failed http calls.</p> <p>When a call failed by any errors, cinderclient will retry the call up to the specified times after sleeping a few seconds.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A positive integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>

Configuration option = Default value	Type	Description
<b>cinder_os_region_name = None</b>	string value	<p>Region name to lookup cinder service from the service catalog.</p> <p>This is used only when <b>cinder_catalog_info</b> is used for determining the endpoint. If set, the lookup for cinder endpoint by this node is filtered to the specified region. It is useful when multiple regions are listed in the catalog. If this is not set, the endpoint is looked up from every region.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A string that is a valid region name.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>cinder_catalog_info</code></li> </ul>
<b>cinder_state_transition_timeout = 300</b>	integer value	<p>Time period, in seconds, to wait for a cinder volume transition to complete.</p> <p>When the cinder volume is created, deleted, or attached to the glance node to read/write the volume data, the volume's state is changed. For example, the newly created volume status changes from <b>creating</b> to <b>available</b> after the creation process is completed. This specifies the maximum time to wait for the status change. If a timeout occurs while waiting, or the status is changed to an unexpected value (e.g. <b>error</b>), the image creation fails.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A positive integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>

Configuration option = Default value	Type	Description
<b>cinder_store_auth_address = None</b>	string value	<p>The address where the cinder authentication service is listening.</p> <p>When all of <b>cinder_store_auth_address</b>, <b>cinder_store_user_name</b>, <b>cinder_store_project_name</b>, and <b>cinder_store_password</b> options are specified, the specified values are always used for the authentication. This is useful to hide the image volumes from users by storing them in a project/tenant specific to the image service. It also enables users to share the image volume among other projects under the control of glance's ACL.</p> <p>If either of these options are not set, the cinder endpoint is looked up from the service catalog, and current context's user and project are used.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid authentication service address, for example: <b>http://openstack.example.org/identity/v2.0</b></li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● cinder_store_user_name</li> <li>● cinder_store_password</li> <li>● cinder_store_project_name</li> </ul>
<b>cinder_store_password = None</b>	string value	<p>Password for the user authenticating against cinder.</p> <p>This must be used with all the following related options. If any of these are not specified, the user of the current context is used.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid password for the user specified by <b>cinder_store_user_name</b></li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● cinder_store_auth_address</li> <li>● cinder_store_user_name</li> <li>● cinder_store_project_name</li> </ul>

Configuration option = Default value	Type	Description
<b>cinder_store_project_name = None</b>	string value	<p>Project name where the image volume is stored in cinder.</p> <p>If this configuration option is not set, the project in current context is used.</p> <p>This must be used with all the following related options. If any of these are not specified, the project of the current context is used.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid project name</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>cinder_store_auth_address</b></li> <li>● <b>cinder_store_user_name</b></li> <li>● <b>cinder_store_password</b></li> </ul>
<b>cinder_store_user_name = None</b>	string value	<p>User name to authenticate against cinder.</p> <p>This must be used with all the following related options. If any of these are not specified, the user of the current context is used.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid user name</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● cinder_store_auth_address</li> <li>● cinder_store_password</li> <li>● cinder_store_project_name</li> </ul>

Configuration option = Default value	Type	Description
<b>cinder_volume_type = None</b>	string value	<p>Volume type that will be used for volume creation in cinder.</p> <p>Some cinder backends can have several volume types to optimize storage usage. Adding this option allows an operator to choose a specific volume type in cinder that can be optimized for images.</p> <p>If this is not set, then the default volume type specified in the cinder configuration will be used for volume creation.</p> <p>Possible values:</p> <ul style="list-style-type: none"><li>• A valid volume type from cinder</li></ul> <p>Related options:</p> <ul style="list-style-type: none"><li>• None</li></ul>

Configuration option = Default value	Type	Description
<b>default_store = file</b>	string value	<p>The default scheme to use for storing images.</p> <p>Provide a string value representing the default scheme to use for storing images. If not set, Glance uses <b>file</b> as the default scheme to store images with the <b>file</b> store.</p> <div data-bbox="815 510 922 703" style="display: inline-block; vertical-align: top;">  </div> <p><b>NOTE</b></p> <p>The value given for this configuration option must be a valid scheme for a store registered with the <b>stores</b> configuration option.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● file</li> <li>● filesystem</li> <li>● http</li> <li>● https</li> <li>● swift</li> <li>● swift+http</li> <li>● swift+https</li> <li>● swift+config</li> <li>● rbd</li> <li>● sheepdog</li> <li>● cinder</li> <li>● vsphere</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● stores</li> </ul>



Configuration option = Default value	Type	Description
<b>default_swift_reference = ref1</b>	string value	<p>Reference to default Swift account/backing store parameters.</p> <p>Provide a string value representing a reference to the default set of parameters required for using swift account/backing store for image storage. The default reference value for this configuration option is <i>ref1</i>. This configuration option dereferences the parameters and facilitates image storage in Swift storage backend every time a new image is added.</p> <p>Possible values:</p> <ul style="list-style-type: none"><li>• A valid string value</li></ul> <p>Related options:</p> <ul style="list-style-type: none"><li>• None</li></ul>

Configuration option = Default value	Type	Description
<b>filesystem_store_datadir</b> = <b>/var/lib/glance/images</b>	string value	<p>Directory to which the filesystem backend store writes images.</p> <p>Upon start up, Glance creates the directory if it doesn't already exist and verifies write access to the user under which <b>glance-api</b> runs. If the write access isn't available, a <b>BadStoreConfiguration</b> exception is raised and the filesystem store may not be available for adding new images.</p> <div style="display: flex; align-items: flex-start;"> <div style="width: 40px; height: 100px; border: 1px solid black; background: repeating-linear-gradient(45deg, transparent, transparent 2px, black 2px, black 4px); margin-right: 10px;"></div> <div> <p><b>NOTE</b></p> <p>This directory is used only when filesystem store is used as a storage backend. Either <b>filesystem_store_datadir</b> or <b>filesystem_store_datadirs</b> option must be specified in <b>glance-api.conf</b>. If both options are specified, a <b>BadStoreConfiguration</b> will be raised and the filesystem store may not be available for adding new images.</p> </div> </div> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid path to a directory</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>filesystem_store_datadirs</b></li> <li>● <b>filesystem_store_file_perm</b></li> </ul>

Configuration option = Default value	Type	Description
<b>filesystem_store_datadirs</b> = None	multi valued	<p>List of directories and their priorities to which the filesystem backend store writes images.</p> <p>The filesystem store can be configured to store images in multiple directories as opposed to using a single directory specified by the <b>filesystem_store_datadir</b> configuration option. When using multiple directories, each directory can be given an optional priority to specify the preference order in which they should be used. Priority is an integer that is concatenated to the directory path with a colon where a higher value indicates higher priority. When two directories have the same priority, the directory with most free space is used. When no priority is specified, it defaults to zero.</p> <p>More information on configuring filesystem store with multiple store directories can be found at <a href="http://docs.openstack.org/developer/glance/configuring.html">http://docs.openstack.org/developer/glance/configuring.html</a></p> <div data-bbox="815 1039 922 1480" style="background: repeating-linear-gradient(45deg, transparent, transparent 2px, gray 2px, gray 4px); width: 67px; height: 200px; margin-bottom: 10px;"></div> <p><b>NOTE</b></p> <p>This directory is used only when filesystem store is used as a storage backend. Either <b>filesystem_store_datadir</b> or <b>filesystem_store_datadirs</b> option must be specified in <b>glance-api.conf</b>. If both options are specified, a <b>BadStoreConfiguration</b> will be raised and the filesystem store may not be available for adding new images.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● List of strings of the following form:             <ul style="list-style-type: none"> <li>○ <b>&lt;a valid directory path&gt;:&lt;optional integer priority&gt;</b></li> </ul> </li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>filesystem_store_datadir</b></li> <li>● <b>filesystem_store_file_perm</b></li> </ul>

Configuration option = Default value	Type	Description
<b>filesystem_store_file_permissions = 0</b>	integer value	<p>File access permissions for the image files.</p> <p>Set the intended file access permissions for image data. This provides a way to enable other services, e.g. Nova, to consume images directly from the filesystem store. The users running the services that are intended to be given access to could be made a member of the group that owns the files created. Assigning a value less than or equal to zero for this configuration option signifies that no changes be made to the default permissions. This value will be decoded as an octal digit.</p> <p>For more information, please refer the documentation at <a href="http://docs.openstack.org/developer/glance/configuring.html">http://docs.openstack.org/developer/glance/configuring.html</a></p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid file access permission</li> <li>● Zero</li> <li>● Any negative integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>filesystem_store_metadata_file = None</b>	string value	<p>Filesystem store metadata file.</p> <p>The path to a file which contains the metadata to be returned with any location associated with the filesystem store. The file must contain a valid JSON object. The object should contain the keys <b>id</b> and <b>mountpoint</b>. The value for both keys should be a string.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid path to the store metadata file</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>

Configuration option = Default value	Type	Description
<b>http_proxy_information =</b> <b>{}</b>	dict value	<p>The http/https proxy information to be used to connect to the remote server.</p> <p>This configuration option specifies the http/https proxy information that should be used to connect to the remote server. The proxy information should be a key value pair of the scheme and proxy, for example, http:10.0.0.1:3128. You can also specify proxies for multiple schemes by separating the key value pairs with a comma, for example, http:10.0.0.1:3128, https:10.0.0.1:1080.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A comma separated list of scheme:proxy pairs as described above</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<b>https_ca_certificates_file</b> <b>= None</b>	string value	<p>Path to the CA bundle file.</p> <p>This configuration option enables the operator to use a custom Certificate Authority file to verify the remote server certificate. If this option is set, the <b>https_insecure</b> option will be ignored and the CA file specified will be used to authenticate the server certificate and establish a secure connection to the server.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A valid path to a CA file</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• https_insecure</li> </ul>

Configuration option = Default value	Type	Description
<b>https_insecure = True</b>	boolean value	<p>Set verification of the remote server certificate.</p> <p>This configuration option takes in a boolean value to determine whether or not to verify the remote server certificate. If set to True, the remote server certificate is not verified. If the option is set to False, then the default CA truststore is used for verification.</p> <p>This option is ignored if <b>https_ca_certificates_file</b> is set. The remote server certificate will then be verified using the file specified using the <b>https_ca_certificates_file</b> option.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• True</li> <li>• False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• https_ca_certificates_file</li> </ul>
<b>rados_connect_timeout = 0</b>	integer value	<p>Timeout value for connecting to Ceph cluster.</p> <p>This configuration option takes in the timeout value in seconds used when connecting to the Ceph cluster i.e. it sets the time to wait for glance-api before closing the connection. This prevents glance-api hangups during the connection to RBD. If the value for this option is set to less than or equal to 0, no timeout is set and the default librados value is used.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>• Any integer value</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>

Configuration option = Default value	Type	Description
<b>rbd_store_ceph_conf = /etc/ceph/ceph.conf</b>	string value	<p>Ceph configuration file path.</p> <p>This configuration option takes in the path to the Ceph configuration file to be used. If the value for this option is not set by the user or is set to None, librados will locate the default configuration file which is located at /etc/ceph/ceph.conf. If using Cephx authentication, this file should include a reference to the right keyring in a client.&lt;USER&gt; section</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>● A valid path to a configuration file</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● rbd_store_user</li> </ul>
<b>rbd_store_chunk_size = 8</b>	integer value	<p>Size, in megabytes, to chunk RADOS images into.</p> <p>Provide an integer value representing the size in megabytes to chunk Glance images into. The default chunk size is 8 megabytes. For optimal performance, the value should be a power of two.</p> <p>When Ceph's RBD object storage system is used as the storage backend for storing Glance images, the images are chunked into objects of the size set using this option. These chunked objects are then stored across the distributed block data store to use for Glance.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>● Any positive integer value</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>

Configuration option = Default value	Type	Description
<b>rbd_store_pool = images</b>	string value	<p>RADOS pool in which images are stored.</p> <p>When RBD is used as the storage backend for storing Glance images, the images are stored by means of logical grouping of the objects (chunks of images) into a <b>pool</b>. Each pool is defined with the number of placement groups it can contain. The default pool that is used is <i>images</i>.</p> <p>More information on the RBD storage backend can be found here: <a href="http://ceph.com/planet/how-data-is-stored-in-ceph-cluster/">http://ceph.com/planet/how-data-is-stored-in-ceph-cluster/</a></p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>● A valid pool name</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>rbd_store_user = None</b>	string value	<p>RADOS user to authenticate as.</p> <p>This configuration option takes in the RADOS user to authenticate as. This is only needed when RADOS authentication is enabled and is applicable only if the user is using Cephx authentication. If the value for this option is not set by the user or is set to None, a default value will be chosen, which will be based on the client. section in <code>rbd_store_ceph_conf</code>.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>● A valid RADOS user</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>rbd_store_ceph_conf</code></li> </ul>



Configuration option = Default value	Type	Description
<b>rootwrap_config = /etc/glance/rootwrap.conf</b>	string value	<p>Path to the rootwrap configuration file to use for running commands as root.</p> <p>The cinder store requires root privileges to operate the image volumes (for connecting to iSCSI/FC volumes and reading/writing the volume data, etc.). The configuration file should allow the required commands by cinder store and os-brick library.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Path to the rootwrap config file</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>sheepdog_store_address = 127.0.0.1</b>	unknown value	<p>Address to bind the Sheepdog daemon to.</p> <p>Provide a string value representing the address to bind the Sheepdog daemon to. The default address set for the <i>sheep</i> is 127.0.0.1.</p> <p>The Sheepdog daemon, also called <i>sheep</i>, manages the storage in the distributed cluster by writing objects across the storage network. It identifies and acts on the messages directed to the address set using <b>sheepdog_store_address</b> option to store chunks of Glance images.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid IPv4 address</li> <li>● A valid IPv6 address</li> <li>● A valid hostname</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● sheepdog_store_port</li> </ul>


Configuration option = Default value	Type	Description
<b>sheepdog_store_chunk_size = 64</b>	integer value	<p>Chunk size for images to be stored in Sheepdog data store.</p> <p>Provide an integer value representing the size in mebibyte (1048576 bytes) to chunk Glance images into. The default chunk size is 64 mebibytes.</p> <p>When using Sheepdog distributed storage system, the images are chunked into objects of this size and then stored across the distributed data store to use for Glance.</p> <p>Chunk sizes, if a power of two, help avoid fragmentation and enable improved performance.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Positive integer value representing size in mebibytes.</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>sheepdog_store_port = 7000</b>	port value	<p>Port number on which the sheep daemon will listen.</p> <p>Provide an integer value representing a valid port number on which you want the Sheepdog daemon to listen on. The default port is 7000.</p> <p>The Sheepdog daemon, also called <i>sheep</i>, manages the storage in the distributed cluster by writing objects across the storage network. It identifies and acts on the messages it receives on the port number set using <b>sheepdog_store_port</b> option to store chunks of Glance images.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid port number (0 to 65535)</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● <code>sheepdog_store_address</code></li> </ul>

Configuration option = Default value	Type	Description
<b>store_capabilities_update_min_interval = 0</b>	integer value	<p>Minimum interval in seconds to execute updating dynamic storage capabilities based on current backend status.</p> <p>Provide an integer value representing time in seconds to set the minimum interval before an update of dynamic storage capabilities for a storage backend can be attempted. Setting <b>store_capabilities_update_min_interval</b> does not mean updates occur periodically based on the set interval. Rather, the update is performed at the elapse of this interval set, if an operation of the store is triggered.</p> <p>By default, this option is set to zero and is disabled. Provide an integer value greater than zero to enable this option.</p> <div data-bbox="815 898 922 1155" style="float: left; width: 60px; height: 115px; background: repeating-linear-gradient(45deg, transparent, transparent 2px, #ccc 2px, #ccc 4px); border: 1px solid #ccc; margin-bottom: 10px;"></div> <p><b>NOTE</b></p> <p>For more information on store capabilities and their updates, please visit:  <a href="https://specs.openstack.org/openstack/glance-specs/specs/kilo/store-capabilities.html">https://specs.openstack.org/openstack/glance-specs/specs/kilo/store-capabilities.html</a></p> <p>For more information on setting up a particular store in your deployment and help with the usage of this feature, please contact the storage driver maintainers listed here:  <a href="http://docs.openstack.org/developer/glance_store_drivers/index.html">http://docs.openstack.org/developer/glance_store_drivers/index.html</a></p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Zero</li> <li>● Positive integer</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>

Configuration option = Default value	Type	Description
<b>stores = [u'file', u'http']</b>	list value	<p>List of enabled Glance stores.</p> <p>Register the storage backends to use for storing disk images as a comma separated list. The default stores enabled for storing disk images with Glance are <b>file</b> and <b>http</b>.</p> <p>Possible values:</p> <ul style="list-style-type: none"><li>● A comma separated list that could include:<ul style="list-style-type: none"><li>○ file</li><li>○ http</li><li>○ swift</li><li>○ rbd</li><li>○ sheepdog</li><li>○ cinder</li><li>○ vmware</li></ul></li></ul> <p>Related Options:</p> <ul style="list-style-type: none"><li>● default_store</li></ul>


Configuration option = Default value	Type	Description
<b>swift_buffer_on_upload = False</b>	boolean value	<p>Buffer image segments before upload to Swift.</p> <p>Provide a boolean value to indicate whether or not Glance should buffer image data to disk while uploading to swift. This enables Glance to resume uploads on error.</p> <p>NOTES: When enabling this option, one should take great care as this increases disk usage on the API node. Be aware that depending upon how the file system is configured, the disk space used for buffering may decrease the actual disk space available for the glance image cache. Disk utilization will cap according to the following equation: <b>(swift_store_large_object_chunk_size * workers * 1000)</b></p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• True</li> <li>• False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• swift_upload_buffer_dir</li> </ul>
<b>swift_store_admin_tenant s = []</b>	list value	<p>List of tenants that will be granted admin access.</p> <p>This is a list of tenants that will be granted read/write access on all Swift containers created by Glance in multi-tenant mode. The default value is an empty list.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A comma separated list of strings representing UUIDs of Keystone projects/tenants</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<b>swift_store_auth_address = None</b>	string value	<p>The address where the Swift authentication service is listening.</p>

Configuration option = Default value	Type	Description
<b>swift_store_auth_insecure = False</b>	boolean value	<p>Set verification of the server certificate.</p> <p>This boolean determines whether or not to verify the server certificate. If this option is set to True, swiftclient won't check for a valid SSL certificate when authenticating. If the option is set to False, then the default CA truststore is used for verification.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• True</li> <li>• False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• swift_store_cacert</li> </ul>
<b>swift_store_auth_version = 2</b>	string value	<p>Version of the authentication service to use. Valid versions are 2 and 3 for keystone and 1 (deprecated) for swauth and rackspace.</p>
<b>swift_store_cacert = None</b>	string value	<p>Path to the CA bundle file.</p> <p>This configuration option enables the operator to specify the path to a custom Certificate Authority file for SSL verification when connecting to Swift.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A valid path to a CA file</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• swift_store_auth_insecure</li> </ul>

Configuration option = Default value	Type	Description
<b>swift_store_config_file =</b> <b>None</b>	string value	<p>Absolute path to the file containing the swift account(s) configurations.</p> <p>Include a string value representing the path to a configuration file that has references for each of the configured Swift account(s)/backing stores. By default, no file path is specified and customized Swift referencing is disabled. Configuring this option is highly recommended while using Swift storage backend for image storage as it avoids storage of credentials in the database.</p> <div data-bbox="815 689 922 882" style="display: inline-block; vertical-align: top;">  </div> <p><b>NOTE</b></p> <p>Please do not configure this option if you have set <b>swift_store_multi_tenant</b> to <b>True</b>.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● String value representing an absolute path on the glance-api node</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● swift_store_multi_tenant</li> </ul>


Configuration option = Default value	Type	Description
<b>swift_store_container = glance</b>	string value	<p>Name of single container to store images/name prefix for multiple containers</p> <p>When a single container is being used to store images, this configuration option indicates the container within the Glance account to be used for storing all images. When multiple containers are used to store images, this will be the name prefix for all containers. Usage of single/multiple containers can be controlled using the configuration option <b>swift_store_multiple_containers_seed</b>.</p> <p>When using multiple containers, the containers will be named after the value set for this configuration option with the first N chars of the image UUID as the suffix delimited by an underscore (where N is specified by <b>swift_store_multiple_containers_seed</b>).</p> <p>Example: if the seed is set to 3 and <code>swift_store_container = glance</code>, then an image with UUID <b>fdae39a1-bac5-4238-aba4-69bcc726e848</b> would be placed in the container <b>glance_fda</b>. All dashes in the UUID are included when creating the container name but do not count toward the character limit, so when N=10 the container name would be <b>glance_fdae39a1-ba</b>.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● If using single container, this configuration option can be any string that is a valid swift container name in Glance's Swift account</li> <li>● If using multiple containers, this configuration option can be any string as long as it satisfies the container naming rules enforced by Swift. The value of <b>swift_store_multiple_containers_seed</b> should be taken into account as well.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>swift_store_multiple_containers_seed</b></li> <li>● <b>swift_store_multi_tenant</b></li> <li>● <b>swift_store_create_container_on_put</b></li> </ul>




Configuration option = Default value	Type	Description
<b>swift_store_create_container_on_put = False</b>	boolean value	<p>Create container, if it doesn't already exist, when uploading image.</p> <p>At the time of uploading an image, if the corresponding container doesn't exist, it will be created provided this configuration option is set to True. By default, it won't be created. This behavior is applicable for both single and multiple containers mode.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>swift_store_endpoint = None</b>	string value	<p>The URL endpoint to use for Swift backend storage.</p> <p>Provide a string value representing the URL endpoint to use for storing Glance images in Swift store. By default, an endpoint is not set and the storage URL returned by <b>auth</b> is used. Setting an endpoint with <b>swift_store_endpoint</b> overrides the storage URL and is used for Glance image storage.</p> <div style="display: flex; align-items: flex-start; margin-top: 10px;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p>The URL should include the path up to, but excluding the container. The location of an object is obtained by appending the container and object to the configured URL.</p> </div> </div> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● String value representing a valid URL path up to a Swift container</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>

Configuration option = Default value	Type	Description
<b>swift_store_endpoint_type = publicURL</b>	string value	<p>Endpoint Type of Swift service.</p> <p>This string value indicates the endpoint type to use to fetch the Swift endpoint. The endpoint type determines the actions the user will be allowed to perform, for instance, reading and writing to the Store. This setting is only used if <code>swift_store_auth_version</code> is greater than 1.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● <code>publicURL</code></li> <li>● <code>adminURL</code></li> <li>● <code>internalURL</code></li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>swift_store_endpoint</code></li> </ul>
<b>swift_store_expire_soon_interval = 60</b>	integer value	<p>Time in seconds defining the size of the window in which a new token may be requested before the current token is due to expire.</p> <p>Typically, the Swift storage driver fetches a new token upon the expiration of the current token to ensure continued access to Swift. However, some Swift transactions (like uploading image segments) may not recover well if the token expires on the fly.</p> <p>Hence, by fetching a new token before the current token expiration, we make sure that the token does not expire or is close to expiry before a transaction is attempted. By default, the Swift storage driver requests for a new token 60 seconds or less before the current token expiration.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Zero</li> <li>● Positive integer value</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>swift_store_key = None</b>	string value	Auth key for the user authenticating against the Swift authentication service.


Configuration option = Default value	Type	Description
<b>swift_store_large_object_chunk_size = 200</b>	integer value	<p>The maximum size, in MB, of the segments when image data is segmented.</p> <p>When image data is segmented to upload images that are larger than the limit enforced by the Swift cluster, image data is broken into segments that are no bigger than the size specified by this configuration option. Refer to <b>swift_store_large_object_size</b> for more detail.</p> <p>For example: if <b>swift_store_large_object_size</b> is 5GB and <b>swift_store_large_object_chunk_size</b> is 1GB, an image of size 6.2GB will be segmented into 7 segments where the first six segments will be 1GB in size and the seventh segment will be 0.2GB.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A positive integer that is less than or equal to the large object limit enforced by Swift cluster in consideration.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <b>swift_store_large_object_size</b></li> </ul>

Configuration option = Default value	Type	Description
<b>swift_store_large_object_size = 5120</b>	integer value	<p>The size threshold, in MB, after which Glance will start segmenting image data.</p> <p>Swift has an upper limit on the size of a single uploaded object. By default, this is 5GB. To upload objects bigger than this limit, objects are segmented into multiple smaller objects that are tied together with a manifest file. For more detail, refer to <a href="http://docs.openstack.org/developer/swift/overview_large_objects.html">http://docs.openstack.org/developer/swift/overview_large_objects.html</a></p> <p>This configuration option specifies the size threshold over which the Swift driver will start segmenting image data into multiple smaller files. Currently, the Swift driver only supports creating Dynamic Large Objects.</p> <div data-bbox="815 864 922 1057" style="border: 1px solid black; padding: 5px; width: fit-content;">  </div> <p><b>NOTE</b></p> <p>This should be set by taking into account the large object limit enforced by the Swift cluster in consideration.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A positive integer that is less than or equal to the large object limit enforced by the Swift cluster in consideration.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>swift_store_large_object_chunk_size</b></li> </ul>

Configuration option = Default value	Type	Description
<b>swift_store_multi_tenant</b> <b>= False</b>	boolean value	<p>Store images in tenant's Swift account.</p> <p>This enables multi-tenant storage mode which causes Glance images to be stored in tenant specific Swift accounts. If this is disabled, Glance stores all images in its own account. More details multi-tenant store can be found at <a href="https://wiki.openstack.org/wiki/GlanceSwiftTenantSpecificStorage">https://wiki.openstack.org/wiki/GlanceSwiftTenantSpecificStorage</a></p> <div data-bbox="815 613 922 808" style="float: left; margin-right: 10px;"> </div> <p><b>NOTE</b></p> <p>If using multi-tenant swift store, please make sure that you do not set a swift configuration file with the <i>swift_store_config_file</i> option.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>swift_store_config_file</code></li> </ul>


Configuration option = Default value	Type	Description
<b>swift_store_multiple_containers_seed = 0</b>	integer value	<p>Seed indicating the number of containers to use for storing images.</p> <p>When using a single-tenant store, images can be stored in one or more than one containers. When set to 0, all images will be stored in one single container. When set to an integer value between 1 and 32, multiple containers will be used to store images. This configuration option will determine how many containers are created. The total number of containers that will be used is equal to <math>16^N</math>, so if this config option is set to 2, then <math>16^2=256</math> containers will be used to store images.</p> <p>Please refer to <b>swift_store_container</b> for more detail on the naming convention. More detail about using multiple containers can be found at <a href="https://specs.openstack.org/openstack/glance-specs/specs/kilo/swift-store-multiple-containers.html">https://specs.openstack.org/openstack/glance-specs/specs/kilo/swift-store-multiple-containers.html</a></p> <div data-bbox="815 1003 922 1137" style="display: inline-block; vertical-align: top;">  </div> <p><b>NOTE</b></p> <p>This is used only when <code>swift_store_multi_tenant</code> is disabled.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A non-negative integer less than or equal to 32</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>swift_store_container</b></li> <li>● <b>swift_store_multi_tenant</b></li> <li>● <b>swift_store_create_container_on_put</b></li> </ul>

Configuration option = Default value	Type	Description
<b>swift_store_region = None</b>	string value	<p>The region of Swift endpoint to use by Glance.</p> <p>Provide a string value representing a Swift region where Glance can connect to for image storage. By default, there is no region set.</p> <p>When Glance uses Swift as the storage backend to store images for a specific tenant that has multiple endpoints, setting of a Swift region with <b>swift_store_region</b> allows Glance to connect to Swift in the specified region as opposed to a single region connectivity.</p> <p>This option can be configured for both single-tenant and multi-tenant storage.</p> <div data-bbox="815 819 922 1048" style="float: left; margin-right: 10px;"> </div> <p><b>NOTE</b></p> <p>Setting the region with <b>swift_store_region</b> is tenant-specific and is necessary <b>only if</b> the tenant has multiple endpoints across different regions.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A string value representing a valid Swift region.</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>

Configuration option = Default value	Type	Description
<b>swift_store_retry_get_count = 0</b>	integer value	<p>The number of times a Swift download will be retried before the request fails.</p> <p>Provide an integer value representing the number of times an image download must be retried before erroring out. The default value is zero (no retry on a failed image download). When set to a positive integer value, <b>swift_store_retry_get_count</b> ensures that the download is attempted this many more times upon a download failure before sending an error message.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Zero</li> <li>• Positive integer value</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<b>swift_store_service_type = object-store</b>	string value	<p>Type of Swift service to use.</p> <p>Provide a string value representing the service type to use for storing images while using Swift backend storage. The default service type is set to <b>object-store</b>.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p>If <b>swift_store_auth_version</b> is set to 2, the value for this configuration option needs to be <b>object-store</b>. If using a higher version of Keystone or a different auth scheme, this option may be modified.</p> </div> </div> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A string representing a valid service type for Swift storage.</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>



Configuration option = Default value	Type	Description
<b>swift_store_ssl_compression = True</b>	boolean value	<p>SSL layer compression for HTTPS Swift requests.</p> <p>Provide a boolean value to determine whether or not to compress HTTPS Swift requests for images at the SSL layer. By default, compression is enabled.</p> <p>When using Swift as the backend store for Glance image storage, SSL layer compression of HTTPS Swift requests can be set using this option. If set to False, SSL layer compression of HTTPS Swift requests is disabled. Disabling this option may improve performance for images which are already in a compressed format, for example, qcow2.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>

Configuration option = Default value	Type	Description
<b>swift_store_use_trusts = True</b>	boolean value	<p>Use trusts for multi-tenant Swift store.</p> <p>This option instructs the Swift store to create a trust for each add/get request when the multi-tenant store is in use. Using trusts allows the Swift store to avoid problems that can be caused by an authentication token expiring during the upload or download of data.</p> <p>By default, <b>swift_store_use_trusts</b> is set to <b>True</b>(use of trusts is enabled). If set to <b>False</b>, a user token is used for the Swift connection instead, eliminating the overhead of trust creation.</p> <p> <b>NOTE</b></p> <p>This option is considered only when <b>swift_store_multi_tenant</b> is set to <b>True</b></p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• True</li> <li>• False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• swift_store_multi_tenant</li> </ul>
<b>swift_store_user = None</b>	string value	The user to authenticate against the Swift authentication service.

Configuration option = Default value	Type	Description
<b>swift_upload_buffer_dir = None</b>	string value	<p>Directory to buffer image segments before upload to Swift.</p> <p>Provide a string value representing the absolute path to the directory on the glance node where image segments will be buffered briefly before they are uploaded to swift.</p> <p>NOTES: * This is required only when the configuration option <b>swift_buffer_on_upload</b> is set to True. * This directory should be provisioned keeping in mind the <b>swift_store_large_object_chunk_size</b> and the maximum number of images that could be uploaded simultaneously by a given glance node.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● String value representing an absolute directory path</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>swift_buffer_on_upload</code></li> <li>● <code>swift_store_large_object_chunk_size</code></li> </ul>
<b>vmware_api_retry_count = 10</b>	integer value	<p>The number of VMware API retries.</p> <p>This configuration option specifies the number of times the VMware ESX/VC server API must be retried upon connection related issues or server API call overload. It is not possible to specify <i>retry forever</i>.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>● Any positive integer value</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>

Configuration option = Default value	Type	Description
<b>vmware_ca_file = None</b>	string value	<p>Absolute path to the CA bundle file.</p> <p>This configuration option enables the operator to use a custom Certificate Authority File to verify the ESX/vCenter certificate.</p> <p>If this option is set, the "vmware_insecure" option will be ignored and the CA file specified will be used to authenticate the ESX/vCenter server certificate and establish a secure connection to the server.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>Any string that is a valid absolute path to a CA file</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>vmware_insecure</li> </ul>
<b>vmware_datastores = None</b>	multi valued	<p>The datastores where the image can be stored.</p> <p>This configuration option specifies the datastores where the image can be stored in the VMWare store backend. This option may be specified multiple times for specifying multiple datastores. The datastore name should be specified after its datacenter path, separated by ":". An optional weight may be given after the datastore name, separated again by ":" to specify the priority. Thus, the required format becomes &lt;datacenter_path&gt;:&lt;datastore_name&gt;:&lt;optional_weight&gt;.</p> <p>When adding an image, the datastore with highest weight will be selected, unless there is not enough free space available in cases where the image size is already known. If no weight is given, it is assumed to be zero and the directory will be considered for selection last. If multiple datastores have the same weight, then the one with the most free space available is selected.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>Any string of the format: &lt;datacenter_path&gt;:&lt;datastore_name&gt;:&lt;optional_weight&gt;</li> </ul> <p>Related options: * None</p>

Configuration option = Default value	Type	Description
<b>vmware_insecure = False</b>	boolean value	<p>Set verification of the ESX/vCenter server certificate.</p> <p>This configuration option takes a boolean value to determine whether or not to verify the ESX/vCenter server certificate. If this option is set to True, the ESX/vCenter server certificate is not verified. If this option is set to False, then the default CA truststore is used for verification.</p> <p>This option is ignored if the "vmware_ca_file" option is set. In that case, the ESX/vCenter server certificate will then be verified using the file specified using the "vmware_ca_file" option .</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● vmware_ca_file</li> </ul>
<b>vmware_server_host = None</b>	unknown value	<p>Address of the ESX/ESXi or vCenter Server target system.</p> <p>This configuration option sets the address of the ESX/ESXi or vCenter Server target system. This option is required when using the VMware storage backend. The address can contain an IP address (127.0.0.1) or a DNS name (www.my-domain.com).</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>● A valid IPv4 or IPv6 address</li> <li>● A valid DNS name</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● vmware_server_username</li> <li>● vmware_server_password</li> </ul>

Configuration option = Default value	Type	Description
<b>vmware_server_password = None</b>	string value	<p>Server password.</p> <p>This configuration option takes the password for authenticating with the VMware ESX/ESXi or vCenter Server. This option is required when using the VMware storage backend.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>Any string that is a password corresponding to the username specified using the "vmware_server_username" option</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>vmware_server_host</li> <li>vmware_server_username</li> </ul>
<b>vmware_server_username = None</b>	string value	<p>Server username.</p> <p>This configuration option takes the username for authenticating with the VMware ESX/ESXi or vCenter Server. This option is required when using the VMware storage backend.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>Any string that is the username for a user with appropriate privileges</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>vmware_server_host</li> <li>vmware_server_password</li> </ul>

Configuration option = Default value	Type	Description
<b>vmware_store_image_dir</b> <b>= /openstack_glance</b>	string value	<p>The directory where the glance images will be stored in the datastore.</p> <p>This configuration option specifies the path to the directory where the glance images will be stored in the VMware datastore. If this option is not set, the default directory where the glance images are stored is <code>openstack_glance</code>.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>Any string that is a valid path to a directory</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>None</li> </ul>
<b>vmware_task_poll_interval</b> <b>= 5</b>	integer value	<p>Interval in seconds used for polling remote tasks invoked on VMware ESX/VC server.</p> <p>This configuration option takes in the sleep time in seconds for polling an on-going async task as part of the VMWare ESX/VC server API call.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>Any positive integer value</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>None</li> </ul>

#### 5.3.4. oslo\_concurrency

The following table outlines the options available under the **[oslo\_concurrency]** group in the `/etc/glance/glance-scrubber.conf` file.

Table 5.36. oslo\_concurrency

Configuration option = Default value	Type	Description
<b>disable_process_locking</b> <b>= False</b>	boolean value	Enables or disables inter-process locks.

Configuration option = Default value	Type	Description
<b>lock_path = None</b>	string value	Directory to use for lock files. For security, the specified directory should only be writable by the user running the processes that need locking. Defaults to environment variable OSLO_LOCK_PATH. If external locks are used, a lock path must be set.

### 5.3.5. oslo\_policy

The following table outlines the options available under the **[oslo\_policy]** group in the `/etc/glance/glance-scrubber.conf` file.

Table 5.37. oslo\_policy

Configuration option = Default value	Type	Description
<b>enforce_scope = False</b>	boolean value	This option controls whether or not to enforce scope when evaluating policies. If <b>True</b> , the scope of the token used in the request is compared to the <b>scope_types</b> of the policy being enforced. If the scopes do not match, an <b>InvalidScope</b> exception will be raised. If <b>False</b> , a message will be logged informing operators that policies are being invoked with mismatching scope.
<b>policy_default_rule = default</b>	string value	Default rule. Enforced when a requested rule is not found.
<b>policy_dirs = [u'policy.d']</b>	multi valued	Directories where policy configuration files are stored. They can be relative to any directory in the search path defined by the <code>config_dir</code> option, or absolute paths. The file defined by <code>policy_file</code> must exist for these directories to be searched. Missing or empty directories are ignored.
<b>policy_file = policy.json</b>	string value	The file that defines policies.
<b>remote_content_type = application/x-www-form-urlencoded</b>	string value	Content Type to send and receive data for REST based policy check
<b>remote_ssl_ca_cert_file = None</b>	string value	Absolute path to ca cert file for REST based policy check



Configuration option = Default value	Type	Description
<b>remote_ssl_client_cert_file = None</b>	string value	Absolute path to client cert for REST based policy check
<b>remote_ssl_client_key_file = None</b>	string value	Absolute path client key file REST based policy check
<b>remote_ssl_verify_server_cert = False</b>	boolean value	server identity verification for REST based policy check

## 5.4. GLANCE-CACHE.CONF

This section contains options for the `/etc/glance/glance-cache.conf` file.

### 5.4.1. DEFAULT


The following table outlines the options available under the **[DEFAULT]** group in the `/etc/glance/glance-cache.conf` file.

Table 5.38. DEFAULT


Configuration option = Default value	Type	Description
<b>admin_password = None</b>	string value	The administrators password. If "use_user_token" is not in effect, then admin credentials can be specified.
<b>admin_tenant_name = None</b>	string value	The tenant name of the administrative user. If "use_user_token" is not in effect, then admin tenant name can be specified.
<b>admin_user = None</b>	string value	The administrators user name. If "use_user_token" is not in effect, then admin credentials can be specified.

Configuration option = Default value	Type	Description
<b>allow_additional_image_properties = True</b>	boolean value	<p>Allow users to add additional/custom properties to images.</p> <p>Glance defines a standard set of properties (in its schema) that appear on every image. These properties are also known as <b>base properties</b>. In addition to these properties, Glance allows users to add custom properties to images. These are known as <b>additional properties</b>.</p> <p>By default, this configuration option is set to <b>True</b> and users are allowed to add additional properties. The number of additional properties that can be added to an image can be controlled via <b>image_property_quota</b> configuration option.</p> <p>Possible values:</p> <ul style="list-style-type: none"><li>• True</li><li>• False</li></ul> <p>Related options:</p> <ul style="list-style-type: none"><li>• image_property_quota</li></ul>

Configuration option = Default value	Type	Description
<b>api_limit_max = 1000</b>	integer value	<p>Maximum number of results that could be returned by a request.</p> <p>As described in the help text of <b>limit_param_default</b>, some requests may return multiple results. The number of results to be returned are governed either by the <b>limit</b> parameter in the request or the <b>limit_param_default</b> configuration option. The value in either case, can't be greater than the absolute maximum defined by this configuration option. Anything greater than this value is trimmed down to the maximum value defined here.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p>Setting this to a very large value may slow down database queries and increase response times. Setting this to a very low value may result in poor user experience.</p> </div> </div> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any positive integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>limit_param_default</code></li> </ul>
<b>auth_region = None</b>	string value	The region for the authentication service. If "use_user_token" is not in effect and using keystone auth, then region name can be specified.
<b>auth_strategy = noauth</b>	string value	The strategy to use for authentication. If "use_user_token" is not in effect, then auth strategy can be specified.
<b>auth_url = None</b>	string value	The URL to the keystone service. If "use_user_token" is not in effect and using keystone auth, then URL of keystone can be specified.

Configuration option = Default value	Type	Description
<b>data_api = glance.db.sqlalchemy.api</b>	string value	<p>Python module path of data access API.</p> <p>Specifies the path to the API to use for accessing the data model. This option determines how the image catalog data will be accessed.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● <code>glance.db.sqlalchemy.api</code></li> <li>● <code>glance.db.registry.api</code></li> <li>● <code>glance.db.simple.api</code></li> </ul> <p>If this option is set to <b>glance.db.sqlalchemy.api</b> then the image catalog data is stored in and read from the database via the SQLAlchemy Core and ORM APIs.</p> <p>Setting this option to <b>glance.db.registry.api</b> will force all database access requests to be routed through the Registry service. This avoids data access from the Glance API nodes for an added layer of security, scalability and manageability.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p>In v2 OpenStack Images API, the registry service is optional. In order to use the Registry API in v2, the option <b>enable_v2_registry</b> must be set to <b>True</b>.</p> </div> </div> <p>Finally, when this configuration option is set to <b>glance.db.simple.api</b>, image catalog data is stored in and read from an in-memory data structure. This is primarily used for testing.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>enable_v2_api</code></li> <li>● <code>enable_v2_registry</code></li> </ul>
<b>debug = False</b>	boolean value	If set to true, the logging level will be set to DEBUG instead of the default INFO level.

Configuration option = Default value	Type	Description
<b>default_log_levels =</b> <b>[u'amqp=WARN',</b> <b>u'amqplib=WARN',</b> <b>u'boto=WARN',</b> <b>u'qpid=WARN',</b> <b>u'sqlalchemy=WARN',</b> <b>u'suds=INFO',</b> <b>u'oslo.messaging=INFO',</b> <b>u'oslo_messaging=INFO',</b> <b>u'iso8601=WARN',</b> <b>u'requests.packages.urlli</b> <b>b3.connectionpool=WAR</b> <b>N',</b> <b>u'urllib3.connectionpool=</b> <b>WARN',</b> <b>u'websocket=WARN',</b> <b>u'requests.packages.urlli</b> <b>b3.util.retry=WARN',</b> <b>u'urllib3.util.retry=WARN',</b> <b>u'keystonemiddleware=W</b> <b>ARN',</b> <b>u'routes.middleware=WA</b> <b>RN', u'stevedore=WARN',</b> <b>u'taskflow=WARN',</b> <b>u'keystoneauth=WARN',</b> <b>u'oslo.cache=INFO',</b> <b>u'dogpile.core.dogpile=IN</b> <b>FO']</b>	list value	List of package logging levels in logger=LEVEL pairs. This option is ignored if log_config_append is set.

Configuration option = Default value	Type	Description
<b>digest_algorithm = sha256</b>	string value	<p>Digest algorithm to use for digital signature.</p> <p>Provide a string value representing the digest algorithm to use for generating digital signatures. By default, <b>sha256</b> is used.</p> <p>To get a list of the available algorithms supported by the version of OpenSSL on your platform, run the command: <b>openssl list-message-digest-algorithms</b>. Examples are <i>sha1</i>, <i>sha256</i>, and <i>sha512</i>.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p><b>digest_algorithm</b> is not related to Glance's image signing and verification. It is only used to sign the universally unique identifier (UUID) as a part of the certificate file and key file validation.</p> </div> </div> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● An OpenSSL message digest algorithm identifier</li> </ul> <p>Relation options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>enable_image_import = True</b>	boolean value	<p>Enables the Image Import workflow introduced in Pike</p> <p>As <code>[DEFAULT]/node_staging_uri</code> is required for the Image Import, it's disabled per default in Pike, enabled per default in Queens and removed in Rocky. This allows Glance to operate with previous version configs upon upgrade.</p> <p>Setting this option to False will disable the endpoints related to Image Import Refactoring work.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>[DEFAULT]/node_staging_uri</code></li> </ul>

Configuration option = Default value	Type	Description
<b>enable_v1_api = True</b>	boolean value	<p>Deploy the v1 OpenStack Images API.</p> <p>When this option is set to <b>True</b>, Glance service will respond to requests on registered endpoints conforming to the v1 OpenStack Images API.</p> <p>NOTES:</p> <ul style="list-style-type: none"> <li>● If this option is enabled, then <b>enable_v1_registry</b> must also be set to <b>True</b> to enable mandatory usage of Registry service with v1 API.</li> <li>● If this option is disabled, then the <b>enable_v1_registry</b> option, which is enabled by default, is also recommended to be disabled.</li> <li>● This option is separate from <b>enable_v2_api</b>, both v1 and v2 OpenStack Images API can be deployed independent of each other.</li> <li>● If deploying only the v2 Images API, this option, which is enabled by default, should be disabled.</li> </ul> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● enable_v1_registry</li> <li>● enable_v2_api</li> </ul>

Configuration option = Default value	Type	Description
<b>enable_v1_registry = True</b>	boolean value	<p>Deploy the v1 API Registry service.</p> <p>When this option is set to <b>True</b>, the Registry service will be enabled in Glance for v1 API requests.</p> <p>NOTES:</p> <ul style="list-style-type: none"><li>● Use of Registry is mandatory in v1 API, so this option must be set to <b>True</b> if the <b>enable_v1_api</b> option is enabled.</li><li>● If deploying only the v2 OpenStack Images API, this option, which is enabled by default, should be disabled.</li></ul> <p>Possible values:</p> <ul style="list-style-type: none"><li>● True</li><li>● False</li></ul> <p>Related options:</p> <ul style="list-style-type: none"><li>● enable_v1_api</li></ul>




Configuration option = Default value	Type	Description
<b>enable_v2_api = True</b>	boolean value	<p>Deploy the v2 OpenStack Images API.</p> <p>When this option is set to <b>True</b>, Glance service will respond to requests on registered endpoints conforming to the v2 OpenStack Images API.</p> <p>NOTES:</p> <ul style="list-style-type: none"> <li>● If this option is disabled, then the <b>enable_v2_registry</b> option, which is enabled by default, is also recommended to be disabled.</li> <li>● This option is separate from <b>enable_v1_api</b>, both v1 and v2 OpenStack Images API can be deployed independent of each other.</li> <li>● If deploying only the v1 Images API, this option, which is enabled by default, should be disabled.</li> </ul> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● enable_v2_registry</li> <li>● enable_v1_api</li> </ul>


Configuration option = Default value	Type	Description
<b>enable_v2_registry = True</b>	boolean value	<p>Deploy the v2 API Registry service.</p> <p>When this option is set to <b>True</b>, the Registry service will be enabled in Glance for v2 API requests.</p> <p>NOTES:</p> <ul style="list-style-type: none"> <li>● Use of Registry is optional in v2 API, so this option must only be enabled if both <b>enable_v2_api</b> is set to <b>True</b> and the <b>data_api</b> option is set to <b>glance.db.registry.api</b>.</li> <li>● If deploying only the v1 OpenStack Images API, this option, which is enabled by default, should be disabled.</li> </ul> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● enable_v2_api</li> <li>● data_api</li> </ul>
<b>enabled_import_methods = [u'glance-direct', u'web-download']</b>	list value	<p>List of enabled Image Import Methods</p> <p>Both <i>glance-direct</i> and <i>web-download</i> are enabled by default.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● [DEFAULT]/node_staging_uri</li> <li>● [DEFAULT]/enable_image_import</li> </ul>
<b>fatal_deprecations = False</b>	boolean value	Enables or disables fatal status of deprecations.

Configuration option = Default value	Type	Description
<b>image_cache_dir = None</b>	string value	<p>Base directory for image cache.</p> <p>This is the location where image data is cached and served out of. All cached images are stored directly under this directory. This directory also contains three subdirectories, namely, <b>incomplete</b>, <b>invalid</b> and <b>queue</b>.</p> <p>The <b>incomplete</b> subdirectory is the staging area for downloading images. An image is first downloaded to this directory. When the image download is successful it is moved to the base directory. However, if the download fails, the partially downloaded image file is moved to the <b>invalid</b> subdirectory.</p> <p>The <b>queue`subdirectory is used for queuing images for download. This is used primarily by the cache-prefetcher, which can be scheduled as a periodic task like cache-pruner and cache-cleaner, to cache images ahead of their usage. Upon receiving the request to cache an image, Glance touches a file in the `queue directory with the image id as the file name. The cache-prefetcher, when running, polls for the files in <b>queue</b> directory and starts downloading them in the order they were created. When the download is successful, the zero-sized file is deleted from the <b>queue</b> directory. If the download fails, the zero-sized file remains and it'll be retried the next time cache-prefetcher runs.</b></p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid path</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>image_cache_sqlite_db</b></li> </ul>

Configuration option = Default value	Type	Description
<b>image_cache_driver = sqlite</b>	string value	<p>The driver to use for image cache management.</p> <p>This configuration option provides the flexibility to choose between the different image-cache drivers available. An image-cache driver is responsible for providing the essential functions of image-cache like write images to/read images from cache, track age and usage of cached images, provide a list of cached images, fetch size of the cache, queue images for caching and clean up the cache, etc.</p> <p>The essential functions of a driver are defined in the base class <b>glance.image_cache.drivers.base.Driver</b>. All image-cache drivers (existing and prospective) must implement this interface. Currently available drivers are <b>sqlite</b> and <b>xattr</b>. These drivers primarily differ in the way they store the information about cached images:</p> <ul style="list-style-type: none"> <li>• The <b>sqlite</b> driver uses a sqlite database (which sits on every glance node locally) to track the usage of cached images.</li> <li>• The <b>xattr</b> driver uses the extended attributes of files to store this information. It also requires a filesystem that sets <b>atime</b> on the files when accessed.</li> </ul> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• sqlite</li> <li>• xattr</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>

Configuration option = Default value	Type	Description
<b>image_cache_max_size = 10737418240</b>	integer value	<p>The upper limit on cache size, in bytes, after which the cache-pruner cleans up the image cache.</p> <div data-bbox="815 371 922 779" style="border: 1px solid black; padding: 5px; width: fit-content;">  </div> <p><b>NOTE</b></p> <p>This is just a threshold for cache-pruner to act upon. It is NOT a hard limit beyond which the image cache would never grow. In fact, depending on how often the cache-pruner runs and how quickly the cache fills, the image cache can far exceed the size specified here very easily. Hence, care must be taken to appropriately schedule the cache-pruner and in setting this limit.</p> <p>Glance caches an image when it is downloaded. Consequently, the size of the image cache grows over time as the number of downloads increases. To keep the cache size from becoming unmanageable, it is recommended to run the cache-pruner as a periodic task. When the cache pruner is kicked off, it compares the current size of image cache and triggers a cleanup if the image cache grew beyond the size specified here. After the cleanup, the size of cache is less than or equal to size specified here.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any non-negative integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>

Configuration option = Default value	Type	Description
<b>image_cache_sqlite_db = cache.db</b>	string value	<p>The relative path to sqlite file database that will be used for image cache management.</p> <p>This is a relative path to the sqlite file database that tracks the age and usage statistics of image cache. The path is relative to image cache base directory, specified by the configuration option <b>image_cache_dir</b>.</p> <p>This is a lightweight database with just one table.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A valid relative path to sqlite file database</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <b>image_cache_dir</b></li> </ul>
<b>image_cache_stall_time = 86400</b>	integer value	<p>The amount of time, in seconds, an incomplete image remains in the cache.</p> <p>Incomplete images are images for which download is in progress. Please see the description of configuration option <b>image_cache_dir</b> for more detail. Sometimes, due to various reasons, it is possible the download may hang and the incompletely downloaded image remains in the <b>incomplete</b> directory. This configuration option sets a time limit on how long the incomplete images should remain in the <b>incomplete</b> directory before they are cleaned up. Once an incomplete image spends more time than is specified here, it'll be removed by cache-cleaner on its next run.</p> <p>It is recommended to run cache-cleaner as a periodic task on the Glance API nodes to keep the incomplete images from occupying disk space.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Any non-negative integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>


Configuration option = Default value	Type	Description
<b>image_location_quota = 10</b>	integer value	<p>Maximum number of locations allowed on an image.</p> <p>Any negative value is interpreted as unlimited.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<b>image_member_quota = 128</b>	integer value	<p>Maximum number of image members per image.</p> <p>This limits the maximum of users an image can be shared with. Any negative value is interpreted as unlimited.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<b>image_property_quota = 128</b>	integer value	<p>Maximum number of properties allowed on an image.</p> <p>This enforces an upper limit on the number of additional properties an image can have. Any negative value is interpreted as unlimited.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p>This won't have any impact if additional properties are disabled. Please refer to <b>allow_additional_image_properties</b>.</p> </div> </div> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <b>allow_additional_image_properties</b></li> </ul>




Configuration option = Default value	Type	Description
<b>image_size_cap = 1099511627776</b>	integer value	<p>Maximum size of image a user can upload in bytes.</p> <p>An image upload greater than the size mentioned here would result in an image creation failure. This configuration option defaults to 1099511627776 bytes (1 TiB).</p> <p>NOTES:</p> <ul style="list-style-type: none"> <li>• This value should only be increased after careful consideration and must be set less than or equal to 8 EiB (9223372036854775808).</li> <li>• This value must be set with careful consideration of the backend storage capacity. Setting this to a very low value may result in a large number of image failures. And, setting this to a very large value may result in faster consumption of storage. Hence, this must be set according to the nature of images created and storage capacity available.</li> </ul> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Any positive number less than or equal to 9223372036854775808</li> </ul>
<b>image_tag_quota = 128</b>	integer value	<p>Maximum number of tags allowed on an image.</p> <p>Any negative value is interpreted as unlimited.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<code>`instance_format = [instance: %(uuid)s] `</code>	string value	The format for an instance that is passed with the log message.
<code>`instance_uuid_format = [instance: %(uuid)s] `</code>	string value	The format for an instance UUID that is passed with the log message.




Configuration option = Default value	Type	Description
<b>limit_param_default = 25</b>	integer value	<p>The default number of results to return for a request.</p> <p>Responses to certain API requests, like list images, may return multiple items. The number of results returned can be explicitly controlled by specifying the <b>limit</b> parameter in the API request. However, if a <b>limit</b> parameter is not specified, this configuration value will be used as the default number of results to be returned for any API request.</p> <p>NOTES:</p> <ul style="list-style-type: none"> <li>• The value of this configuration option may not be greater than the value specified by <b>api_limit_max</b>.</li> <li>• Setting this to a very large value may slow down database queries and increase response times. Setting this to a very low value may result in poor user experience.</li> </ul> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Any positive integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <code>api_limit_max</code></li> </ul>
<b>log-config-append = None</b>	string value	<p>The name of a logging configuration file. This file is appended to any existing logging configuration files. For details about logging configuration files, see the Python logging module documentation. Note that when logging configuration files are used then all logging configuration is set in the configuration file and other logging configuration options are ignored (for example, <code>logging_context_format_string</code>).</p>
<b>log-date-format = %Y-%m-%d %H:%M:%S</b>	string value	<p>Defines the format string for <code>%(asctime)s</code> in log records. Default: <code>%(default)s</code>. This option is ignored if <code>log_config_append</code> is set.</p>
<b>log-dir = None</b>	string value	<p>(Optional) The base directory used for relative <code>log_file</code> paths. This option is ignored if <code>log_config_append</code> is set.</p>
<b>log-file = None</b>	string value	<p>(Optional) Name of log file to send logging output to. If no default is set, logging will go to <code>stderr</code> as defined by <code>use_stderr</code>. This option is ignored if <code>log_config_append</code> is set.</p>


Configuration option = Default value	Type	Description
<b>logging_context_format_string</b> = <code>%(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [%(request_id)s %(user_identity)s] %(instance)s%(message)s</code>	string value	Format string to use for log messages with context.
<b>logging_debug_format_s uffix</b> = <code>%(funcName)s %(pathname)s:%(lineno)d</code>	string value	Additional data to append to log message when logging level for the message is DEBUG.
<b>logging_default_format_s tring</b> = <code>%(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [-] %(instance)s%(message)s</code>	string value	Format string to use for log messages when context is undefined.
<b>logging_exception_prefix</b> = <code>%(asctime)s.%(msecs)03d %(process)d ERROR %(name)s %(instance)s</code>	string value	Prefix each line of exception output with this format.
<b>logging_user_identity_for mat</b> = <code>%(user)s %(tenant)s %(domain)s %(user_domain)s %(project_domain)s</code>	string value	Defines the format string for <code>%(user_identity)s</code> that is used in <code>logging_context_format_string</code> .


Configuration option = Default value	Type	Description
<b>metadata_encryption_key</b> = None	string value	<p>AES key for encrypting store location metadata.</p> <p>Provide a string value representing the AES cipher to use for encrypting Glance store metadata.</p> <p> <b>NOTE</b></p> <p>The AES key to use must be set to a random string of length 16, 24 or 32 bytes.</p> <p>Possible values:</p> <ul style="list-style-type: none"><li>• String value representing a valid AES key</li></ul> <p>Related options:</p> <ul style="list-style-type: none"><li>• None</li></ul>

Configuration option = Default value	Type	Description
<b>node_staging_uri = file:///tmp/staging/</b>	string value	<p>The URL provides location where the temporary data will be stored</p> <p>This option is for Glance internal use only. Glance will save the image data uploaded by the user to <i>staging</i> endpoint during the image import process.</p> <p>This option does not change the <i>staging</i> API endpoint by any means.</p> <p> <b>NOTE</b></p> <p>It is discouraged to use same path as [task]/work_dir</p> <p> <b>NOTE</b></p> <p><i>file://&lt;absolute-directory-path&gt;</i> is the only option api_image_import flow will support for now.</p> <p> <b>NOTE</b></p> <p>The staging path must be on shared filesystem available to all Glance API nodes.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● String starting with <i>file://</i> followed by absolute FS path</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● [task]/work_dir</li> <li>● [DEFAULT]/enable_image_import (<b>deprecated</b>)</li> </ul>
<b>publish_errors = False</b>	boolean value	Enables or disables publication of error events.

Configuration option = Default value	Type	Description
<b>pydev_worker_debug_host = None</b>	unknown value	<p>Host address of the pydev server.</p> <p>Provide a string value representing the hostname or IP of the pydev server to use for debugging. The pydev server listens for debug connections on this address, facilitating remote debugging in Glance.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Valid hostname</li> <li>● Valid IP address</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>pydev_worker_debug_port = 5678</b>	port value	<p>Port number that the pydev server will listen on.</p> <p>Provide a port number to bind the pydev server to. The pydev process accepts debug connections on this port and facilitates remote debugging in Glance.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid port number</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>rate_limit_burst = 0</b>	integer value	Maximum number of logged messages per <code>rate_limit_interval</code> .
<b>rate_limit_except_level = CRITICAL</b>	string value	Log level name used by rate limiting: CRITICAL, ERROR, INFO, WARNING, DEBUG or empty string. Logs with level greater or equal to <code>rate_limit_except_level</code> are not filtered. An empty string means that all levels are filtered.
<b>rate_limit_interval = 0</b>	integer value	Interval, number of seconds, of log rate limiting.

Configuration option = Default value	Type	Description
<b>registry_client_ca_file = None</b>	string value	<p data-bbox="815 255 1342 286">Absolute path to the Certificate Authority file.</p> <p data-bbox="815 322 1401 456">Provide a string value representing a valid absolute path to the certificate authority file to use for establishing a secure connection to the registry server.</p> <div data-bbox="815 510 922 857" style="border: 1px solid black; padding: 5px; width: fit-content;">  </div> <p data-bbox="1002 517 1091 548"><b>NOTE</b></p> <p data-bbox="1002 584 1426 857">This option must be set if <b>registry_client_protocol</b> is set to <b>https</b>. Alternatively, the <code>GLANCE_CLIENT_CA_FILE</code> environment variable may be set to a filepath of the CA file. This option is ignored if the <b>registry_client_insecure</b> option is set to <b>True</b>.</p> <p data-bbox="815 913 999 945">Possible values:</p> <ul data-bbox="882 978 1394 1037" style="list-style-type: none"> <li>● String value representing a valid absolute path to the CA file.</li> </ul> <p data-bbox="815 1070 1007 1102">Related options:</p> <ul data-bbox="882 1135 1193 1227" style="list-style-type: none"> <li>● <code>registry_client_protocol</code></li> <li>● <code>registry_client_insecure</code></li> </ul>

Configuration option = Default value	Type	Description
<b>registry_client_cert_file = None</b>	string value	<p>Absolute path to the certificate file.</p> <p>Provide a string value representing a valid absolute path to the certificate file to use for establishing a secure connection to the registry server.</p> <div data-bbox="815 472 922 730" style="display: inline-block; vertical-align: top;">  </div> <p><b>NOTE</b></p> <p>This option must be set if <b>registry_client_protocol</b> is set to <b>https</b>. Alternatively, the <code>GLANCE_CLIENT_CERT_FILE</code> environment variable may be set to a filepath of the certificate file.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● String value representing a valid absolute path to the certificate file.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>registry_client_protocol</code></li> </ul>
<b>registry_client_insecure = False</b>	boolean value	<p>Set verification of the registry server certificate.</p> <p>Provide a boolean value to determine whether or not to validate SSL connections to the registry server. By default, this option is set to <b>False</b> and the SSL connections are validated.</p> <p>If set to <b>True</b>, the connection to the registry server is not validated via a certifying authority and the <b>registry_client_ca_file</b> option is ignored. This is the registry's equivalent of specifying <code>--insecure</code> on the command line using <code>glanceclient</code> for the API.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>registry_client_protocol</code></li> <li>● <code>registry_client_ca_file</code></li> </ul>

Configuration option = Default value	Type	Description
<b>registry_client_key_file = None</b>	string value	<p>Absolute path to the private key file.</p> <p>Provide a string value representing a valid absolute path to the private key file to use for establishing a secure connection to the registry server.</p> <p> <b>NOTE</b></p> <p>This option must be set if <b>registry_client_protocol</b> is set to <b>https</b>. Alternatively, the <code>GLANCE_CLIENT_KEY_FILE</code> environment variable may be set to a filepath of the key file.</p> <p>Possible values:</p> <ul style="list-style-type: none"><li>● String value representing a valid absolute path to the key file.</li></ul> <p>Related options:</p> <ul style="list-style-type: none"><li>● <code>registry_client_protocol</code></li></ul>



Configuration option = Default value	Type	Description
<b>registry_client_protocol = http</b>	string value	<p>Protocol to use for communication with the registry server.</p> <p>Provide a string value representing the protocol to use for communication with the registry server. By default, this option is set to <b>http</b> and the connection is not secure.</p> <p>This option can be set to <b>https</b> to establish a secure connection to the registry server. In this case, provide a key to use for the SSL connection using the <b>registry_client_key_file</b> option. Also include the CA file and cert file using the options <b>registry_client_ca_file</b> and <b>registry_client_cert_file</b> respectively.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● http</li> <li>● https</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● registry_client_key_file</li> <li>● registry_client_cert_file</li> <li>● registry_client_ca_file</li> </ul>
<b>registry_client_timeout = 600</b>	integer value	<p>Timeout value for registry requests.</p> <p>Provide an integer value representing the period of time in seconds that the API server will wait for a registry request to complete. The default value is 600 seconds.</p> <p>A value of 0 implies that a request will never timeout.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Zero</li> <li>● Positive integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>

Configuration option = Default value	Type	Description
<b>registry_host = 0.0.0.0</b>	unknown value	Address the registry server is hosted on.  Possible values: <ul style="list-style-type: none"><li>• A valid IP or hostname</li></ul> Related options: <ul style="list-style-type: none"><li>• None</li></ul>
<b>registry_port = 9191</b>	port value	Port the registry server is listening on.  Possible values: <ul style="list-style-type: none"><li>• A valid port number</li></ul> Related options: <ul style="list-style-type: none"><li>• None</li></ul>

Configuration option = Default value	Type	Description
<b>show_image_direct_url = False</b>	boolean value	<p>Show direct image location when returning an image.</p> <p>This configuration option indicates whether to show the direct image location when returning image details to the user. The direct image location is where the image data is stored in backend storage. This image location is shown under the image property <b>direct_url</b>.</p> <p>When multiple image locations exist for an image, the best location is displayed based on the location strategy indicated by the configuration option <b>location_strategy</b>.</p> <p>NOTES:</p> <ul style="list-style-type: none"> <li>● Revealing image locations can present a GRAVE SECURITY RISK as image locations can sometimes include credentials. Hence, this is set to <b>False</b> by default. Set this to <b>True</b> with EXTREME CAUTION and ONLY IF you know what you are doing!</li> <li>● If an operator wishes to avoid showing any image location(s) to the user, then both this option and <b>show_multiple_locations</b> MUST be set to <b>False</b>.</li> </ul> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● show_multiple_locations</li> <li>● location_strategy</li> </ul>

Configuration option = Default value	Type	Description
<b>show_multiple_locations = False</b>	boolean value	<p>Show all image locations when returning an image.</p> <p>This configuration option indicates whether to show all the image locations when returning image details to the user. When multiple image locations exist for an image, the locations are ordered based on the location strategy indicated by the configuration opt <b>location_strategy</b>. The image locations are shown under the image property <b>locations</b>.</p> <p>NOTES:</p> <ul style="list-style-type: none"> <li>● Revealing image locations can present a GRAVE SECURITY RISK as image locations can sometimes include credentials. Hence, this is set to <b>False</b> by default. Set this to <b>True</b> with EXTREME CAUTION and ONLY IF you know what you are doing!</li> <li>● See <a href="https://wiki.openstack.org/wiki/OSSN/OSSN-0065">https://wiki.openstack.org/wiki/OSSN/OSSN-0065</a> for more information.</li> <li>● If an operator wishes to avoid showing any image location(s) to the user, then both this option and <b>show_image_direct_url</b> MUST be set to <b>False</b>.</li> </ul> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● show_image_direct_url</li> <li>● location_strategy</li> </ul>
<b>syslog-log-facility = LOG_USER</b>	string value	Syslog facility to receive log lines. This option is ignored if log_config_append is set.
<b>use-journal = False</b>	boolean value	Enable journald for logging. If running in a systemd environment you may wish to enable journal support. Doing so will use the journal native protocol which includes structured metadata in addition to log messages. This option is ignored if log_config_append is set.
<b>use-json = False</b>	boolean value	Use JSON formatting for logging. This option is ignored if log_config_append is set.

Configuration option = Default value	Type	Description
<b>use-syslog = False</b>	boolean value	Use syslog for logging. Existing syslog format is DEPRECATED and will be changed later to honor RFC5424. This option is ignored if log_config_append is set.
<b>use_stderr = False</b>	boolean value	Log output to standard error. This option is ignored if log_config_append is set.
<b>use_user_token = True</b>	boolean value	Whether to pass through the user token when making requests to the registry. To prevent failures with token expiration during big files upload, it is recommended to set this parameter to False. If "use_user_token" is not in effect, then admin credentials can be specified.
<b>user_storage_quota = 0</b>	string value	<p>Maximum amount of image storage per tenant.</p> <p>This enforces an upper limit on the cumulative storage consumed by all images of a tenant across all stores. This is a per-tenant limit.</p> <p>The default unit for this configuration option is Bytes. However, storage units can be specified using case-sensitive literals <b>B</b>, <b>KB</b>, <b>MB</b>, <b>GB</b> and <b>TB</b> representing Bytes, KiloBytes, MegaBytes, GigaBytes and TeraBytes respectively. Note that there should not be any space between the value and unit. Value <b>0</b> signifies no quota enforcement. Negative values are invalid and result in errors.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A string that is a valid concatenation of a non-negative integer representing the storage value and an optional string literal representing storage units as mentioned above.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<b>watch-log-file = False</b>	boolean value	Uses logging handler designed to watch file system. When log file is moved or removed this handler will open a new log file with specified path instantaneously. It makes sense only if log_file option is specified and Linux platform is used. This option is ignored if log_config_append is set.

## 5.4.2. glance\_store

The following table outlines the options available under the **[glance\_store]** group in the `/etc/glance/glance-cache.conf` file.

Table 5.39. glance\_store

Configuration option = Default value	Type	Description
<b>cinder_api_insecure = False</b>	boolean value	<p>Allow to perform insecure SSL requests to cinder.</p> <p>If this option is set to True, HTTPS endpoint connection is verified using the CA certificates file specified by <b>cinder_ca_certificates_file</b> option.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• True</li> <li>• False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• cinder_ca_certificates_file</li> </ul>
<b>cinder_ca_certificates_file = None</b>	string value	<p>Location of a CA certificates file used for cinder client requests.</p> <p>The specified CA certificates file, if set, is used to verify cinder connections via HTTPS endpoint. If the endpoint is HTTP, this value is ignored.</p> <p><b>cinder_api_insecure</b> must be set to <b>True</b> to enable the verification.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Path to a ca certificates file</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• cinder_api_insecure</li> </ul>

Configuration option = Default value	Type	Description
<b>cinder_catalog_info = volume2::publicURL</b>	string value	<p>Information to match when looking for cinder in the service catalog.</p> <p>When the <b>cinder_endpoint_template</b> is not set and any of <b>cinder_store_auth_address</b>, <b>cinder_store_user_name</b>, <b>cinder_store_project_name</b>, <b>cinder_store_password</b> is not set, cinder store uses this information to lookup cinder endpoint from the service catalog in the current context. <b>cinder_os_region_name</b>, if set, is taken into consideration to fetch the appropriate endpoint.</p> <p>The service catalog can be listed by the <b>openstack catalog list</b> command.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A string of of the following form: <b>&lt;service_type&gt;:&lt;service_name&gt;: &lt;interface&gt;</b> At least <b>service_type</b> and <b>interface</b> should be specified. <b>service_name</b> can be omitted.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● cinder_os_region_name</li> <li>● cinder_endpoint_template</li> <li>● cinder_store_auth_address</li> <li>● cinder_store_user_name</li> <li>● cinder_store_project_name</li> <li>● cinder_store_password</li> </ul>

Configuration option = Default value	Type	Description
<b>cinder_endpoint_template = None</b>	string value	<p>Override service catalog lookup with template for cinder endpoint.</p> <p>When this option is set, this value is used to generate cinder endpoint, instead of looking up from the service catalog. This value is ignored if <b>cinder_store_auth_address</b>, <b>cinder_store_user_name</b>, <b>cinder_store_project_name</b>, and <b>cinder_store_password</b> are specified.</p> <p>If this configuration option is set, <b>cinder_catalog_info</b> will be ignored.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● URL template string for cinder endpoint, where <b>%%(tenant)s</b> is replaced with the current tenant (project) name. For example: <b>http://cinder.openstack.example.org/v2/%%(tenant)s</b></li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● cinder_store_auth_address</li> <li>● cinder_store_user_name</li> <li>● cinder_store_project_name</li> <li>● cinder_store_password</li> <li>● cinder_catalog_info</li> </ul>
<b>cinder_http_retries = 3</b>	integer value	<p>Number of cinderclient retries on failed http calls.</p> <p>When a call failed by any errors, cinderclient will retry the call up to the specified times after sleeping a few seconds.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A positive integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>




Configuration option = Default value	Type	Description
<b>cinder_os_region_name = None</b>	string value	<p>Region name to lookup cinder service from the service catalog.</p> <p>This is used only when <b>cinder_catalog_info</b> is used for determining the endpoint. If set, the lookup for cinder endpoint by this node is filtered to the specified region. It is useful when multiple regions are listed in the catalog. If this is not set, the endpoint is looked up from every region.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A string that is a valid region name.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>cinder_catalog_info</code></li> </ul>
<b>cinder_state_transition_timeout = 300</b>	integer value	<p>Time period, in seconds, to wait for a cinder volume transition to complete.</p> <p>When the cinder volume is created, deleted, or attached to the glance node to read/write the volume data, the volume's state is changed. For example, the newly created volume status changes from <b>creating</b> to <b>available</b> after the creation process is completed. This specifies the maximum time to wait for the status change. If a timeout occurs while waiting, or the status is changed to an unexpected value (e.g. <b>error</b>), the image creation fails.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A positive integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>


Configuration option = Default value	Type	Description
<b>cinder_store_auth_addresses = None</b>	string value	<p>The address where the cinder authentication service is listening.</p> <p>When all of <b>cinder_store_auth_address</b>, <b>cinder_store_user_name</b>, <b>cinder_store_project_name</b>, and <b>cinder_store_password</b> options are specified, the specified values are always used for the authentication. This is useful to hide the image volumes from users by storing them in a project/tenant specific to the image service. It also enables users to share the image volume among other projects under the control of glance's ACL.</p> <p>If either of these options are not set, the cinder endpoint is looked up from the service catalog, and current context's user and project are used.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid authentication service address, for example: <b>http://openstack.example.org/identity/v2.0</b></li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● cinder_store_user_name</li> <li>● cinder_store_password</li> <li>● cinder_store_project_name</li> </ul>
<b>cinder_store_password = None</b>	string value	<p>Password for the user authenticating against cinder.</p> <p>This must be used with all the following related options. If any of these are not specified, the user of the current context is used.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid password for the user specified by <b>cinder_store_user_name</b></li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● cinder_store_auth_address</li> <li>● cinder_store_user_name</li> <li>● cinder_store_project_name</li> </ul>

Configuration option = Default value	Type	Description
<b>cinder_store_project_name = None</b>	string value	<p>Project name where the image volume is stored in cinder.</p> <p>If this configuration option is not set, the project in current context is used.</p> <p>This must be used with all the following related options. If any of these are not specified, the project of the current context is used.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid project name</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>cinder_store_auth_address</b></li> <li>● <b>cinder_store_user_name</b></li> <li>● <b>cinder_store_password</b></li> </ul>
<b>cinder_store_user_name = None</b>	string value	<p>User name to authenticate against cinder.</p> <p>This must be used with all the following related options. If any of these are not specified, the user of the current context is used.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid user name</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● cinder_store_auth_address</li> <li>● cinder_store_password</li> <li>● cinder_store_project_name</li> </ul>

Configuration option = Default value	Type	Description
<b>cinder_volume_type = None</b>	string value	<p>Volume type that will be used for volume creation in cinder.</p> <p>Some cinder backends can have several volume types to optimize storage usage. Adding this option allows an operator to choose a specific volume type in cinder that can be optimized for images.</p> <p>If this is not set, then the default volume type specified in the cinder configuration will be used for volume creation.</p> <p>Possible values:</p> <ul style="list-style-type: none"><li>• A valid volume type from cinder</li></ul> <p>Related options:</p> <ul style="list-style-type: none"><li>• None</li></ul>

Configuration option = Default value	Type	Description
<b>default_store = file</b>	string value	<p>The default scheme to use for storing images.</p> <p>Provide a string value representing the default scheme to use for storing images. If not set, Glance uses <b>file</b> as the default scheme to store images with the <b>file</b> store.</p> <div data-bbox="815 510 922 703" style="display: inline-block; vertical-align: top;">  </div> <p><b>NOTE</b></p> <p>The value given for this configuration option must be a valid scheme for a store registered with the <b>stores</b> configuration option.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● file</li> <li>● filesystem</li> <li>● http</li> <li>● https</li> <li>● swift</li> <li>● swift+http</li> <li>● swift+https</li> <li>● swift+config</li> <li>● rbd</li> <li>● sheepdog</li> <li>● cinder</li> <li>● vsphere</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● stores</li> </ul>

Configuration option = Default value	Type	Description
<b>default_swift_reference = ref1</b>	string value	<p>Reference to default Swift account/backing store parameters.</p> <p>Provide a string value representing a reference to the default set of parameters required for using swift account/backing store for image storage. The default reference value for this configuration option is <i>ref1</i>. This configuration option dereferences the parameters and facilitates image storage in Swift storage backend every time a new image is added.</p> <p>Possible values:</p> <ul style="list-style-type: none"><li>• A valid string value</li></ul> <p>Related options:</p> <ul style="list-style-type: none"><li>• None</li></ul>

Configuration option = Default value	Type	Description
<b>filesystem_store_datadir</b> = <b>/var/lib/glance/images</b>	string value	<p>Directory to which the filesystem backend store writes images.</p> <p>Upon start up, Glance creates the directory if it doesn't already exist and verifies write access to the user under which <b>glance-api</b> runs. If the write access isn't available, a <b>BadStoreConfiguration</b> exception is raised and the filesystem store may not be available for adding new images.</p> <div data-bbox="815 616 922 1055" style="border: 1px solid black; padding: 5px; width: fit-content;">  </div> <p><b>NOTE</b></p> <p>This directory is used only when filesystem store is used as a storage backend. Either <b>filesystem_store_datadir</b> or <b>filesystem_store_datadirs</b> option must be specified in <b>glance-api.conf</b>. If both options are specified, a <b>BadStoreConfiguration</b> will be raised and the filesystem store may not be available for adding new images.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid path to a directory</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>filesystem_store_datadirs</b></li> <li>● <b>filesystem_store_file_perm</b></li> </ul>

Configuration option = Default value	Type	Description
<b>filesystem_store_datadirs</b> = None	multi valued	<p>List of directories and their priorities to which the filesystem backend store writes images.</p> <p>The filesystem store can be configured to store images in multiple directories as opposed to using a single directory specified by the <b>filesystem_store_datadir</b> configuration option. When using multiple directories, each directory can be given an optional priority to specify the preference order in which they should be used. Priority is an integer that is concatenated to the directory path with a colon where a higher value indicates higher priority. When two directories have the same priority, the directory with most free space is used. When no priority is specified, it defaults to zero.</p> <p>More information on configuring filesystem store with multiple store directories can be found at <a href="http://docs.openstack.org/developer/glance/configuring.html">http://docs.openstack.org/developer/glance/configuring.html</a></p> <div style="display: flex; align-items: flex-start;"> <div style="width: 40px; height: 100px; background: repeating-linear-gradient(45deg, transparent, transparent 2px, #ccc 2px, #ccc 4px); margin-right: 10px;"></div> <div> <p><b>NOTE</b></p> <p>This directory is used only when filesystem store is used as a storage backend. Either <b>filesystem_store_datadir</b> or <b>filesystem_store_datadirs</b> option must be specified in <b>glance-api.conf</b>. If both options are specified, a <b>BadStoreConfiguration</b> will be raised and the filesystem store may not be available for adding new images.</p> </div> </div> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● List of strings of the following form:             <ul style="list-style-type: none"> <li>○ <b>&lt;a valid directory path&gt;:&lt;optional integer priority&gt;</b></li> </ul> </li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>filesystem_store_datadir</b></li> <li>● <b>filesystem_store_file_perm</b></li> </ul>



Configuration option = filesystem_store_file_per m = 0 Default value	Type integer value	Description File access permissions for the image files.
		<p>Set the intended file access permissions for image data. This provides a way to enable other services, e.g. Nova, to consume images directly from the filesystem store. The users running the services that are intended to be given access to could be made a member of the group that owns the files created. Assigning a value less than or equal to zero for this configuration option signifies that no changes be made to the default permissions. This value will be decoded as an octal digit.</p> <p>For more information, please refer the documentation at <a href="http://docs.openstack.org/developer/glance/configuring.html">http://docs.openstack.org/developer/glance/configuring.html</a></p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid file access permission</li> <li>● Zero</li> <li>● Any negative integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<p><b>filesystem_store_metadata_file = None</b></p>	<p>string value</p>	<p>Filesystem store metadata file.</p> <p>The path to a file which contains the metadata to be returned with any location associated with the filesystem store. The file must contain a valid JSON object. The object should contain the keys <b>id</b> and <b>mountpoint</b>. The value for both keys should be a string.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid path to the store metadata file</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>

Configuration option = Default value	Type	Description
<b>http_proxy_information =</b> {}	dict value	<p>The http/https proxy information to be used to connect to the remote server.</p> <p>This configuration option specifies the http/https proxy information that should be used to connect to the remote server. The proxy information should be a key value pair of the scheme and proxy, for example, http:10.0.0.1:3128. You can also specify proxies for multiple schemes by separating the key value pairs with a comma, for example, http:10.0.0.1:3128, https:10.0.0.1:1080.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A comma separated list of scheme:proxy pairs as described above</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<b>https_ca_certificates_file = None</b>	string value	<p>Path to the CA bundle file.</p> <p>This configuration option enables the operator to use a custom Certificate Authority file to verify the remote server certificate. If this option is set, the <b>https_insecure</b> option will be ignored and the CA file specified will be used to authenticate the server certificate and establish a secure connection to the server.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A valid path to a CA file</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• https_insecure</li> </ul>


Configuration option = Default value	Type	Description
<b>https_insecure = True</b>	boolean value	<p>Set verification of the remote server certificate.</p> <p>This configuration option takes in a boolean value to determine whether or not to verify the remote server certificate. If set to True, the remote server certificate is not verified. If the option is set to False, then the default CA truststore is used for verification.</p> <p>This option is ignored if <b>https_ca_certificates_file</b> is set. The remote server certificate will then be verified using the file specified using the <b>https_ca_certificates_file</b> option.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● https_ca_certificates_file</li> </ul>
<b>rados_connect_timeout = 0</b>	integer value	<p>Timeout value for connecting to Ceph cluster.</p> <p>This configuration option takes in the timeout value in seconds used when connecting to the Ceph cluster i.e. it sets the time to wait for glance-api before closing the connection. This prevents glance-api hangups during the connection to RBD. If the value for this option is set to less than or equal to 0, no timeout is set and the default librados value is used.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>● Any integer value</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>

Configuration option = Default value	Type	Description
<b>rbd_store_ceph_conf = /etc/ceph/ceph.conf</b>	string value	<p>Ceph configuration file path.</p> <p>This configuration option takes in the path to the Ceph configuration file to be used. If the value for this option is not set by the user or is set to None, librados will locate the default configuration file which is located at /etc/ceph/ceph.conf. If using Cephx authentication, this file should include a reference to the right keyring in a client.&lt;USER&gt; section</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>• A valid path to a configuration file</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• rbd_store_user</li> </ul>
<b>rbd_store_chunk_size = 8</b>	integer value	<p>Size, in megabytes, to chunk RADOS images into.</p> <p>Provide an integer value representing the size in megabytes to chunk Glance images into. The default chunk size is 8 megabytes. For optimal performance, the value should be a power of two.</p> <p>When Ceph's RBD object storage system is used as the storage backend for storing Glance images, the images are chunked into objects of the size set using this option. These chunked objects are then stored across the distributed block data store to use for Glance.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>• Any positive integer value</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>

Configuration option = Default value	Type	Description
<b>rbd_store_pool = images</b>	string value	<p>RADOS pool in which images are stored.</p> <p>When RBD is used as the storage backend for storing Glance images, the images are stored by means of logical grouping of the objects (chunks of images) into a <b>pool</b>. Each pool is defined with the number of placement groups it can contain. The default pool that is used is <i>images</i>.</p> <p>More information on the RBD storage backend can be found here: <a href="http://ceph.com/planet/how-data-is-stored-in-ceph-cluster/">http://ceph.com/planet/how-data-is-stored-in-ceph-cluster/</a></p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>● A valid pool name</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>rbd_store_user = None</b>	string value	<p>RADOS user to authenticate as.</p> <p>This configuration option takes in the RADOS user to authenticate as. This is only needed when RADOS authentication is enabled and is applicable only if the user is using Cephx authentication. If the value for this option is not set by the user or is set to None, a default value will be chosen, which will be based on the client. section in <code>rbd_store_ceph_conf</code>.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>● A valid RADOS user</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>rbd_store_ceph_conf</code></li> </ul>

Configuration option = Default value	Type	Description
<b>rootwrap_config = /etc/glance/rootwrap.conf</b>	string value	<p>Path to the rootwrap configuration file to use for running commands as root.</p> <p>The cinder store requires root privileges to operate the image volumes (for connecting to iSCSI/FC volumes and reading/writing the volume data, etc.). The configuration file should allow the required commands by cinder store and os-brick library.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Path to the rootwrap config file</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>sheepdog_store_address = 127.0.0.1</b>	unknown value	<p>Address to bind the Sheepdog daemon to.</p> <p>Provide a string value representing the address to bind the Sheepdog daemon to. The default address set for the <i>sheep</i> is 127.0.0.1.</p> <p>The Sheepdog daemon, also called <i>sheep</i>, manages the storage in the distributed cluster by writing objects across the storage network. It identifies and acts on the messages directed to the address set using <b>sheepdog_store_address</b> option to store chunks of Glance images.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid IPv4 address</li> <li>● A valid IPv6 address</li> <li>● A valid hostname</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● sheepdog_store_port</li> </ul>

Configuration option = Default value	Type	Description
<b>sheepdog_store_chunk_size = 64</b>	integer value	<p>Chunk size for images to be stored in Sheepdog data store.</p> <p>Provide an integer value representing the size in mebibyte (1048576 bytes) to chunk Glance images into. The default chunk size is 64 mebibytes.</p> <p>When using Sheepdog distributed storage system, the images are chunked into objects of this size and then stored across the distributed data store to use for Glance.</p> <p>Chunk sizes, if a power of two, help avoid fragmentation and enable improved performance.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Positive integer value representing size in mebibytes.</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>
<b>sheepdog_store_port = 7000</b>	port value	<p>Port number on which the sheep daemon will listen.</p> <p>Provide an integer value representing a valid port number on which you want the Sheepdog daemon to listen on. The default port is 7000.</p> <p>The Sheepdog daemon, also called <i>sheep</i>, manages the storage in the distributed cluster by writing objects across the storage network. It identifies and acts on the messages it receives on the port number set using <b>sheepdog_store_port</b> option to store chunks of Glance images.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid port number (0 to 65535)</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● <code>sheepdog_store_address</code></li> </ul>


Configuration option = Default value	Type	Description
<b>store_capabilities_update_min_interval = 0</b>	integer value	<p>Minimum interval in seconds to execute updating dynamic storage capabilities based on current backend status.</p> <p>Provide an integer value representing time in seconds to set the minimum interval before an update of dynamic storage capabilities for a storage backend can be attempted. Setting <b>store_capabilities_update_min_interval</b> does not mean updates occur periodically based on the set interval. Rather, the update is performed at the elapse of this interval set, if an operation of the store is triggered.</p> <p>By default, this option is set to zero and is disabled. Provide an integer value greater than zero to enable this option.</p> <div data-bbox="815 898 922 1155" style="border: 1px solid black; padding: 5px; width: fit-content;">  </div> <p><b>NOTE</b></p> <p>For more information on store capabilities and their updates, please visit:  <a href="https://specs.openstack.org/openstack/glance-specs/specs/kilo/store-capabilities.html">https://specs.openstack.org/openstack/glance-specs/specs/kilo/store-capabilities.html</a></p> <p>For more information on setting up a particular store in your deployment and help with the usage of this feature, please contact the storage driver maintainers listed here:  <a href="http://docs.openstack.org/developer/glance_store/drivers/index.html">http://docs.openstack.org/developer/glance_store/drivers/index.html</a></p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Zero</li> <li>● Positive integer</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>




Configuration option = Default value	Type	Description
<b>stores = [u'file', u'http']</b>	list value	<p>List of enabled Glance stores.</p> <p>Register the storage backends to use for storing disk images as a comma separated list. The default stores enabled for storing disk images with Glance are <b>file</b> and <b>http</b>.</p> <p>Possible values:</p> <ul style="list-style-type: none"><li>● A comma separated list that could include:<ul style="list-style-type: none"><li>○ file</li><li>○ http</li><li>○ swift</li><li>○ rbd</li><li>○ sheepdog</li><li>○ cinder</li><li>○ vmware</li></ul></li></ul> <p>Related Options:</p> <ul style="list-style-type: none"><li>● default_store</li></ul>

Configuration option = Default value	Type	Description
<b>swift_buffer_on_upload = False</b>	boolean value	<p>Buffer image segments before upload to Swift.</p> <p>Provide a boolean value to indicate whether or not Glance should buffer image data to disk while uploading to swift. This enables Glance to resume uploads on error.</p> <p>NOTES: When enabling this option, one should take great care as this increases disk usage on the API node. Be aware that depending upon how the file system is configured, the disk space used for buffering may decrease the actual disk space available for the glance image cache. Disk utilization will cap according to the following equation: <b>(swift_store_large_object_chunk_size * workers * 1000)</b></p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• True</li> <li>• False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• swift_upload_buffer_dir</li> </ul>
<b>swift_store_admin_tenant s = []</b>	list value	<p>List of tenants that will be granted admin access.</p> <p>This is a list of tenants that will be granted read/write access on all Swift containers created by Glance in multi-tenant mode. The default value is an empty list.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A comma separated list of strings representing UUIDs of Keystone projects/tenants</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<b>swift_store_auth_address = None</b>	string value	<p>The address where the Swift authentication service is listening.</p>

Configuration option = Default value	Type	Description
<b>swift_store_auth_insecure = False</b>	boolean value	<p>Set verification of the server certificate.</p> <p>This boolean determines whether or not to verify the server certificate. If this option is set to True, swiftclient won't check for a valid SSL certificate when authenticating. If the option is set to False, then the default CA truststore is used for verification.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• True</li> <li>• False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• swift_store_cacert</li> </ul>
<b>swift_store_auth_version = 2</b>	string value	<p>Version of the authentication service to use. Valid versions are 2 and 3 for keystone and 1 (deprecated) for swauth and rackspace.</p>
<b>swift_store_cacert = None</b>	string value	<p>Path to the CA bundle file.</p> <p>This configuration option enables the operator to specify the path to a custom Certificate Authority file for SSL verification when connecting to Swift.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A valid path to a CA file</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• swift_store_auth_insecure</li> </ul>

Configuration option = Default value	Type	Description
<b>swift_store_config_file =</b> <b>None</b>	string value	<p>Absolute path to the file containing the swift account(s) configurations.</p> <p>Include a string value representing the path to a configuration file that has references for each of the configured Swift account(s)/backing stores. By default, no file path is specified and customized Swift referencing is disabled. Configuring this option is highly recommended while using Swift storage backend for image storage as it avoids storage of credentials in the database.</p> <div data-bbox="815 689 922 884" style="display: inline-block; vertical-align: top;">  </div> <p><b>NOTE</b></p> <p>Please do not configure this option if you have set <b>swift_store_multi_tenant</b> to <b>True</b>.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● String value representing an absolute path on the glance-api node</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● swift_store_multi_tenant</li> </ul>


Configuration option = Default value	Type	Description
<b>swift_store_container = glance</b>	string value	<p>Name of single container to store images/name prefix for multiple containers</p> <p>When a single container is being used to store images, this configuration option indicates the container within the Glance account to be used for storing all images. When multiple containers are used to store images, this will be the name prefix for all containers. Usage of single/multiple containers can be controlled using the configuration option <b>swift_store_multiple_containers_seed</b>.</p> <p>When using multiple containers, the containers will be named after the value set for this configuration option with the first N chars of the image UUID as the suffix delimited by an underscore (where N is specified by <b>swift_store_multiple_containers_seed</b>).</p> <p>Example: if the seed is set to 3 and <code>swift_store_container = glance</code>, then an image with UUID <b>fdae39a1-bac5-4238-aba4-69bcc726e848</b> would be placed in the container <b>glance_fda</b>. All dashes in the UUID are included when creating the container name but do not count toward the character limit, so when N=10 the container name would be <b>glance_fdae39a1-ba</b>.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● If using single container, this configuration option can be any string that is a valid swift container name in Glance's Swift account</li> <li>● If using multiple containers, this configuration option can be any string as long as it satisfies the container naming rules enforced by Swift. The value of <b>swift_store_multiple_containers_seed</b> should be taken into account as well.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>swift_store_multiple_containers_seed</b></li> <li>● <b>swift_store_multi_tenant</b></li> <li>● <b>swift_store_create_container_on_put</b></li> </ul>


Configuration option = Default value	Type	Description
<b>swift_store_create_container_on_put = False</b>	boolean value	<p>Create container, if it doesn't already exist, when uploading image.</p> <p>At the time of uploading an image, if the corresponding container doesn't exist, it will be created provided this configuration option is set to True. By default, it won't be created. This behavior is applicable for both single and multiple containers mode.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• True</li> <li>• False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<b>swift_store_endpoint = None</b>	string value	<p>The URL endpoint to use for Swift backend storage.</p> <p>Provide a string value representing the URL endpoint to use for storing Glance images in Swift store. By default, an endpoint is not set and the storage URL returned by <b>auth</b> is used. Setting an endpoint with <b>swift_store_endpoint</b> overrides the storage URL and is used for Glance image storage.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p>The URL should include the path up to, but excluding the container. The location of an object is obtained by appending the container and object to the configured URL.</p> </div> </div> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• String value representing a valid URL path up to a Swift container</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>


Configuration option = Default value	Type	Description
<b>swift_store_endpoint_type = publicURL</b>	string value	<p>Endpoint Type of Swift service.</p> <p>This string value indicates the endpoint type to use to fetch the Swift endpoint. The endpoint type determines the actions the user will be allowed to perform, for instance, reading and writing to the Store. This setting is only used if <code>swift_store_auth_version</code> is greater than 1.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● publicURL</li> <li>● adminURL</li> <li>● internalURL</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● swift_store_endpoint</li> </ul>
<b>swift_store_expire_soon_interval = 60</b>	integer value	<p>Time in seconds defining the size of the window in which a new token may be requested before the current token is due to expire.</p> <p>Typically, the Swift storage driver fetches a new token upon the expiration of the current token to ensure continued access to Swift. However, some Swift transactions (like uploading image segments) may not recover well if the token expires on the fly.</p> <p>Hence, by fetching a new token before the current token expiration, we make sure that the token does not expire or is close to expiry before a transaction is attempted. By default, the Swift storage driver requests for a new token 60 seconds or less before the current token expiration.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Zero</li> <li>● Positive integer value</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>

Configuration option = Default value	Type	Description
<b>swift_store_key = None</b>	string value	Auth key for the user authenticating against the Swift authentication service.
<b>swift_store_large_object_chunk_size = 200</b>	integer value	<p>The maximum size, in MB, of the segments when image data is segmented.</p> <p>When image data is segmented to upload images that are larger than the limit enforced by the Swift cluster, image data is broken into segments that are no bigger than the size specified by this configuration option. Refer to <b>swift_store_large_object_size</b> for more detail.</p> <p>For example: if <b>swift_store_large_object_size</b> is 5GB and <b>swift_store_large_object_chunk_size</b> is 1GB, an image of size 6.2GB will be segmented into 7 segments where the first six segments will be 1GB in size and the seventh segment will be 0.2GB.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A positive integer that is less than or equal to the large object limit enforced by Swift cluster in consideration.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <b>swift_store_large_object_size</b></li> </ul>



Configuration option = Default value	Type	Description
<b>swift_store_large_object_size = 5120</b>	integer value	<p>The size threshold, in MB, after which Glance will start segmenting image data.</p> <p>Swift has an upper limit on the size of a single uploaded object. By default, this is 5GB. To upload objects bigger than this limit, objects are segmented into multiple smaller objects that are tied together with a manifest file. For more detail, refer to <a href="http://docs.openstack.org/developer/swift/overview_large_objects.html">http://docs.openstack.org/developer/swift/overview_large_objects.html</a></p> <p>This configuration option specifies the size threshold over which the Swift driver will start segmenting image data into multiple smaller files. Currently, the Swift driver only supports creating Dynamic Large Objects.</p> <div data-bbox="815 862 922 1055" style="display: inline-block; vertical-align: top;">  </div> <p><b>NOTE</b></p> <p>This should be set by taking into account the large object limit enforced by the Swift cluster in consideration.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A positive integer that is less than or equal to the large object limit enforced by the Swift cluster in consideration.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>swift_store_large_object_chunk_size</b></li> </ul>


Configuration option = Default value	Type	Description
<b>swift_store_multi_tenant</b> = <b>False</b>	boolean value	<p>Store images in tenant's Swift account.</p> <p>This enables multi-tenant storage mode which causes Glance images to be stored in tenant specific Swift accounts. If this is disabled, Glance stores all images in its own account. More details multi-tenant store can be found at <a href="https://wiki.openstack.org/wiki/GlanceSwiftTenantSpecificStorage">https://wiki.openstack.org/wiki/GlanceSwiftTenantSpecificStorage</a></p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p>If using multi-tenant swift store, please make sure that you do not set a swift configuration file with the <i>swift_store_config_file</i> option.</p> </div> </div> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>swift_store_config_file</code></li> </ul>

Configuration option = Default value	Type	Description
<b>swift_store_multiple_containers_seed = 0</b>	integer value	<p>Seed indicating the number of containers to use for storing images.</p> <p>When using a single-tenant store, images can be stored in one or more than one containers. When set to 0, all images will be stored in one single container. When set to an integer value between 1 and 32, multiple containers will be used to store images. This configuration option will determine how many containers are created. The total number of containers that will be used is equal to <math>16^N</math>, so if this config option is set to 2, then <math>16^2=256</math> containers will be used to store images.</p> <p>Please refer to <b>swift_store_container</b> for more detail on the naming convention. More detail about using multiple containers can be found at <a href="https://specs.openstack.org/openstack/glance-specs/specs/kilo/swift-store-multiple-containers.html">https://specs.openstack.org/openstack/glance-specs/specs/kilo/swift-store-multiple-containers.html</a></p> <div data-bbox="815 1003 922 1137" style="display: inline-block; vertical-align: top;">  </div> <p><b>NOTE</b></p> <p>This is used only when <code>swift_store_multi_tenant</code> is disabled.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A non-negative integer less than or equal to 32</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>swift_store_container</b></li> <li>● <b>swift_store_multi_tenant</b></li> <li>● <b>swift_store_create_container_on_put</b></li> </ul>

Configuration option = Default value	Type	Description
<b>swift_store_region = None</b>	string value	<p>The region of Swift endpoint to use by Glance.</p> <p>Provide a string value representing a Swift region where Glance can connect to for image storage. By default, there is no region set.</p> <p>When Glance uses Swift as the storage backend to store images for a specific tenant that has multiple endpoints, setting of a Swift region with <b>swift_store_region</b> allows Glance to connect to Swift in the specified region as opposed to a single region connectivity.</p> <p>This option can be configured for both single-tenant and multi-tenant storage.</p> <div data-bbox="815 819 922 1048" style="float: left; margin-right: 10px;"> </div> <p><b>NOTE</b></p> <p>Setting the region with <b>swift_store_region</b> is tenant-specific and is necessary <b>only if</b> the tenant has multiple endpoints across different regions.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A string value representing a valid Swift region.</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>

Configuration option = Default value	Type	Description
<b>swift_store_retry_get_count = 0</b>	integer value	<p>The number of times a Swift download will be retried before the request fails.</p> <p>Provide an integer value representing the number of times an image download must be retried before erroring out. The default value is zero (no retry on a failed image download). When set to a positive integer value, <b>swift_store_retry_get_count</b> ensures that the download is attempted this many more times upon a download failure before sending an error message.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Zero</li> <li>• Positive integer value</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>
<b>swift_store_service_type = object-store</b>	string value	<p>Type of Swift service to use.</p> <p>Provide a string value representing the service type to use for storing images while using Swift backend storage. The default service type is set to <b>object-store</b>.</p> <div data-bbox="815 1294 922 1547" style="border: 1px solid black; padding: 5px; width: fit-content;">  </div> <p><b>NOTE</b></p> <p>If <b>swift_store_auth_version</b> is set to 2, the value for this configuration option needs to be <b>object-store</b>. If using a higher version of Keystone or a different auth scheme, this option may be modified.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A string representing a valid service type for Swift storage.</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li>• None</li> </ul>

Configuration option = Default value	Type	Description
<b>swift_store_ssl_compression = True</b>	boolean value	<p>SSL layer compression for HTTPS Swift requests.</p> <p>Provide a boolean value to determine whether or not to compress HTTPS Swift requests for images at the SSL layer. By default, compression is enabled.</p> <p>When using Swift as the backend store for Glance image storage, SSL layer compression of HTTPS Swift requests can be set using this option. If set to False, SSL layer compression of HTTPS Swift requests is disabled. Disabling this option may improve performance for images which are already in a compressed format, for example, qcow2.</p> <p>Possible values:</p> <ul style="list-style-type: none"><li>• True</li><li>• False</li></ul> <p>Related Options:</p> <ul style="list-style-type: none"><li>• None</li></ul>

Configuration option = Default value	Type	Description
<b>swift_store_use_trusts = True</b>	boolean value	<p>Use trusts for multi-tenant Swift store.</p> <p>This option instructs the Swift store to create a trust for each add/get request when the multi-tenant store is in use. Using trusts allows the Swift store to avoid problems that can be caused by an authentication token expiring during the upload or download of data.</p> <p>By default, <b>swift_store_use_trusts</b> is set to <b>True</b>(use of trusts is enabled). If set to <b>False</b>, a user token is used for the Swift connection instead, eliminating the overhead of trust creation.</p> <div data-bbox="815 752 922 920" style="display: inline-block; vertical-align: top;">  </div> <p><b>NOTE</b></p> <p>This option is considered only when <b>swift_store_multi_tenant</b> is set to <b>True</b></p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● swift_store_multi_tenant</li> </ul>
<b>swift_store_user = None</b>	string value	<p>The user to authenticate against the Swift authentication service.</p>

Configuration option = Default value	Type	Description
<b>swift_upload_buffer_dir = None</b>	string value	<p>Directory to buffer image segments before upload to Swift.</p> <p>Provide a string value representing the absolute path to the directory on the glance node where image segments will be buffered briefly before they are uploaded to swift.</p> <p>NOTES: * This is required only when the configuration option <b>swift_buffer_on_upload</b> is set to True. * This directory should be provisioned keeping in mind the <b>swift_store_large_object_chunk_size</b> and the maximum number of images that could be uploaded simultaneously by a given glance node.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● String value representing an absolute directory path</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>swift_buffer_on_upload</code></li> <li>● <code>swift_store_large_object_chunk_size</code></li> </ul>
<b>vmware_api_retry_count = 10</b>	integer value	<p>The number of VMware API retries.</p> <p>This configuration option specifies the number of times the VMware ESX/VC server API must be retried upon connection related issues or server API call overload. It is not possible to specify <i>retry forever</i>.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>● Any positive integer value</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● None</li> </ul>



Configuration option = Default value	Type	Description
<b>vmware_ca_file = None</b>	string value	<p>Absolute path to the CA bundle file.</p> <p>This configuration option enables the operator to use a custom Certificate Authority File to verify the ESX/vCenter certificate.</p> <p>If this option is set, the "vmware_insecure" option will be ignored and the CA file specified will be used to authenticate the ESX/vCenter server certificate and establish a secure connection to the server.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>Any string that is a valid absolute path to a CA file</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>vmware_insecure</li> </ul>
<b>vmware_datastores = None</b>	multi valued	<p>The datastores where the image can be stored.</p> <p>This configuration option specifies the datastores where the image can be stored in the VMWare store backend. This option may be specified multiple times for specifying multiple datastores. The datastore name should be specified after its datacenter path, separated by ":". An optional weight may be given after the datastore name, separated again by ":" to specify the priority. Thus, the required format becomes &lt;datacenter_path&gt;:&lt;datastore_name&gt;:&lt;optional_weight&gt;.</p> <p>When adding an image, the datastore with highest weight will be selected, unless there is not enough free space available in cases where the image size is already known. If no weight is given, it is assumed to be zero and the directory will be considered for selection last. If multiple datastores have the same weight, then the one with the most free space available is selected.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>Any string of the format: &lt;datacenter_path&gt;:&lt;datastore_name&gt;:&lt;optional_weight&gt;</li> </ul> <p>Related options: * None</p>

Configuration option = Default value	Type	Description
<b>vmware_insecure = False</b>	boolean value	<p>Set verification of the ESX/vCenter server certificate.</p> <p>This configuration option takes a boolean value to determine whether or not to verify the ESX/vCenter server certificate. If this option is set to True, the ESX/vCenter server certificate is not verified. If this option is set to False, then the default CA truststore is used for verification.</p> <p>This option is ignored if the "vmware_ca_file" option is set. In that case, the ESX/vCenter server certificate will then be verified using the file specified using the "vmware_ca_file" option .</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>● True</li> <li>● False</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● vmware_ca_file</li> </ul>
<b>vmware_server_host = None</b>	unknown value	<p>Address of the ESX/ESXi or vCenter Server target system.</p> <p>This configuration option sets the address of the ESX/ESXi or vCenter Server target system. This option is required when using the VMware storage backend. The address can contain an IP address (127.0.0.1) or a DNS name (www.my-domain.com).</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>● A valid IPv4 or IPv6 address</li> <li>● A valid DNS name</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● vmware_server_username</li> <li>● vmware_server_password</li> </ul>

Configuration option = Default value	Type	Description
<b>vmware_server_password = None</b>	string value	<p>Server password.</p> <p>This configuration option takes the password for authenticating with the VMware ESX/ESXi or vCenter Server. This option is required when using the VMware storage backend.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>Any string that is a password corresponding to the username specified using the "vmware_server_username" option</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>vmware_server_host</li> <li>vmware_server_username</li> </ul>
<b>vmware_server_username = None</b>	string value	<p>Server username.</p> <p>This configuration option takes the username for authenticating with the VMware ESX/ESXi or vCenter Server. This option is required when using the VMware storage backend.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>Any string that is the username for a user with appropriate privileges</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>vmware_server_host</li> <li>vmware_server_password</li> </ul>

Configuration option = Default value	Type	Description
<b>vmware_store_image_dir</b> <b>= /openstack_glance</b>	string value	<p>The directory where the glance images will be stored in the datastore.</p> <p>This configuration option specifies the path to the directory where the glance images will be stored in the VMware datastore. If this option is not set, the default directory where the glance images are stored is <code>openstack_glance</code>.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>Any string that is a valid path to a directory</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>None</li> </ul>
<b>vmware_task_poll_interval</b> <b>= 5</b>	integer value	<p>Interval in seconds used for polling remote tasks invoked on VMware ESX/VC server.</p> <p>This configuration option takes in the sleep time in seconds for polling an on-going async task as part of the VMWare ESX/VC server API call.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>Any positive integer value</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>None</li> </ul>

### 5.4.3. oslo\_policy

The following table outlines the options available under the **[oslo\_policy]** group in the `/etc/glance/glance-cache.conf` file.

Table 5.40. oslo\_policy

Configuration option = Default value	Type	Description
<b>enforce_scope = False</b>	boolean value	<p>This option controls whether or not to enforce scope when evaluating policies. If <b>True</b>, the scope of the token used in the request is compared to the <b>scope_types</b> of the policy being enforced. If the scopes do not match, an <b>InvalidScope</b> exception will be raised. If <b>False</b>, a message will be logged informing operators that policies are being invoked with mismatching scope.</p>

Configuration option = Default value	Type	Description
<b>policy_default_rule = default</b>	string value	Default rule. Enforced when a requested rule is not found.
<b>policy_dirs = [u'policy.d']</b>	multi valued	Directories where policy configuration files are stored. They can be relative to any directory in the search path defined by the config_dir option, or absolute paths. The file defined by policy_file must exist for these directories to be searched. Missing or empty directories are ignored.
<b>policy_file = policy.json</b>	string value	The file that defines policies.
<b>remote_content_type = application/x-www-form-urlencoded</b>	string value	Content Type to send and receive data for REST based policy check
<b>remote_ssl_ca_cert_file = None</b>	string value	Absolute path to ca cert file for REST based policy check
<b>remote_ssl_client_cert_file = None</b>	string value	Absolute path to client cert for REST based policy check
<b>remote_ssl_client_key_file = None</b>	string value	Absolute path client key file REST based policy check
<b>remote_ssl_verify_server_cert = False</b>	boolean value	server identity verification for REST based policy check

## CHAPTER 6. HEAT

The following chapter contains information about the configuration options in the **heat** service.

### 6.1. HEAT.CONF

This section contains options for the `/etc/heat/heat.conf` file.

#### 6.1.1. DEFAULT

The following table outlines the options available under the **[DEFAULT]** group in the `/etc/heat/heat.conf` file.

Table 6.1. DEFAULT

Configuration option = Default value	Type	Description
<b>action_retry_limit = 5</b>	integer value	Number of times to retry to bring a resource to a non-error state. Set to 0 to disable retries.
<b>auth_encryption_key = notgood but just long enough i t</b>	string value	Key used to encrypt authentication info in the database. Length of this key must be 32 characters.
<b>backdoor_port = None</b>	string value	Enable eventlet backdoor. Acceptable values are 0, <port>, and <start>:<end>, where 0 results in listening on a random tcp port number; <port> results in listening on the specified port number (and not enabling backdoor if that port is in use); and <start>:<end> results in listening on the smallest unused port number within the specified range of port numbers. The chosen port is displayed in the service's log file.
<b>backdoor_socket = None</b>	string value	Enable eventlet backdoor, using the provided path as a unix socket that can receive connections. This option is mutually exclusive with <code>backdoor_port</code> in that only one should be provided. If both are provided then the existence of this option overrides the usage of that option.
<b>client_retry_limit = 2</b>	integer value	Number of times to retry when a client encounters an expected intermittent error. Set to 0 to disable retries.
<b>cloud_backend = heat.engine.clients.OpenStackClients</b>	string value	Fully qualified class name to use as a client backend.
<b>conn_pool_min_size = 2</b>	integer value	The pool size limit for connections expiration policy

Configuration option = Default value	Type	Description
<b>conn_pool_ttl = 1200</b>	integer value	The time-to-live in sec of idle connections in the pool
<b>control_exchange = openstack</b>	string value	The default exchange under which topics are scoped. May be overridden by an exchange name specified in the transport_url option.
<b>convergence_engine = True</b>	boolean value	Enables engine with convergence architecture. All stacks with this option will be created using convergence engine.
<b>debug = False</b>	boolean value	If set to true, the logging level will be set to DEBUG instead of the default INFO level.
<b>default_deployment_signal_transport = CFN_SIGNAL</b>	string value	Template default for how the server should signal to heat with the deployment output values. CFN_SIGNAL will allow an HTTP POST to a CFN keypair signed URL (requires enabled heat-api-cfn). TEMP_URL_SIGNAL will create a Swift TempURL to be signaled via HTTP PUT (requires object-store endpoint which supports TempURL). HEAT_SIGNAL will allow calls to the Heat API resource-signal using the provided keystone credentials. ZAQAR_SIGNAL will create a dedicated zaqar queue to be signaled using the provided keystone credentials.

Configuration option = Default value	Type	Description
<b>default_log_levels =</b> [u'amqp=WARN', u'amqplib=WARN', u'boto=WARN', u'qpid=WARN', u'sqlalchemy=WARN', u'suds=INFO', u'oslo.messaging=INFO', u'oslo_messaging=INFO', u'iso8601=WARN', u'requests.packages.urlli b3.connectionpool=WAR N', u'urllib3.connectionpool= WARN', u'websocket=WARN', u'requests.packages.urlli b3.util.retry=WARN', u'urllib3.util.retry=WARN', u'keystonemiddleware=W ARN', u'routes.middleware=WA RN', u'stevedore=WARN', u'taskflow=WARN', u'keystoneauth=WARN', u'oslo.cache=INFO', u'dogpile.core.dogpile=IN FO']	list value	List of package logging levels in logger=LEVEL pairs. This option is ignored if log_config_append is set.
<b>default_notification_level = INFO</b>	string value	Default notification level for outgoing notifications.
<b>default_publisher_id = None</b>	string value	Default publisher_id for outgoing notifications.
<b>default_software_config_t ransport = POLL_SERVER_CFN</b>	string value	Template default for how the server should receive the metadata required for software configuration. POLL_SERVER_CFN will allow calls to the cfn API action DescribeStackResource authenticated with the provided keypair (requires enabled heat-api-cfn). POLL_SERVER_HEAT will allow calls to the Heat API resource-show using the provided keystone credentials (requires keystone v3 API, and configured stack_user_* config options). POLL_TEMP_URL will create and populate a Swift TempURL with metadata for polling (requires object-store endpoint which supports TempURL).ZAQAR_MESSAGE will create a dedicated zaqar queue and post the metadata for polling.



Configuration option = Default value	Type	Description
<b>default_user_data_format = HEAT_CFNTOOLS</b>	string value	Template default for how the user_data should be formatted for the server. For HEAT_CFNTOOLS, the user_data is bundled as part of the heat-cfn-tools cloud-init boot configuration data. For RAW the user_data is passed to Nova unmodified. For SOFTWARE_CONFIG user_data is bundled as part of the software config data, and metadata is derived from any associated SoftwareDeployment resources.
<b>deferred_auth_method = trusts</b>	string value	Select deferred auth method, stored password or trusts.
<b>enable_cloud_watch_lite = False</b>	boolean value	Enable the legacy OS::Heat::CWLiteAlarm resource.
<b>enable_stack_abandon = False</b>	boolean value	Enable the preview Stack Abandon feature.
<b>enable_stack_adopt = False</b>	boolean value	Enable the preview Stack Adopt feature.
<b>encrypt_parameters_and_properties = False</b>	boolean value	Encrypt template parameters that were marked as hidden and also all the resource properties before storing them in database.
<b>engine_life_check_timeout = 2</b>	integer value	RPC timeout for the engine liveness check that is used for stack locking.
<b>environment_dir = /etc/heat/environment.d</b>	string value	The directory to search for environment files.
<b>error_wait_time = 240</b>	integer value	The amount of time in seconds after an error has occurred that tasks may continue to run before being cancelled.
<b>event_purge_batch_size = 200</b>	integer value	Controls how many events will be pruned whenever a stack's events are purged. Set this lower to keep more events at the expense of more frequent purges.
<b>executor_thread_pool_size = 64</b>	integer value	Size of executor thread pool when executor is threading or eventlet.
<b>fatal_deprecations = False</b>	boolean value	Enables or disables fatal status of deprecations.

Configuration option = Default value	Type	Description
<b>graceful_shutdown_timeout = 60</b>	integer value	Specify a timeout after which a gracefully shutdown server will exit. Zero value means endless wait.
<b>heat_metadata_server_url = None</b>	string value	URL of the Heat metadata server. NOTE: Setting this is only needed if you require instances to use a different endpoint than in the keystone catalog
<b>heat_stack_user_role = heat_stack_user</b>	string value	Keystone role for heat template-defined users.
<b>heat_waitcondition_server_url = None</b>	string value	URL of the Heat waitcondition server.
<code>heat_watch_server_url = `</code>	string value	URL of the Heat CloudWatch server.
<b>hidden_stack_tags = [u'data-processing-cluster']</b>	list value	Stacks containing these tag names will be hidden. Multiple tags should be given in a comma-delimited list (eg. hidden_stack_tags=hide_me,me_too).
<b>host = ccosp-undercloud.usersys.redhat.com</b>	string value	Name of the engine node. This can be an opaque identifier. It is not necessarily a hostname, FQDN, or IP address.
<b>instance_connection_https_validate_certificates = 1</b>	string value	Instance connection to CFN/CW API validate certs if SSL is used.
<b>instance_connection_is_secure = 0</b>	string value	Instance connection to CFN/CW API via https.
<code>instance_format = [instance: %(uuid)s] `</code>	string value	The format for an instance that is passed with the log message.
<code>instance_uuid_format = [instance: %(uuid)s] `</code>	string value	The format for an instance UUID that is passed with the log message.
<b>keystone_backend = heat.engine.clients.os.keystone.heat_keystoneclient.KsClientWrapper</b>	string value	Fully qualified class name to use as a keystone backend.
<b>loadbalancer_template = None</b>	string value	Custom template for the built-in loadbalancer nested stack.

Configuration option = Default value	Type	Description
<b>log-config-append = None</b>	string value	The name of a logging configuration file. This file is appended to any existing logging configuration files. For details about logging configuration files, see the Python logging module documentation. Note that when logging configuration files are used then all logging configuration is set in the configuration file and other logging configuration options are ignored (for example, <code>logging_context_format_string</code> ).
<b>log-date-format = %Y-%m-%d %H:%M:%S</b>	string value	Defines the format string for <code>%(asctime)s</code> in log records. Default: <code>%(default)s</code> . This option is ignored if <code>log_config_append</code> is set.
<b>log-dir = None</b>	string value	(Optional) The base directory used for relative <code>log_file</code> paths. This option is ignored if <code>log_config_append</code> is set.
<b>log-file = None</b>	string value	(Optional) Name of log file to send logging output to. If no default is set, logging will go to <code>stderr</code> as defined by <code>use_stderr</code> . This option is ignored if <code>log_config_append</code> is set.
<b>log_options = True</b>	boolean value	Enables or disables logging values of all registered options when starting a service (at <code>DEBUG</code> level).
<b>logging_context_format_string = %(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [%(request_id)s %(user_identity)s] %(instance)s%(message)s</b>	string value	Format string to use for log messages with context.
<b>logging_debug_format_s uffix = %(funcName)s % (pathname)s:%(lineno)d</b>	string value	Additional data to append to log message when logging level for the message is <code>DEBUG</code> .
<b>logging_default_format_s tring = %(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [-] %(instance)s%(message)s</b>	string value	Format string to use for log messages when context is undefined.
<b>logging_exception_prefix = %(asctime)s.%(msecs)03d %(process)d ERROR %(name)s % (instance)s</b>	string value	Prefix each line of exception output with this format.

Configuration option = Default value	Type	Description
<b>logging_user_identity_for_mat = %(user)s %(tenant)s %(domain)s %(user_domain)s %(project_domain)s</b>	string value	Defines the format string for %(user_identity)s that is used in logging_context_format_string.
<b>max_events_per_stack = 1000</b>	integer value	Rough number of maximum events that will be available per stack. Actual number of events can be a bit higher since purge checks take place randomly 200/event_purge_batch_size percent of the time. Older events are deleted when events are purged. Set to 0 for unlimited events per stack.
<b>max_interface_check_attempts = 10</b>	integer value	Number of times to check whether an interface has been attached or detached.
<b>max_json_body_size = 1048576</b>	integer value	Maximum raw byte size of JSON request body. Should be larger than max_template_size.
<b>max_nested_stack_depth = 5</b>	integer value	Maximum depth allowed when using nested stacks.
<b>max_resources_per_stack = 1000</b>	integer value	Maximum resources allowed per top-level stack. -1 stands for unlimited.
<b>max_server_name_length = 53</b>	integer value	Maximum length of a server name to be used in nova.
<b>max_stacks_per_tenant = 100</b>	integer value	Maximum number of stacks any one tenant may have active at one time.
<b>max_template_size = 524288</b>	integer value	Maximum raw byte size of any template.
<b>num_engine_workers = None</b>	integer value	Number of heat-engine processes to fork and run. Will default to either to 4 or number of CPUs on the host, whichever is greater.
<b>observe_on_update = False</b>	boolean value	On update, enables heat to collect existing resource properties from reality and converge to updated template.
<b>onready = None</b>	string value	Deprecated.
<b>periodic_interval = 60</b>	integer value	Seconds between running periodic tasks.

Configuration option = Default value	Type	Description
<b>plugin_dirs =</b> [u'/usr/lib64/heat', u'/usr/lib/heat', u'/usr/local/lib/heat', u'/usr/local/lib64/heat']	list value	List of directories to search for plug-ins.
<b>publish_errors = False</b>	boolean value	Enables or disables publication of error events.
<b>rate_limit_burst = 0</b>	integer value	Maximum number of logged messages per rate_limit_interval.
<b>rate_limit_except_level = CRITICAL</b>	string value	Log level name used by rate limiting: CRITICAL, ERROR, INFO, WARNING, DEBUG or empty string. Logs with level greater or equal to rate_limit_except_level are not filtered. An empty string means that all levels are filtered.
<b>rate_limit_interval = 0</b>	integer value	Interval, number of seconds, of log rate limiting.
<code>reauthentication_auth_method = `</code>	string value	Allow reauthentication on token expiry, such that long-running tasks may complete. Note this defeats the expiry of any provided user tokens.
<b>region_name_for_services = None</b>	string value	Default region name used to get services endpoints.
<b>rpc_ack_timeout_base = 15</b>	integer value	Number of seconds to wait for an ack from a cast/call. After each retry attempt this timeout is multiplied by some specified multiplier.
<b>rpc_ack_timeout_multiplier = 2</b>	integer value	Number to multiply base ack timeout by after each retry attempt.
<b>rpc_backend = rabbit</b>	string value	The messaging driver to use, defaults to rabbit. Other drivers include amqp and zmq.
<b>rpc_conn_pool_size = 30</b>	integer value	Size of RPC connection pool.
<b>rpc_message_ttl = 300</b>	integer value	Expiration timeout in seconds of a sent/received message after which it is not tracked anymore by a client/server.
<b>rpc_poll_timeout = 1</b>	integer value	The default number of seconds that poll should wait. Poll raises timeout exception when timeout expired.
<b>rpc_response_timeout = 60</b>	integer value	Seconds to wait for a response from a call.

Configuration option = Default value	Type	Description
<b>rpc_retry_attempts = 3</b>	integer value	Default number of message sending attempts in case of any problems occurred: positive value N means at most N retries, 0 means no retries, None or -1 (or any other negative values) mean to retry forever. This option is used only if acknowledgments are enabled.
<b>rpc_thread_pool_size = 100</b>	integer value	Maximum number of (green) threads to work concurrently.
<b>rpc_use_acks = False</b>	boolean value	Wait for message acknowledgements from receivers. This mechanism works only via proxy without PUB/SUB.
<b>rpc_zmq_bind_address = *</b>	string value	ZeroMQ bind address. Should be a wildcard (*), an ethernet interface, or IP. The "host" option should point or resolve to this address.
<b>rpc_zmq_bind_port_retries = 100</b>	integer value	Number of retries to find free port number before fail with ZMQBindError.
<b>rpc_zmq_contexts = 1</b>	integer value	Number of ZeroMQ contexts, defaults to 1.
<b>rpc_zmq_host = ccsosp-undercloud.usersys.redhat.com</b>	string value	Name of this node. Must be a valid hostname, FQDN, or IP address. Must match "host" option, if running Nova.
<b>rpc_zmq_ipc_dir = /var/run/openstack</b>	string value	Directory for holding IPC sockets.
<b>rpc_zmq_matchmaker = redis</b>	string value	MatchMaker driver.
<b>rpc_zmq_max_port = 65536</b>	integer value	Maximal port number for random ports range.
<b>rpc_zmq_min_port = 49153</b>	port value	Minimal port number for random ports range.
<b>rpc_zmq_serialization = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>rpc_zmq_topic_backlog = None</b>	integer value	Maximum number of ingress messages to locally buffer per topic. Default is unlimited.
<b>run_external_periodic_tasks = True</b>	boolean value	Some periodic tasks can be run in a separate process. Should we run them here?

Configuration option = Default value	Type	Description
<b>stack_action_timeout = 3600</b>	integer value	Timeout in seconds for stack action (ie. create or update).
<b>stack_domain_admin = None</b>	string value	Keystone username, a user with roles sufficient to manage users and projects in the stack_user_domain.
<b>stack_domain_admin_password = None</b>	string value	Keystone password for stack_domain_admin user.
<b>stack_scheduler_hints = False</b>	boolean value	When this feature is enabled, scheduler hints identifying the heat stack context of a server or volume resource are passed to the configured schedulers in nova and cinder, for creates done using heat resource types OS::Cinder::Volume, OS::Nova::Server, and AWS::EC2::Instance. heat_root_stack_id will be set to the id of the root stack of the resource, heat_stack_id will be set to the id of the resource's parent stack, heat_stack_name will be set to the name of the resource's parent stack, heat_path_in_stack will be set to a list of comma delimited strings of stackresource name and stackname with list[0] being <i>rootstackname</i> , heat_resource_name will be set to the resource's name, and heat_resource_uuid will be set to the resource's orchestration id.
<b>stack_user_domain_id = None</b>	string value	Keystone domain ID which contains heat template-defined users. If this option is set, stack_user_domain_name option will be ignored.
<b>stack_user_domain_name = None</b>	string value	Keystone domain name which contains heat template-defined users. If <b>stack_user_domain_id</b> option is set, this option is ignored.
<b>subscribe_on = []</b>	list value	List of publisher hosts SubConsumer can subscribe on. This option has higher priority then the default publishers list taken from the matchmaker.
<b>syslog-log-facility = LOG_USER</b>	string value	Syslog facility to receive log lines. This option is ignored if log_config_append is set.
<b>template_dir = /etc/heat/templates</b>	string value	The directory to search for template files.

Configuration option = Default value	Type	Description
<b>transport_url = None</b>	string value	<p>The network address and optional user credentials for connecting to the messaging backend, in URL format. The expected format is:</p> <pre>driver://[user:pass@]host:port[, [userN:passN@]hostN:portN]/virtual_host?query</pre> <p>Example: rabbit://rabbitmq:password@127.0.0.1:5672//</p> <p>For full details on the fields in the URL see the documentation of oslo_messaging.TransportURL at <a href="https://docs.openstack.org/oslo.messaging/latest/reference/transport.html">https://docs.openstack.org/oslo.messaging/latest/reference/transport.html</a></p>
<b>trusts_delegated_roles = []</b>	list value	Subset of trustor roles to be delegated to heat. If left unset, all roles of a user will be delegated to heat when creating a stack.
<b>use-journal = False</b>	boolean value	Enable journald for logging. If running in a systemd environment you may wish to enable journal support. Doing so will use the journal native protocol which includes structured metadata in addition to log messages. This option is ignored if log_config_append is set.
<b>use-json = False</b>	boolean value	Use JSON formatting for logging. This option is ignored if log_config_append is set.
<b>use-syslog = False</b>	boolean value	Use syslog for logging. Existing syslog format is DEPRECATED and will be changed later to honor RFC5424. This option is ignored if log_config_append is set.
<b>use_dynamic_connections = False</b>	boolean value	This option makes direct connections dynamic or static. It makes sense only with use_router_proxy=False which means to use direct connections for direct message types (ignored otherwise).
<b>use_pub_sub = False</b>	boolean value	Use PUB/SUB pattern for fanout methods. PUB/SUB always uses proxy.
<b>use_router_proxy = False</b>	boolean value	Use ROUTER remote proxy.
<b>use_stderr = False</b>	boolean value	Log output to standard error. This option is ignored if log_config_append is set.



Configuration option = Default value	Type	Description
<b>watch-log-file = False</b>	boolean value	Uses logging handler designed to watch file system. When log file is moved or removed this handler will open a new log file with specified path instantaneously. It makes sense only if log_file option is specified and Linux platform is used. This option is ignored if log_config_append is set.
<b>zmq_failover_connections = 2</b>	integer value	How many additional connections to a host will be made for failover reasons. This option is actual only in dynamic connections mode.
<b>zmq_immediate = True</b>	boolean value	This option configures round-robin mode in zmq socket. True means not keeping a queue when server side disconnects. False means to keep queue and messages even if server is disconnected, when the server appears we send all accumulated messages to it.
<b>zmq_linger = -1</b>	integer value	Number of seconds to wait before all pending messages will be sent after closing a socket. The default value of -1 specifies an infinite linger period. The value of 0 specifies no linger period. Pending messages shall be discarded immediately when the socket is closed. Positive values specify an upper bound for the linger period.
<b>zmq_target_expire = 300</b>	integer value	Expiration timeout in seconds of a name service record about existing target (< 0 means no timeout).
<b>zmq_target_update = 180</b>	integer value	Update period in seconds of a name service record about existing target.
<b>zmq_tcp_keepalive = -1</b>	integer value	Enable/disable TCP keepalive (KA) mechanism. The default value of -1 (or any other negative value) means to skip any overrides and leave it to OS default; 0 and 1 (or any other positive value) mean to disable and enable the option respectively.
<b>zmq_tcp_keepalive_cnt = -1</b>	integer value	The number of retransmissions to be carried out before declaring that remote end is not available. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.

Configuration option = Default value	Type	Description
<b>zmq_tcp_keepalive_idle = -1</b>	integer value	The duration between two keepalive transmissions in idle condition. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_intvl = -1</b>	integer value	The duration between two successive keepalive retransmissions, if acknowledgement to the previous keepalive transmission is not received. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.

### 6.1.2. auth\_password

The following table outlines the options available under the **[auth\_password]** group in the **/etc/heat/heat.conf** file.

Table 6.2. auth\_password

Configuration option = Default value	Type	Description
<b>allowed_auth_uris = []</b>	list value	Allowed keystone endpoints for auth_uri when multi_cloud is enabled. At least one endpoint needs to be specified.
<b>multi_cloud = False</b>	boolean value	Allow orchestration of multiple clouds.

### 6.1.3. clients

The following table outlines the options available under the **[clients]** group in the **/etc/heat/heat.conf** file.

Table 6.3. clients

Configuration option = Default value	Type	Description
<b>ca_file = None</b>	string value	Optional CA cert file to use in SSL connections.
<b>cert_file = None</b>	string value	Optional PEM-formatted certificate chain file.
<b>endpoint_type = publicURL</b>	string value	Type of endpoint in Identity service catalog to use for communication with the OpenStack service.

Configuration option = Default value	Type	Description
<b>insecure = False</b>	boolean value	If set, then the server's certificate will not be verified.
<b>key_file = None</b>	string value	Optional PEM-formatted file that contains the private key.

#### 6.1.4. clients\_aodh

The following table outlines the options available under the **[clients\_aodh]** group in the `/etc/heat/heat.conf` file.

Table 6.4. clients\_aodh

Configuration option = Default value	Type	Description
<b>ca_file = None</b>	string value	Optional CA cert file to use in SSL connections.
<b>cert_file = None</b>	string value	Optional PEM-formatted certificate chain file.
<b>endpoint_type = None</b>	string value	Type of endpoint in Identity service catalog to use for communication with the OpenStack service.
<b>insecure = None</b>	boolean value	If set, then the server's certificate will not be verified.
<b>key_file = None</b>	string value	Optional PEM-formatted file that contains the private key.

#### 6.1.5. clients\_barbican

The following table outlines the options available under the **[clients\_barbican]** group in the `/etc/heat/heat.conf` file.

Table 6.5. clients\_barbican

Configuration option = Default value	Type	Description
<b>ca_file = None</b>	string value	Optional CA cert file to use in SSL connections.
<b>cert_file = None</b>	string value	Optional PEM-formatted certificate chain file.
<b>endpoint_type = None</b>	string value	Type of endpoint in Identity service catalog to use for communication with the OpenStack service.
<b>insecure = None</b>	boolean value	If set, then the server's certificate will not be verified.

Configuration option = Default value	Type	Description
<b>key_file = None</b>	string value	Optional PEM-formatted file that contains the private key.

### 6.1.6. clients\_ceilometer

The following table outlines the options available under the **[clients\_ceilometer]** group in the `/etc/heat/heat.conf` file.

Table 6.6. clients\_ceilometer

Configuration option = Default value	Type	Description
<b>ca_file = None</b>	string value	Optional CA cert file to use in SSL connections.
<b>cert_file = None</b>	string value	Optional PEM-formatted certificate chain file.
<b>endpoint_type = None</b>	string value	Type of endpoint in Identity service catalog to use for communication with the OpenStack service.
<b>insecure = None</b>	boolean value	If set, then the server's certificate will not be verified.
<b>key_file = None</b>	string value	Optional PEM-formatted file that contains the private key.

### 6.1.7. clients\_cinder

The following table outlines the options available under the **[clients\_cinder]** group in the `/etc/heat/heat.conf` file.

Table 6.7. clients\_cinder

Configuration option = Default value	Type	Description
<b>ca_file = None</b>	string value	Optional CA cert file to use in SSL connections.
<b>cert_file = None</b>	string value	Optional PEM-formatted certificate chain file.
<b>endpoint_type = None</b>	string value	Type of endpoint in Identity service catalog to use for communication with the OpenStack service.
<b>http_log_debug = False</b>	boolean value	Allow client's debug log output.
<b>insecure = None</b>	boolean value	If set, then the server's certificate will not be verified.

Configuration option = Default value	Type	Description
<b>key_file = None</b>	string value	Optional PEM-formatted file that contains the private key.

### 6.1.8. clients\_designate

The following table outlines the options available under the **[clients\_designate]** group in the **/etc/heat/heat.conf** file.

Table 6.8. clients\_designate

Configuration option = Default value	Type	Description
<b>ca_file = None</b>	string value	Optional CA cert file to use in SSL connections.
<b>cert_file = None</b>	string value	Optional PEM-formatted certificate chain file.
<b>endpoint_type = None</b>	string value	Type of endpoint in Identity service catalog to use for communication with the OpenStack service.
<b>insecure = None</b>	boolean value	If set, then the server's certificate will not be verified.
<b>key_file = None</b>	string value	Optional PEM-formatted file that contains the private key.

### 6.1.9. clients\_glance

The following table outlines the options available under the **[clients\_glance]** group in the **/etc/heat/heat.conf** file.

Table 6.9. clients\_glance

Configuration option = Default value	Type	Description
<b>ca_file = None</b>	string value	Optional CA cert file to use in SSL connections.
<b>cert_file = None</b>	string value	Optional PEM-formatted certificate chain file.
<b>endpoint_type = None</b>	string value	Type of endpoint in Identity service catalog to use for communication with the OpenStack service.
<b>insecure = None</b>	boolean value	If set, then the server's certificate will not be verified.
<b>key_file = None</b>	string value	Optional PEM-formatted file that contains the private key.

Configuration option = Default value	Type	Description
---	------	-------------

### 6.1.10. clients\_heat

The following table outlines the options available under the **[clients\_heat]** group in the **/etc/heat/heat.conf** file.

Table 6.10. clients\_heat

Configuration option = Default value	Type	Description
<b>ca_file = None</b>	string value	Optional CA cert file to use in SSL connections.
<b>cert_file = None</b>	string value	Optional PEM-formatted certificate chain file.
<b>endpoint_type = None</b>	string value	Type of endpoint in Identity service catalog to use for communication with the OpenStack service.
<b>insecure = None</b>	boolean value	If set, then the server's certificate will not be verified.
<b>key_file = None</b>	string value	Optional PEM-formatted file that contains the private key.
<code>`url = `</code>	string value	Optional heat url in format like <a href="http://0.0.0.0:8004/v1/(tenant_ids)">http://0.0.0.0:8004/v1/(tenant_ids)</a> .

### 6.1.11. clients\_keystone

The following table outlines the options available under the **[clients\_keystone]** group in the **/etc/heat/heat.conf** file.

Table 6.11. clients\_keystone

Configuration option = Default value	Type	Description
<code>`auth_uri = `</code>	string value	Unversioned keystone url in format like <a href="http://0.0.0.0:5000">http://0.0.0.0:5000</a> .
<b>ca_file = None</b>	string value	Optional CA cert file to use in SSL connections.
<b>cert_file = None</b>	string value	Optional PEM-formatted certificate chain file.

Configuration option = Default value	Type	Description
<b>endpoint_type = None</b>	string value	Type of endpoint in Identity service catalog to use for communication with the OpenStack service.
<b>insecure = None</b>	boolean value	If set, then the server's certificate will not be verified.
<b>key_file = None</b>	string value	Optional PEM-formatted file that contains the private key.

### 6.1.12. clients\_magnum

The following table outlines the options available under the **[clients\_magnum]** group in the `/etc/heat/heat.conf` file.

Table 6.12. clients\_magnum

Configuration option = Default value	Type	Description
<b>ca_file = None</b>	string value	Optional CA cert file to use in SSL connections.
<b>cert_file = None</b>	string value	Optional PEM-formatted certificate chain file.
<b>endpoint_type = None</b>	string value	Type of endpoint in Identity service catalog to use for communication with the OpenStack service.
<b>insecure = None</b>	boolean value	If set, then the server's certificate will not be verified.
<b>key_file = None</b>	string value	Optional PEM-formatted file that contains the private key.

### 6.1.13. clients\_manila

The following table outlines the options available under the **[clients\_manila]** group in the `/etc/heat/heat.conf` file.

Table 6.13. clients\_manila

Configuration option = Default value	Type	Description
<b>ca_file = None</b>	string value	Optional CA cert file to use in SSL connections.
<b>cert_file = None</b>	string value	Optional PEM-formatted certificate chain file.

Configuration option = Default value	Type	Description
<b>endpoint_type = None</b>	string value	Type of endpoint in Identity service catalog to use for communication with the OpenStack service.
<b>insecure = None</b>	boolean value	If set, then the server's certificate will not be verified.
<b>key_file = None</b>	string value	Optional PEM-formatted file that contains the private key.

#### 6.1.14. clients\_mistral

The following table outlines the options available under the **[clients\_mistral]** group in the `/etc/heat/heat.conf` file.

Table 6.14. clients\_mistral

Configuration option = Default value	Type	Description
<b>ca_file = None</b>	string value	Optional CA cert file to use in SSL connections.
<b>cert_file = None</b>	string value	Optional PEM-formatted certificate chain file.
<b>endpoint_type = None</b>	string value	Type of endpoint in Identity service catalog to use for communication with the OpenStack service.
<b>insecure = None</b>	boolean value	If set, then the server's certificate will not be verified.
<b>key_file = None</b>	string value	Optional PEM-formatted file that contains the private key.

#### 6.1.15. clients\_monasca

The following table outlines the options available under the **[clients\_monasca]** group in the `/etc/heat/heat.conf` file.

Table 6.15. clients\_monasca

Configuration option = Default value	Type	Description
<b>ca_file = None</b>	string value	Optional CA cert file to use in SSL connections.
<b>cert_file = None</b>	string value	Optional PEM-formatted certificate chain file.
<b>endpoint_type = None</b>	string value	Type of endpoint in Identity service catalog to use for communication with the OpenStack service.



Configuration option = Default value	Type	Description
<b>insecure = None</b>	boolean value	If set, then the server's certificate will not be verified.
<b>key_file = None</b>	string value	Optional PEM-formatted file that contains the private key.

### 6.1.16. clients\_neutron

The following table outlines the options available under the **[clients\_neutron]** group in the `/etc/heat/heat.conf` file.

Table 6.16. clients\_neutron

Configuration option = Default value	Type	Description
<b>ca_file = None</b>	string value	Optional CA cert file to use in SSL connections.
<b>cert_file = None</b>	string value	Optional PEM-formatted certificate chain file.
<b>endpoint_type = None</b>	string value	Type of endpoint in Identity service catalog to use for communication with the OpenStack service.
<b>insecure = None</b>	boolean value	If set, then the server's certificate will not be verified.
<b>key_file = None</b>	string value	Optional PEM-formatted file that contains the private key.

### 6.1.17. clients\_nova

The following table outlines the options available under the **[clients\_nova]** group in the `/etc/heat/heat.conf` file.

Table 6.17. clients\_nova

Configuration option = Default value	Type	Description
<b>ca_file = None</b>	string value	Optional CA cert file to use in SSL connections.
<b>cert_file = None</b>	string value	Optional PEM-formatted certificate chain file.
<b>endpoint_type = None</b>	string value	Type of endpoint in Identity service catalog to use for communication with the OpenStack service.

Configuration option = Default value	Type	Description
<b>http_log_debug = False</b>	boolean value	Allow client's debug log output.
<b>insecure = None</b>	boolean value	If set, then the server's certificate will not be verified.
<b>key_file = None</b>	string value	Optional PEM-formatted file that contains the private key.

### 6.1.18. clients\_octavia

The following table outlines the options available under the **[clients\_octavia]** group in the **/etc/heat/heat.conf** file.

Table 6.18. clients\_octavia

Configuration option = Default value	Type	Description
<b>ca_file = None</b>	string value	Optional CA cert file to use in SSL connections.
<b>cert_file = None</b>	string value	Optional PEM-formatted certificate chain file.
<b>endpoint_type = None</b>	string value	Type of endpoint in Identity service catalog to use for communication with the OpenStack service.
<b>insecure = None</b>	boolean value	If set, then the server's certificate will not be verified.
<b>key_file = None</b>	string value	Optional PEM-formatted file that contains the private key.

### 6.1.19. clients\_sahara

The following table outlines the options available under the **[clients\_sahara]** group in the **/etc/heat/heat.conf** file.

Table 6.19. clients\_sahara

Configuration option = Default value	Type	Description
<b>ca_file = None</b>	string value	Optional CA cert file to use in SSL connections.
<b>cert_file = None</b>	string value	Optional PEM-formatted certificate chain file.

Configuration option = Default value	Type	Description
<b>endpoint_type = None</b>	string value	Type of endpoint in Identity service catalog to use for communication with the OpenStack service.
<b>insecure = None</b>	boolean value	If set, then the server's certificate will not be verified.
<b>key_file = None</b>	string value	Optional PEM-formatted file that contains the private key.

### 6.1.20. clients\_senlin

The following table outlines the options available under the **[clients\_senlin]** group in the **/etc/heat/heat.conf** file.

Table 6.20. clients\_senlin

Configuration option = Default value	Type	Description
<b>ca_file = None</b>	string value	Optional CA cert file to use in SSL connections.
<b>cert_file = None</b>	string value	Optional PEM-formatted certificate chain file.
<b>endpoint_type = None</b>	string value	Type of endpoint in Identity service catalog to use for communication with the OpenStack service.
<b>insecure = None</b>	boolean value	If set, then the server's certificate will not be verified.
<b>key_file = None</b>	string value	Optional PEM-formatted file that contains the private key.

### 6.1.21. clients\_swift

The following table outlines the options available under the **[clients\_swift]** group in the **/etc/heat/heat.conf** file.

Table 6.21. clients\_swift

Configuration option = Default value	Type	Description
<b>ca_file = None</b>	string value	Optional CA cert file to use in SSL connections.
<b>cert_file = None</b>	string value	Optional PEM-formatted certificate chain file.

Configuration option = Default value	Type	Description
<b>endpoint_type = None</b>	string value	Type of endpoint in Identity service catalog to use for communication with the OpenStack service.
<b>insecure = None</b>	boolean value	If set, then the server's certificate will not be verified.
<b>key_file = None</b>	string value	Optional PEM-formatted file that contains the private key.

### 6.1.22. clients\_trove

The following table outlines the options available under the **[clients\_trove]** group in the `/etc/heat/heat.conf` file.

Table 6.22. clients\_trove

Configuration option = Default value	Type	Description
<b>ca_file = None</b>	string value	Optional CA cert file to use in SSL connections.
<b>cert_file = None</b>	string value	Optional PEM-formatted certificate chain file.
<b>endpoint_type = None</b>	string value	Type of endpoint in Identity service catalog to use for communication with the OpenStack service.
<b>insecure = None</b>	boolean value	If set, then the server's certificate will not be verified.
<b>key_file = None</b>	string value	Optional PEM-formatted file that contains the private key.

### 6.1.23. clients\_zaqar

The following table outlines the options available under the **[clients\_zaqar]** group in the `/etc/heat/heat.conf` file.

Table 6.23. clients\_zaqar

Configuration option = Default value	Type	Description
<b>ca_file = None</b>	string value	Optional CA cert file to use in SSL connections.
<b>cert_file = None</b>	string value	Optional PEM-formatted certificate chain file.
<b>endpoint_type = None</b>	string value	Type of endpoint in Identity service catalog to use for communication with the OpenStack service.

Configuration option = Default value	Type	Description
<b>insecure = None</b>	boolean value	If set, then the server's certificate will not be verified.
<b>key_file = None</b>	string value	Optional PEM-formatted file that contains the private key.

### 6.1.24. cors

The following table outlines the options available under the **[cors]** group in the `/etc/heat/heat.conf` file.

Table 6.24. cors

Configuration option = Default value	Type	Description
<b>allow_credentials = True</b>	boolean value	Indicate that the actual request can include user credentials
<b>allow_headers = [u'X-Auth-Token', u'X-Identity-Status', u'X-Roles', u'X-Service-Catalog', u'X-User-Id', u'X-Tenant-Id', u'X-OpenStack-Request-ID']</b>	list value	Indicate which header field names may be used during the actual request.
<b>allow_methods = [u'GET', u'PUT', u'POST', u'DELETE', u'PATCH']</b>	list value	Indicate which methods can be used during the actual request.
<b>allowed_origin = None</b>	list value	Indicate whether this resource may be shared with the domain received in the requests "origin" header. Format: "<protocol>://<host>[:<port>]", no trailing slash. Example: <a href="https://horizon.example.com">https://horizon.example.com</a>
<b>expose_headers = [u'X-Auth-Token', u'X-Subject-Token', u'X-Service-Token', u'X-OpenStack-Request-ID']</b>	list value	Indicate which headers are safe to expose to the API. Defaults to HTTP Simple Headers.
<b>max_age = 3600</b>	integer value	Maximum cache age of CORS preflight requests.

### 6.1.25. database

The following table outlines the options available under the **[database]** group in the `/etc/heat/heat.conf` file.

Table 6.25. database

Configuration option = Default value	Type	Description
<b>backend = sqlalchemy</b>	string value	The back end to use for the database.
<b>connection = None</b>	string value	The SQLAlchemy connection string to use to connect to the database.
<b>connection_debug = 0</b>	integer value	Verbosity of SQL debugging information: 0=None, 100=Everything.
<b>connection_recycle_time = 3600</b>	integer value	Connections which have been present in the connection pool longer than this number of seconds will be replaced with a new one the next time they are checked out from the pool.
<b>connection_trace = False</b>	boolean value	Add Python stack traces to SQL as comment strings.
<b>db_inc_retry_interval = True</b>	boolean value	If True, increases the interval between retries of a database operation up to <code>db_max_retry_interval</code> .
<b>db_max_retries = 20</b>	integer value	Maximum retries in case of connection error or deadlock error before error is raised. Set to -1 to specify an infinite retry count.
<b>db_max_retry_interval = 10</b>	integer value	If <code>db_inc_retry_interval</code> is set, the maximum seconds between retries of a database operation.
<b>db_retry_interval = 1</b>	integer value	Seconds between retries of a database transaction.
<b>max_overflow = 50</b>	integer value	If set, use this value for <code>max_overflow</code> with SQLAlchemy.
<b>max_pool_size = 5</b>	integer value	Maximum number of SQL connections to keep open in a pool. Setting a value of 0 indicates no limit.
<b>max_retries = 10</b>	integer value	Maximum number of database connection retries during startup. Set to -1 to specify an infinite retry count.
<b>min_pool_size = 1</b>	integer value	Minimum number of SQL connections to keep open in a pool.
<b>mysql_enable_ndb = False</b>	boolean value	If True, transparently enables support for handling MySQL Cluster (NDB).

Configuration option = Default value	Type	Description
<b>mysql_sql_mode = TRADITIONAL</b>	string value	The SQL mode to be used for MySQL sessions. This option, including the default, overrides any server-set SQL mode. To use whatever SQL mode is set by the server configuration, set this to no value. Example: <code>mysql_sql_mode=</code>
<b>pool_timeout = None</b>	integer value	If set, use this value for <code>pool_timeout</code> with SQLAlchemy.
<b>retry_interval = 10</b>	integer value	Interval between retries of opening a SQL connection.
<b>slave_connection = None</b>	string value	The SQLAlchemy connection string to use to connect to the slave database.
<b>sqlite_synchronous = True</b>	boolean value	If True, SQLite uses synchronous mode.
<b>use_db_reconnect = False</b>	boolean value	Enable the experimental use of database reconnect on connection lost.

### 6.1.26. ec2authtoken

The following table outlines the options available under the **[ec2authtoken]** group in the `/etc/heat/heat.conf` file.

Table 6.26. ec2authtoken

Configuration option = Default value	Type	Description
<b>allowed_auth_uris = []</b>	list value	Allowed keystone endpoints for <code>auth_uri</code> when <code>multi_cloud</code> is enabled. At least one endpoint needs to be specified.
<b>auth_uri = None</b>	string value	Authentication Endpoint URI.
<b>ca_file = None</b>	string value	Optional CA cert file to use in SSL connections.
<b>cert_file = None</b>	string value	Optional PEM-formatted certificate chain file.
<b>insecure = False</b>	boolean value	If set, then the server's certificate will not be verified.
<b>key_file = None</b>	string value	Optional PEM-formatted file that contains the private key.

Configuration option = Default value	Type	Description
<b>multi_cloud = False</b>	boolean value	Allow orchestration of multiple clouds.

### 6.1.27. eventlet\_opts

The following table outlines the options available under the **[eventlet\_opts]** group in the **/etc/heat/heat.conf** file.

Table 6.27. eventlet\_opts

Configuration option = Default value	Type	Description
<b>client_socket_timeout = 900</b>	integer value	Timeout for client connections' socket operations. If an incoming connection is idle for this number of seconds it will be closed. A value of 0 means wait forever.
<b>wsgi_keep_alive = True</b>	boolean value	If False, closes the client socket connection explicitly.

### 6.1.28. healthcheck

The following table outlines the options available under the **[healthcheck]** group in the **/etc/heat/heat.conf** file.

Table 6.28. healthcheck

Configuration option = Default value	Type	Description
<b>backends = []</b>	list value	Additional backends that can perform health checks and report that information back as part of a request.
<b>detailed = False</b>	boolean value	Show more detailed information as part of the response
<b>disable_by_file_path = None</b>	string value	Check the presence of a file to determine if an application is running on a port. Used by DisableByFileHealthcheck plugin.
<b>disable_by_file_paths = []</b>	list value	Check the presence of a file based on a port to determine if an application is running on a port. Expects a "port:path" list of strings. Used by DisableByFilesPortsHealthcheck plugin.
<b>path = /healthcheck</b>	string value	The path to respond to healthcheck requests on.



### 6.1.29. heat\_api

The following table outlines the options available under the **[heat\_api]** group in the `/etc/heat/heat.conf` file.

Table 6.29. heat\_api

Configuration option = Default value	Type	Description
<b>backlog = 4096</b>	integer value	Number of backlog requests to configure the socket with.
<b>bind_host = 0.0.0.0</b>	IP address value	Address to bind the server. Useful when selecting a particular network interface.
<b>bind_port = 8004</b>	port value	The port on which the server will listen.
<b>cert_file = None</b>	string value	Location of the SSL certificate file to use for SSL mode.
<b>key_file = None</b>	string value	Location of the SSL key file to use for enabling SSL mode.
<b>max_header_line = 16384</b>	integer value	Maximum line size of message headers to be accepted. <code>max_header_line</code> may need to be increased when using large tokens (typically those generated by the Keystone v3 API with big service catalogs).
<b>tcp_keepidle = 600</b>	integer value	The value for the socket option <code>TCP_KEEPIDLE</code> . This is the time in seconds that the connection must be idle before TCP starts sending keepalive probes.
<b>workers = 0</b>	integer value	Number of workers for Heat service. Default value 0 means, that service will start number of workers equal number of cores on server.

### 6.1.30. heat\_api\_cfn

The following table outlines the options available under the **[heat\_api\_cfn]** group in the `/etc/heat/heat.conf` file.

Table 6.30. heat\_api\_cfn

Configuration option = Default value	Type	Description
<b>backlog = 4096</b>	integer value	Number of backlog requests to configure the socket with.

Configuration option = Default value	Type	Description
<b>bind_host = 0.0.0.0</b>	IP address value	Address to bind the server. Useful when selecting a particular network interface.
<b>bind_port = 8000</b>	port value	The port on which the server will listen.
<b>cert_file = None</b>	string value	Location of the SSL certificate file to use for SSL mode.
<b>key_file = None</b>	string value	Location of the SSL key file to use for enabling SSL mode.
<b>max_header_line = 16384</b>	integer value	Maximum line size of message headers to be accepted. max_header_line may need to be increased when using large tokens (typically those generated by the Keystone v3 API with big service catalogs).
<b>tcp_keepidle = 600</b>	integer value	The value for the socket option TCP_KEEPIDLE. This is the time in seconds that the connection must be idle before TCP starts sending keepalive probes.
<b>workers = 1</b>	integer value	Number of workers for Heat service.

### 6.1.31. heat\_api\_cloudwatch

The following table outlines the options available under the **[heat\_api\_cloudwatch]** group in the `/etc/heat/heat.conf` file.

Table 6.31. heat\_api\_cloudwatch

Configuration option = Default value	Type	Description
<b>backlog = 4096</b>	integer value	Number of backlog requests to configure the socket with.
<b>bind_host = 0.0.0.0</b>	IP address value	Address to bind the server. Useful when selecting a particular network interface.
<b>bind_port = 8003</b>	port value	The port on which the server will listen.
<b>cert_file = None</b>	string value	Location of the SSL certificate file to use for SSL mode.
<b>key_file = None</b>	string value	Location of the SSL key file to use for enabling SSL mode.

Configuration option = Default value	Type	Description
<b>max_header_line = 16384</b>	integer value	Maximum line size of message headers to be accepted. max_header_line may need to be increased when using large tokens (typically those generated by the Keystone v3 API with big service catalogs.)
<b>tcp_keepidle = 600</b>	integer value	The value for the socket option TCP_KEEPIDLE. This is the time in seconds that the connection must be idle before TCP starts sending keepalive probes.
<b>workers = 1</b>	integer value	Number of workers for Heat service.

### 6.1.32. keystone\_authtoken

The following table outlines the options available under the **[keystone\_authtoken]** group in the **/etc/heat/heat.conf** file.

Table 6.32. keystone\_authtoken

Configuration option = Default value	Type	Description
<b>auth_section = None</b>	string value	Config Section from which to load plugin specific options
<b>auth_type = None</b>	string value	Authentication type to load
<b>auth_uri = None</b>	string value	Complete "public" Identity API endpoint. This endpoint should not be an "admin" endpoint, as it should be accessible by all end users. Unauthenticated clients are redirected to this endpoint to authenticate. Although this endpoint should ideally be unversioned, client support in the wild varies. If you're using a versioned v2 endpoint here, then this should <b>not</b> be the same endpoint the service user utilizes for validating tokens, because normal end users may not be able to reach that endpoint. This option is deprecated in favor of <code>www_authenticate_uri</code> and will be removed in the S release.
<b>auth_version = None</b>	string value	API version of the admin Identity API endpoint.

Configuration option = Default value	Type	Description
<b>cache = None</b>	string value	Request environment key where the Swift cache object is stored. When <code>auth_token</code> middleware is deployed with a Swift cache, use this option to have the middleware share a caching backend with swift. Otherwise, use the <b><code>memcached_servers</code></b> option instead.
<b>cafile = None</b>	string value	A PEM encoded Certificate Authority to use when verifying HTTPs connections. Defaults to system CAs.
<b>certfile = None</b>	string value	Required if identity server requires client certificate
<b>check_revocations_for_cached = False</b>	boolean value	If true, the revocation list will be checked for cached tokens. This requires that PKI tokens are configured on the identity server.
<b>delay_auth_decision = False</b>	boolean value	Do not handle authorization requests within the middleware, but delegate the authorization decision to downstream WSGI components.
<b>enforce_token_bind = permissive</b>	string value	Used to control the use and type of token binding. Can be set to: "disabled" to not check token binding. "permissive" (default) to validate binding information if the bind type is of a form known to the server and ignore it if not. "strict" like "permissive" but if the bind type is unknown the token will be rejected. "required" any form of token binding is needed to be allowed. Finally the name of a binding method that must be present in tokens.
<b>hash_algorithms = [u'md5']</b>	list value	Hash algorithms to use for hashing PKI tokens. This may be a single algorithm or multiple. The algorithms are those supported by Python standard <code>hashlib.new()</code> . The hashes will be tried in the order given, so put the preferred one first for performance. The result of the first hash will be stored in the cache. This will typically be set to multiple values only while migrating from a less secure algorithm to a more secure one. Once all the old tokens are expired this option should be set to a single value for better performance.
<b>http_connect_timeout = None</b>	integer value	Request timeout value for communicating with Identity API server.

Configuration option = Default value	Type	Description
<b>http_request_max_retries = 3</b>	integer value	How many times are we trying to reconnect when communicating with Identity API Server.
<b>include_service_catalog = True</b>	boolean value	(Optional) Indicate whether to set the X-Service-Catalog header. If False, middleware will not ask for service catalog on token validation and will not set the X-Service-Catalog header.
<b>insecure = False</b>	boolean value	Verify HTTPS connections.
<b>keyfile = None</b>	string value	Required if identity server requires client certificate
<b>memcache_pool_connection_timeout = 10</b>	integer value	(Optional) Number of seconds that an operation will wait to get a memcached client connection from the pool.
<b>memcache_pool_dead_retry = 300</b>	integer value	(Optional) Number of seconds memcached server is considered dead before it is tried again.
<b>memcache_pool_maxsize = 10</b>	integer value	(Optional) Maximum total number of open connections to every memcached server.
<b>memcache_pool_socket_timeout = 3</b>	integer value	(Optional) Socket timeout in seconds for communicating with a memcached server.
<b>memcache_pool_unused_timeout = 60</b>	integer value	(Optional) Number of seconds a connection to memcached is held unused in the pool before it is closed.
<b>memcache_secret_key = None</b>	string value	(Optional, mandatory if memcache_security_strategy is defined) This string is used for key derivation.
<b>memcache_security_strategy = None</b>	string value	(Optional) If defined, indicate whether token data should be authenticated or authenticated and encrypted. If MAC, token data is authenticated (with HMAC) in the cache. If ENCRYPT, token data is encrypted and authenticated in the cache. If the value is not one of these options or empty, auth_token will raise an exception on initialization.
<b>memcache_use_advanced_pool = False</b>	boolean value	(Optional) Use the advanced (eventlet safe) memcached client pool. The advanced pool will only work under python 2.x.
<b>memcached_servers = None</b>	list value	Optionally specify a list of memcached server(s) to use for caching. If left undefined, tokens will instead be cached in-process.

Configuration option = Default value	Type	Description
<b>region_name = None</b>	string value	The region in which the identity server can be found.
<b>revocation_cache_time = 10</b>	integer value	Determines the frequency at which the list of revoked tokens is retrieved from the Identity service (in seconds). A high number of revocation events combined with a low cache duration may significantly reduce performance. Only valid for PKI tokens. This option has been deprecated in the Ocata release and will be removed in the P release.
<b>service_token_roles = [u'service']</b>	list value	A choice of roles that must be present in a service token. Service tokens are allowed to request that an expired token can be used and so this check should tightly control that only actual services should be sending this token. Roles here are applied as an ANY check so any role in this list must be present. For backwards compatibility reasons this currently only affects the allow_expired check.
<b>service_token_roles_required = False</b>	boolean value	For backwards compatibility reasons we must let valid service tokens pass that don't pass the service_token_roles check as valid. Setting this true will become the default in a future release and should be enabled if possible.
<b>signing_dir = None</b>	string value	Directory used to cache files related to PKI tokens. This option has been deprecated in the Ocata release and will be removed in the P release.
<b>token_cache_time = 300</b>	integer value	In order to prevent excessive effort spent validating tokens, the middleware caches previously-seen tokens for a configurable duration (in seconds). Set to -1 to disable caching completely.
<b>www_authenticate_uri = None</b>	string value	Complete "public" Identity API endpoint. This endpoint should not be an "admin" endpoint, as it should be accessible by all end users. Unauthenticated clients are redirected to this endpoint to authenticate. Although this endpoint should ideally be unversioned, client support in the wild varies. If you're using a versioned v2 endpoint here, then this should <b>not</b> be the same endpoint the service user utilizes for validating tokens, because normal end users may not be able to reach that endpoint.

### 6.1.33. matchmaker\_redis

The following table outlines the options available under the **[matchmaker\_redis]** group in the `/etc/heat/heat.conf` file.

Table 6.33. `matchmaker_redis`

Configuration option = Default value	Type	Description
<b>check_timeout = 20000</b>	integer value	Time in ms to wait before the transaction is killed.
<b>host = 127.0.0.1</b>	string value	Host to locate redis.
<code>`password = `</code>	string value	Password for Redis server (optional).
<b>port = 6379</b>	port value	Use this port to connect to redis host.
<b>sentinel_group_name = oslo-messaging-zeromq</b>	string value	Redis replica set name.
<b>sentinel_hosts = []</b>	list value	List of Redis Sentinel hosts (fault tolerance mode), e.g., [host:port, host1:port ... ]
<b>socket_timeout = 10000</b>	integer value	Timeout in ms on blocking socket operations.
<b>wait_timeout = 2000</b>	integer value	Time in ms to wait between connection attempts.

### 6.1.34. noauth

The following table outlines the options available under the **[noauth]** group in the `/etc/heat/heat.conf` file.

Table 6.34. `noauth`

Configuration option = Default value	Type	Description
<code>`token_response = `</code>	string value	JSON file containing the content returned by the noauth middleware.

### 6.1.35. oslo\_messaging\_amqp

The following table outlines the options available under the **[oslo\_messaging\_amqp]** group in the `/etc/heat/heat.conf` file.

Table 6.35. `oslo_messaging_amqp`

Configuration option = Default value	Type	Description
<b>addressing_mode = dynamic</b>	string value	Indicates the addressing mode used by the driver. Permitted values: <i>legacy</i> - use legacy non-routable addressing <i>routable</i> - use routable addresses <i>dynamic</i> - use legacy addresses if the message bus does not support routing otherwise use routable addressing
<b>allow_insecure_clients = False</b>	boolean value	Accept clients using either SSL or plain TCP
<b>anycast_address = anycast</b>	string value	Appended to the address prefix when sending to a group of consumers. Used by the message bus to identify messages that should be delivered in a round-robin fashion across consumers.
<b>broadcast_prefix = broadcast</b>	string value	address prefix used when broadcasting to all servers
<b>connection_retry_backoff = 2</b>	integer value	Increase the <code>connection_retry_interval</code> by this many seconds after each unsuccessful failover attempt.
<b>connection_retry_interval = 1</b>	integer value	Seconds to pause before attempting to re-connect.
<b>connection_retry_interval _max = 30</b>	integer value	Maximum limit for <code>connection_retry_interval</code> + <code>connection_retry_backoff</code>
<b>container_name = None</b>	string value	Name for the AMQP container. must be globally unique. Defaults to a generated UUID
<b>default_notification_exch ange = None</b>	string value	Exchange name used in notification addresses. Exchange name resolution precedence: Target.exchange if set else default_notification_exchange if set else control_exchange if set else <i>notify</i>
<b>default_notify_timeout = 30</b>	integer value	The deadline for a sent notification message delivery. Only used when caller does not provide a timeout expiry.
<b>default_reply_retry = 0</b>	integer value	The maximum number of attempts to re-send a reply message which failed due to a recoverable error.
<b>default_reply_timeout = 30</b>	integer value	The deadline for an rpc reply message delivery.



Configuration option = Default value	Type	Description
<b>default_rpc_exchange = None</b>	string value	Exchange name used in RPC addresses. Exchange name resolution precedence: Target.exchange if set else default_rpc_exchange if set else control_exchange if set else <i>rpc</i>
<b>default_send_timeout = 30</b>	integer value	The deadline for an rpc cast or call message delivery. Only used when caller does not provide a timeout expiry.
<b>default_sender_link_timeout = 600</b>	integer value	The duration to schedule a purge of idle sender links. Detach link after expiry.
<b>group_request_prefix = unicast</b>	string value	address prefix when sending to any server in group
<b>idle_timeout = 0</b>	integer value	Timeout for inactive connections (in seconds)
<b>link_retry_delay = 10</b>	integer value	Time to pause between re-connecting an AMQP 1.0 link that failed due to a recoverable error.
<b>multicast_address = multicast</b>	string value	Appended to the address prefix when sending a fanout message. Used by the message bus to identify fanout messages.
<b>notify_address_prefix = openstack.org/om/notify</b>	string value	Address prefix for all generated Notification addresses
<b>notify_server_credit = 100</b>	integer value	Window size for incoming Notification messages
<b>password = `</b>	string value	Password for message broker authentication
<b>pre_settled = [u'rpc-cast', u'rpc-reply']</b>	multi valued	Send messages of this type pre-settled. Pre-settled messages will not receive acknowledgement from the peer. Note well: pre-settled messages may be silently discarded if the delivery fails. Permitted values: <i>rpc-call</i> - send RPC Calls pre-settled <i>rpc-reply</i> - send RPC Replies pre-settled <i>rpc-cast</i> - Send RPC Casts pre-settled <i>notify</i> - Send Notifications pre-settled

Configuration option = Default value	Type	Description
<b>pseudo_vhost = True</b>	boolean value	Enable virtual host support for those message buses that do not natively support virtual hosting (such as qpid). When set to true the virtual host name will be added to all message bus addresses, effectively creating a private <i>subnet</i> per virtual host. Set to False if the message bus supports virtual hosting using the <i>hostname</i> field in the AMQP 1.0 Open performative as the name of the virtual host.
<b>reply_link_credit = 200</b>	integer value	Window size for incoming RPC Reply messages.
<b>rpc_address_prefix = openstack.org/om/rpc</b>	string value	Address prefix for all generated RPC addresses
<b>rpc_server_credit = 100</b>	integer value	Window size for incoming RPC Request messages
<code>`sasl_config_dir = `</code>	string value	Path to directory that contains the SASL configuration
<code>`sasl_config_name = `</code>	string value	Name of configuration file (without .conf suffix)
<code>`sasl_default_realm = `</code>	string value	SASL realm to use if no realm present in username
<code>`sasl_mechanisms = `</code>	string value	Space separated list of acceptable SASL mechanisms
<b>server_request_prefix = exclusive</b>	string value	address prefix used when sending to a specific server
<b>ssl = False</b>	boolean value	Attempt to connect via SSL. If no other ssl-related parameters are given, it will use the system's CA-bundle to verify the server's certificate.
<code>`ssl_ca_file = `</code>	string value	CA certificate PEM file used to verify the server's certificate
<code>`ssl_cert_file = `</code>	string value	Self-identifying certificate PEM file for client authentication
<code>`ssl_key_file = `</code>	string value	Private key PEM file used to sign <code>ssl_cert_file</code> certificate (optional)
<b>ssl_key_password = None</b>	string value	Password for decrypting <code>ssl_key_file</code> (if encrypted)

Configuration option = Default value	Type	Description
<b>ssl_verify_vhost = False</b>	boolean value	By default SSL checks that the name in the server's certificate matches the hostname in the <code>transport_url</code> . In some configurations it may be preferable to use the virtual hostname instead, for example if the server uses the Server Name Indication TLS extension (rfc6066) to provide a certificate per virtual host. Set <code>ssl_verify_vhost</code> to True if the server's SSL certificate uses the virtual host name instead of the DNS name.
<b>trace = False</b>	boolean value	Debug: dump AMQP frames to stdout
<b>unicast_address = unicast</b>	string value	Appended to the address prefix when sending to a particular RPC/Notification server. Used by the message bus to identify messages sent to a single destination.
<code>`username = `</code>	string value	User name for message broker authentication

### 6.1.36. oslo\_messaging\_kafka

The following table outlines the options available under the `[oslo_messaging_kafka]` group in the `/etc/heat/heat.conf` file.

Table 6.36. oslo\_messaging\_kafka

Configuration option = Default value	Type	Description
<b>conn_pool_min_size = 2</b>	integer value	The pool size limit for connections expiration policy
<b>conn_pool_ttl = 1200</b>	integer value	The time-to-live in sec of idle connections in the pool
<b>consumer_group = oslo_messaging_consumer</b>	string value	Group id for Kafka consumer. Consumers in one group will coordinate message consumption
<b>kafka_consumer_timeout = 1.0</b>	floating point value	Default timeout(s) for Kafka consumers
<b>kafka_default_host = localhost</b>	string value	Default Kafka broker Host
<b>kafka_default_port = 9092</b>	port value	Default Kafka broker Port
<b>kafka_max_fetch_bytes = 1048576</b>	integer value	Max fetch bytes of Kafka consumer

Configuration option = Default value	Type	Description
<b>pool_size = 10</b>	integer value	Pool Size for Kafka Consumers
<b>producer_batch_size = 16384</b>	integer value	Size of batch for the producer async send
<b>producer_batch_timeout = 0.0</b>	floating point value	Upper bound on the delay for KafkaProducer batching in seconds

### 6.1.37. oslo\_messaging\_notifications

The following table outlines the options available under the **[oslo\_messaging\_notifications]** group in the `/etc/heat/heat.conf` file.

Table 6.37. oslo\_messaging\_notifications

Configuration option = Default value	Type	Description
<b>driver = []</b>	multi valued	The Drivers(s) to handle sending notifications. Possible values are messaging, messagingv2, routing, log, test, noop
<b>retry = -1</b>	integer value	The maximum number of attempts to re-send a notification message which failed to be delivered due to a recoverable error. 0 - No retry, -1 - indefinite
<b>topics = [u'notifications']</b>	list value	AMQP topic used for OpenStack notifications.
<b>transport_url = None</b>	string value	A URL representing the messaging driver to use for notifications. If not set, we fall back to the same configuration used for RPC.

### 6.1.38. oslo\_messaging\_rabbit

The following table outlines the options available under the **[oslo\_messaging\_rabbit]** group in the `/etc/heat/heat.conf` file.

Table 6.38. oslo\_messaging\_rabbit

Configuration option = Default value	Type	Description
<b>amqp_auto_delete = False</b>	boolean value	Auto-delete queues in AMQP.
<b>amqp_durable_queues = False</b>	boolean value	Use durable queues in AMQP.

Configuration option = Default value	Type	Description
<b>channel_max = None</b>	integer value	Maximum number of channels to allow
<b>connection_factory = single</b>	string value	Connection factory implementation
<b>default_notification_exchange = \${control_exchange}_notification</b>	string value	Exchange name for sending notifications
<b>default_notification_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending notification, -1 means infinite retry.
<b>default_rpc_exchange = \${control_exchange}_rpc</b>	string value	Exchange name for sending RPC messages
<b>default_rpc_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending RPC message, -1 means infinite retry. If actual retry attempts in not 0 the rpc request could be processed more than one time
<b>default_serializer_type = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>fake_rabbit = False</b>	boolean value	Deprecated, use <code>rpc_backend=kombu+memory</code> or <code>rpc_backend=fake</code>
<b>frame_max = None</b>	integer value	The maximum byte size for an AMQP frame
<b>heartbeat_interval = 3</b>	integer value	How often to send heartbeats for consumer's connections
<b>heartbeat_rate = 2</b>	integer value	How often times during the <code>heartbeat_timeout_threshold</code> we check the heartbeat.
<b>heartbeat_timeout_threshold = 60</b>	integer value	Number of seconds after which the Rabbit broker is considered down if heartbeat's keep-alive fails (0 disable the heartbeat). EXPERIMENTAL
<b>host_connection_reconnect_delay = 0.25</b>	floating point value	Set delay for reconnection to some host which has connection error

Configuration option = Default value	Type	Description
<b>kombu_compression = None</b>	string value	EXPERIMENTAL: Possible values are: gzip, bz2. If not set compression will not be used. This option may not be available in future versions.
<b>kombu_failover_strategy = round-robin</b>	string value	Determines how the next RabbitMQ node is chosen in case the one we are currently connected to becomes unavailable. Takes effect only if more than one RabbitMQ node is provided in config.
<b>kombu_missing_consumer_retry_timeout = 60</b>	integer value	How long to wait a missing client before abandoning to send it its replies. This value should not be longer than <code>rpc_response_timeout</code> .
<b>kombu_reconnect_delay = 1.0</b>	floating point value	How long to wait before reconnecting in response to an AMQP consumer cancel notification.
<b>notification_listener_prefetch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to notification listener.
<b>notification_persistence = False</b>	boolean value	Persist notification messages.
<b>notification_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending notification message
<b>pool_max_overflow = 0</b>	integer value	Maximum number of connections to create above <b>pool_max_size</b> .
<b>pool_max_size = 30</b>	integer value	Maximum number of connections to keep queued.
<b>pool_recycle = 600</b>	integer value	Lifetime of a connection (since creation) in seconds or None for no recycling. Expired connections are closed on acquire.
<b>pool_stale = 60</b>	integer value	Threshold at which inactive (since release) connections are considered stale in seconds or None for no staleness. Stale connections are closed on acquire.
<b>pool_timeout = 30</b>	integer value	Default number of seconds to wait for a connections to available

Configuration option = Default value	Type	Description
<b>rabbit_ha_queues = False</b>	boolean value	Try to use HA queues in RabbitMQ (x-ha-policy: all). If you change this option, you must wipe the RabbitMQ database. In RabbitMQ 3.0, queue mirroring is no longer controlled by the x-ha-policy argument when declaring a queue. If you just want to make sure that all queues (except those with auto-generated names) are mirrored across all nodes, run: <code>"rabbitmqctl set_policy HA ^(?!amq\.).* {"ha-mode": "all"}"</code>
<b>rabbit_host = localhost</b>	string value	The RabbitMQ broker address where a single node is used.
<b>rabbit_hosts = [u'\$rabbit_host:\$rabbit_port']</b>	list value	RabbitMQ HA cluster host:port pairs.
<b>rabbit_interval_max = 30</b>	integer value	Maximum interval of RabbitMQ connection retries. Default is 30 seconds.
<b>rabbit_login_method = AMQPLAIN</b>	string value	The RabbitMQ login method.
<b>rabbit_max_retries = 0</b>	integer value	Maximum number of RabbitMQ connection retries. Default is 0 (infinite retry count).
<b>rabbit_password = guest</b>	string value	The RabbitMQ password.
<b>rabbit_port = 5672</b>	port value	The RabbitMQ broker port where a single node is used.
<b>rabbit_qos_prefetch_count = 0</b>	integer value	Specifies the number of messages to prefetch. Setting to zero allows unlimited messages.
<b>rabbit_retry_backoff = 2</b>	integer value	How long to backoff for between retries when connecting to RabbitMQ.
<b>rabbit_retry_interval = 1</b>	integer value	How frequently to retry connecting with RabbitMQ.
<b>rabbit_transient_queues_ttl = 1800</b>	integer value	Positive integer representing duration in seconds for queue TTL (x-expires). Queues which are unused for the duration of the TTL are automatically deleted. The parameter affects only reply and fanout queues.
<b>rabbit_userid = guest</b>	string value	The RabbitMQ userid.
<b>rabbit_virtual_host = /</b>	string value	The RabbitMQ virtual host.

Configuration option = Default value	Type	Description
<b>rpc_listener_prefetch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to rpc listener.
<b>rpc_queue_expiration = 60</b>	integer value	Time to live for rpc queues without consumers in seconds.
<b>rpc_reply_exchange = \${control_exchange}_rpc_reply</b>	string value	Exchange name for receiving RPC replies
<b>rpc_reply_listener_prefetch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to rpc reply listener.
<b>rpc_reply_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending reply. -1 means infinite retry during rpc_timeout
<b>rpc_reply_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending reply.
<b>rpc_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending RPC message
<b>socket_timeout = 0.25</b>	floating point value	Set socket timeout in seconds for connection's socket
<b>ssl = None</b>	boolean value	Enable SSL
<code>`ssl_ca_file = `</code>	string value	SSL certification authority file (valid only if SSL enabled).
<code>`ssl_cert_file = `</code>	string value	SSL cert file (valid only if SSL enabled).
<code>`ssl_key_file = `</code>	string value	SSL key file (valid only if SSL enabled).
<b>ssl_options = None</b>	dict value	Arguments passed to ssl.wrap_socket
<code>`ssl_version = `</code>	string value	SSL version to use (valid only if SSL enabled). Valid values are TLSv1 and SSLv23. SSLv2, SSLv3, TLSv1_1, and TLSv1_2 may be available on some distributions.
<b>tcp_user_timeout = 0.25</b>	floating point value	Set TCP_USER_TIMEOUT in seconds for connection's socket

### 6.1.39. oslo\_messaging\_zmq



The following table outlines the options available under the `[oslo_messaging_zmq]` group in the `/etc/heat/heat.conf` file.

Table 6.39. `oslo_messaging_zmq`

Configuration option = Default value	Type	Description
<code>rpc_ack_timeout_base = 15</code>	integer value	Number of seconds to wait for an ack from a cast/call. After each retry attempt this timeout is multiplied by some specified multiplier.
<code>rpc_ack_timeout_multiplier = 2</code>	integer value	Number to multiply base ack timeout by after each retry attempt.
<code>rpc_message_ttl = 300</code>	integer value	Expiration timeout in seconds of a sent/received message after which it is not tracked anymore by a client/server.
<code>rpc_poll_timeout = 1</code>	integer value	The default number of seconds that poll should wait. Poll raises timeout exception when timeout expired.
<code>rpc_retry_attempts = 3</code>	integer value	Default number of message sending attempts in case of any problems occurred: positive value N means at most N retries, 0 means no retries, None or -1 (or any other negative values) mean to retry forever. This option is used only if acknowledgments are enabled.
<code>rpc_thread_pool_size = 100</code>	integer value	Maximum number of (green) threads to work concurrently.
<code>rpc_use_acks = False</code>	boolean value	Wait for message acknowledgements from receivers. This mechanism works only via proxy without PUB/SUB.
<code>rpc_zmq_bind_address = *</code>	string value	ZeroMQ bind address. Should be a wildcard (*), an ethernet interface, or IP. The "host" option should point or resolve to this address.
<code>rpc_zmq_bind_port_retries = 100</code>	integer value	Number of retries to find free port number before fail with ZMQBindError.
<code>rpc_zmq_contexts = 1</code>	integer value	Number of ZeroMQ contexts, defaults to 1.
<code>rpc_zmq_host = ccso-sp-undercloud.usersys.redhat.com</code>	string value	Name of this node. Must be a valid hostname, FQDN, or IP address. Must match "host" option, if running Nova.
<code>rpc_zmq_ipc_dir = /var/run/openstack</code>	string value	Directory for holding IPC sockets.

Configuration option = Default value	Type	Description
<b>rpc_zmq_matchmaker = redis</b>	string value	MatchMaker driver.
<b>rpc_zmq_max_port = 65536</b>	integer value	Maximal port number for random ports range.
<b>rpc_zmq_min_port = 49153</b>	port value	Minimal port number for random ports range.
<b>rpc_zmq_serialization = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>rpc_zmq_topic_backlog = None</b>	integer value	Maximum number of ingress messages to locally buffer per topic. Default is unlimited.
<b>subscribe_on = []</b>	list value	List of publisher hosts SubConsumer can subscribe on. This option has higher priority then the default publishers list taken from the matchmaker.
<b>use_dynamic_connections = False</b>	boolean value	This option makes direct connections dynamic or static. It makes sense only with <code>use_router_proxy=False</code> which means to use direct connections for direct message types (ignored otherwise).
<b>use_pub_sub = False</b>	boolean value	Use PUB/SUB pattern for fanout methods. PUB/SUB always uses proxy.
<b>use_router_proxy = False</b>	boolean value	Use ROUTER remote proxy.
<b>zmq_failover_connections = 2</b>	integer value	How many additional connections to a host will be made for failover reasons. This option is actual only in dynamic connections mode.
<b>zmq_immediate = True</b>	boolean value	This option configures round-robin mode in zmq socket. True means not keeping a queue when server side disconnects. False means to keep queue and messages even if server is disconnected, when the server appears we send all accumulated messages to it.

Configuration option = Default value	Type	Description
<b>zmq_linger = -1</b>	integer value	Number of seconds to wait before all pending messages will be sent after closing a socket. The default value of -1 specifies an infinite linger period. The value of 0 specifies no linger period. Pending messages shall be discarded immediately when the socket is closed. Positive values specify an upper bound for the linger period.
<b>zmq_target_expire = 300</b>	integer value	Expiration timeout in seconds of a name service record about existing target (< 0 means no timeout).
<b>zmq_target_update = 180</b>	integer value	Update period in seconds of a name service record about existing target.
<b>zmq_tcp_keepalive = -1</b>	integer value	Enable/disable TCP keepalive (KA) mechanism. The default value of -1 (or any other negative value) means to skip any overrides and leave it to OS default; 0 and 1 (or any other positive value) mean to disable and enable the option respectively.
<b>zmq_tcp_keepalive_cnt = -1</b>	integer value	The number of retransmissions to be carried out before declaring that remote end is not available. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_idle = -1</b>	integer value	The duration between two keepalive transmissions in idle condition. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_intvl = -1</b>	integer value	The duration between two successive keepalive retransmissions, if acknowledgement to the previous keepalive transmission is not received. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.

#### 6.1.40. oslo\_middlewares

The following table outlines the options available under the **[oslo\_middlewares]** group in the **/etc/heat/heat.conf** file.

Table 6.40. oslo\_middlewares

Configuration option = Default value	Type	Description
<b>enable_proxy_headers_parsing = False</b>	boolean value	Whether the application is behind a proxy or not. This determines if the middleware should parse the headers or not.
<b>max_request_body_size = 114688</b>	integer value	The maximum body size for each request, in bytes.
<b>secure_proxy_ssl_header = X-Forwarded-Proto</b>	string value	The HTTP Header that will be used to determine what the original request protocol scheme was, even if it was hidden by a SSL termination proxy.

### 6.1.41. oslo\_policy

The following table outlines the options available under the **[oslo\_policy]** group in the `/etc/heat/heat.conf` file.

Table 6.41. oslo\_policy

Configuration option = Default value	Type	Description
<b>enforce_scope = False</b>	boolean value	This option controls whether or not to enforce scope when evaluating policies. If <b>True</b> , the scope of the token used in the request is compared to the <b>scope_types</b> of the policy being enforced. If the scopes do not match, an <b>InvalidScope</b> exception will be raised. If <b>False</b> , a message will be logged informing operators that policies are being invoked with mismatching scope.
<b>policy_default_rule = default</b>	string value	Default rule. Enforced when a requested rule is not found.
<b>policy_dirs = [u'policy.d']</b>	multi valued	Directories where policy configuration files are stored. They can be relative to any directory in the search path defined by the <code>config_dir</code> option, or absolute paths. The file defined by <code>policy_file</code> must exist for these directories to be searched. Missing or empty directories are ignored.
<b>policy_file = policy.json</b>	string value	The file that defines policies.
<b>remote_content_type = application/x-www-form-urlencoded</b>	string value	Content Type to send and receive data for REST based policy check

Configuration option = Default value	Type	Description
<b>remote_ssl_ca_cert_file = None</b>	string value	Absolute path to ca cert file for REST based policy check
<b>remote_ssl_client_cert_file = None</b>	string value	Absolute path to client cert for REST based policy check
<b>remote_ssl_client_key_file = None</b>	string value	Absolute path client key file REST based policy check
<b>remote_ssl_verify_server_cert = False</b>	boolean value	server identity verification for REST based policy check

### 6.1.42. paste\_deploy

The following table outlines the options available under the **[paste\_deploy]** group in the `/etc/heat/heat.conf` file.

Table 6.42. paste\_deploy

Configuration option = Default value	Type	Description
<b>api_paste_config = api-paste.ini</b>	string value	The API paste config file to use.
<b>flavor = None</b>	string value	The flavor to use.

### 6.1.43. profiler

The following table outlines the options available under the **[profiler]** group in the `/etc/heat/heat.conf` file.

Table 6.43. profiler

Configuration option = Default value	Type	Description
---	------	-------------

Configuration option = Default value	Type	Description
<b>connection_string = messaging://</b>	string value	<p>Connection string for a notifier backend. Default value is messaging:// which sets the notifier to oslo_messaging.</p> <p>Examples of possible values:</p> <ul style="list-style-type: none"> <li>● messaging://: use oslo_messaging driver for sending notifications.</li> <li>● mongodb://127.0.0.1:27017 : use mongodb driver for sending notifications.</li> <li>● elasticsearch://127.0.0.1:9200 : use elasticsearch driver for sending notifications.</li> </ul>
<b>enabled = False</b>	boolean value	<p>Enables the profiling for all services on this node. Default value is False (fully disable the profiling feature).</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True: Enables the feature</li> <li>● False: Disables the feature. The profiling cannot be started via this project operations. If the profiling is triggered by another project, this project part will be empty.</li> </ul>
<b>es_doc_type = notification</b>	string value	Document type for notification indexing in elasticsearch.
<b>es_scroll_size = 10000</b>	integer value	Elasticsearch splits large requests in batches. This parameter defines maximum size of each batch (for example: es_scroll_size=10000).
<b>es_scroll_time = 2m</b>	string value	This parameter is a time value parameter (for example: es_scroll_time=2m), indicating for how long the nodes that participate in the search will maintain relevant resources in order to continue and support it.

Configuration option = Default value	Type	Description
<b>hmac_keys = SECRET_KEY</b>	string value	<p>Secret key(s) to use for encrypting context data for performance profiling. This string value should have the following format: &lt;key1&gt;[,&lt;key2&gt;,...&lt;keyn&gt;], where each key is some random string. A user who triggers the profiling via the REST API has to set one of these keys in the headers of the REST API call to include profiling results of this node for this particular project.</p> <p>Both "enabled" flag and "hmac_keys" config options should be set to enable profiling. Also, to generate correct profiling information across all services at least one key needs to be consistent between OpenStack projects. This ensures it can be used from client side to generate the trace, containing information from all possible resources.</p>
<b>sentinel_service_name = mymaster</b>	string value	<p>Redis sentinel uses a service name to identify a master redis service. This parameter defines the name (for example: sentinel_service_name=mymaster).</p>
<b>socket_timeout = 0.1</b>	floating point value	<p>Redis sentinel provides a timeout option on the connections. This parameter defines that timeout (for example: socket_timeout=0.1).</p>
<b>trace_sqlalchemy = False</b>	boolean value	<p>Enables SQL requests profiling in services. Default value is False (SQL requests won't be traced).</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• True: Enables SQL requests profiling. Each SQL query will be part of the trace and can be analyzed by how much time was spent for that.</li> <li>• False: Disables SQL requests profiling. The spent time is only shown on a higher level of operations. Single SQL queries cannot be analyzed this way.</li> </ul>

#### 6.1.44. revision

The following table outlines the options available under the **[revision]** group in the **/etc/heat/heat.conf** file.

**Table 6.44. revision**

Configuration option = Default value	Type	Description
<b>heat_revision = unknown</b>	string value	Heat build revision. If you would prefer to manage your build revision separately, you can move this section to a different file and add it as another config option.

### 6.1.45. ssl

The following table outlines the options available under the **[ssl]** group in the `/etc/heat/heat.conf` file.

Table 6.45. ssl

Configuration option = Default value	Type	Description
<b>ca_file = None</b>	string value	CA certificate file to use to verify connecting clients.
<b>cert_file = None</b>	string value	Certificate file to use when starting the server securely.
<b>ciphers = None</b>	string value	Sets the list of available ciphers. value should be a string in the OpenSSL cipher list format.
<b>key_file = None</b>	string value	Private key file to use when starting the server securely.
<b>version = None</b>	string value	SSL version to use (valid only if SSL enabled). Valid values are TLSv1 and SSLv23. SSLv2, SSLv3, TLSv1_1, and TLSv1_2 may be available on some distributions.

### 6.1.46. trustee

The following table outlines the options available under the **[trustee]** group in the `/etc/heat/heat.conf` file.

Table 6.46. trustee

Configuration option = Default value	Type	Description
<b>auth-url = None</b>	string value	Authentication URL
<b>auth_section = None</b>	string value	Config Section from which to load plugin specific options
<b>auth_type = None</b>	string value	Authentication type to load



Configuration option = Default value	Type	Description
<b>default-domain-id = None</b>	string value	Optional domain ID to use with v3 and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.
<b>default-domain-name = None</b>	string value	Optional domain name to use with v3 API and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.
<b>domain-id = None</b>	string value	Domain ID to scope to
<b>domain-name = None</b>	string value	Domain name to scope to
<b>password = None</b>	string value	User's password
<b>project-domain-id = None</b>	string value	Domain ID containing project
<b>project-domain-name = None</b>	string value	Domain name containing project
<b>project-id = None</b>	string value	Project ID to scope to
<b>project-name = None</b>	string value	Project name to scope to
<b>system-scope = None</b>	string value	Scope for system operations
<b>trust-id = None</b>	string value	Trust ID
<b>user-domain-id = None</b>	string value	User's domain id
<b>user-domain-name = None</b>	string value	User's domain name
<b>user-id = None</b>	string value	User id
<b>username = None</b>	string value	Username

### 6.1.47. volumes

The following table outlines the options available under the **[volumes]** group in the `/etc/heat/heat.conf` file.

Table 6.47. volumes

Configuration option = Default value	Type	Description
<b>backups_enabled = True</b>	boolean value	Indicate if cinder-backup service is enabled. This is a temporary workaround until cinder-backup service becomes discoverable, see LP#1334856.

## CHAPTER 7. IRONIC

The following chapter contains information about the configuration options in the **ironic** service.

### 7.1. IRONIC.CONF

This section contains options for the `/etc/ironic/ironic.conf` file.

#### 7.1.1. DEFAULT

The following table outlines the options available under the **[DEFAULT]** group in the `/etc/ironic/ironic.conf` file.

Table 7.1. DEFAULT

Configuration option = Default value	Type	Description
<b>auth_strategy = keystone</b>	string value	Authentication strategy used by ironic-api. "noauth" should not be used in a production environment because all authentication will be disabled.
<b>backdoor_port = None</b>	string value	Enable eventlet backdoor. Acceptable values are 0, <port>, and <start>:<end>, where 0 results in listening on a random tcp port number; <port> results in listening on the specified port number (and not enabling backdoor if that port is in use); and <start>:<end> results in listening on the smallest unused port number within the specified range of port numbers. The chosen port is displayed in the service's log file.
<b>backdoor_socket = None</b>	string value	Enable eventlet backdoor, using the provided path as a unix socket that can receive connections. This option is mutually exclusive with <code>backdoor_port</code> in that only one should be provided. If both are provided then the existence of this option overrides the usage of that option.
<b>bindir = \$pybasedir/bin</b>	string value	Directory where ironic binaries are installed.
<b>conn_pool_min_size = 2</b>	integer value	The pool size limit for connections expiration policy
<b>conn_pool_ttl = 1200</b>	integer value	The time-to-live in sec of idle connections in the pool
<b>control_exchange = openstack</b>	string value	The default exchange under which topics are scoped. May be overridden by an exchange name specified in the <code>transport_url</code> option.
<b>debug = False</b>	boolean value	If set to true, the logging level will be set to DEBUG instead of the default INFO level.

Configuration option = Default value	Type	Description
<b>debug_tracebacks_in_api</b> = <b>False</b>	boolean value	Return server tracebacks in the API response for any error responses. WARNING: this is insecure and should not be used in a production environment.
<b>default_boot_interface</b> = <b>None</b>	string value	Default boot interface to be used for nodes that do not have boot_interface field set. A complete list of boot interfaces present on your system may be found by enumerating the "ironic.hardware.interfaces.boot" endpoint.
<b>default_console_interface</b> = <b>None</b>	string value	Default console interface to be used for nodes that do not have console_interface field set. A complete list of console interfaces present on your system may be found by enumerating the "ironic.hardware.interfaces.console" endpoint.
<b>default_deploy_interface</b> = <b>None</b>	string value	Default deploy interface to be used for nodes that do not have deploy_interface field set. A complete list of deploy interfaces present on your system may be found by enumerating the "ironic.hardware.interfaces.deploy" endpoint.
<b>default_inspect_interface</b> = <b>None</b>	string value	Default inspect interface to be used for nodes that do not have inspect_interface field set. A complete list of inspect interfaces present on your system may be found by enumerating the "ironic.hardware.interfaces.inspect" endpoint.

Configuration option = Default value	Type	Description
<b>default_log_levels =</b> [u'amqp=WARNING', u'amqplib=WARNING', u'qpid.messaging=INFO', u'oslo_messaging=INFO', u'oslo.messaging=INFO', u'sqlalchemy=WARNING', u'stevedore=INFO', u'eventlet.wsgi.server=IN FO', u'iso8601=WARNING', u'requests=WARNING', u'neutronclient=WARNIN G', u'glanceclient=WARNING' , u'urllib3.connectionpool= WARNING', u'keystonemiddleware.au th_token=INFO', u'keystoneauth.session=I NFO']	list value	List of package logging levels in logger=LEVEL pairs. This option is ignored if log_config_append is set.
<b>default_management_inte rface = None</b>	string value	Default management interface to be used for nodes that do not have management_interface field set. A complete list of management interfaces present on your system may be found by enumerating the "ironic.hardware.interfaces.management" endpoint.
<b>default_network_interface = None</b>	string value	Default network interface to be used for nodes that do not have network_interface field set. A complete list of network interfaces present on your system may be found by enumerating the "ironic.hardware.interfaces.network" endpoint.
<b>default_portgroup_mode = active-backup</b>	string value	Default mode for portgroups. Allowed values can be found in the linux kernel documentation on bonding: <a href="https://www.kernel.org/doc/Documentation/networking/bonding.txt">https://www.kernel.org/doc/Documentation/networking/bonding.txt</a> .
<b>default_power_interface = None</b>	string value	Default power interface to be used for nodes that do not have power_interface field set. A complete list of power interfaces present on your system may be found by enumerating the "ironic.hardware.interfaces.power" endpoint.

Configuration option = Default value	Type	Description
<b>default_raid_interface = None</b>	string value	Default raid interface to be used for nodes that do not have <code>raid_interface</code> field set. A complete list of raid interfaces present on your system may be found by enumerating the <code>"ironic.hardware.interfaces.raid"</code> endpoint.
<b>default_rescue_interface = None</b>	string value	Default rescue interface to be used for nodes that do not have <code>rescue_interface</code> field set. A complete list of rescue interfaces present on your system may be found by enumerating the <code>"ironic.hardware.interfaces.rescue"</code> endpoint.
<b>default_resource_class = None</b>	string value	Resource class to use for new nodes when no resource class is provided in the creation request.
<b>default_storage_interface = None</b>	string value	Default storage interface to be used for nodes that do not have <code>storage_interface</code> field set. A complete list of storage interfaces present on your system may be found by enumerating the <code>"ironic.hardware.interfaces.storage"</code> endpoint.
<b>default_vendor_interface = None</b>	string value	Default vendor interface to be used for nodes that do not have <code>vendor_interface</code> field set. A complete list of vendor interfaces present on your system may be found by enumerating the <code>"ironic.hardware.interfaces.vendor"</code> endpoint.
<b>enabled_boot_interfaces = [u'pxe']</b>	list value	Specify the list of boot interfaces to load during service initialization. Missing boot interfaces, or boot interfaces which fail to initialize, will prevent the <code>ironic-conductor</code> service from starting. At least one boot interface that is supported by each enabled hardware type must be enabled here, or the <code>ironic-conductor</code> service will not start. Must not be an empty list. The default value is a recommended set of production-oriented boot interfaces. A complete list of boot interfaces present on your system may be found by enumerating the <code>"ironic.hardware.interfaces.boot"</code> endpoint. When setting this value, please make sure that every enabled hardware type will have the same set of enabled boot interfaces on every <code>ironic-conductor</code> service.

Configuration option = Default value	Type	Description
<b>enabled_console_interfaces = [u'no-console']</b>	list value	Specify the list of console interfaces to load during service initialization. Missing console interfaces, or console interfaces which fail to initialize, will prevent the ironic-conductor service from starting. At least one console interface that is supported by each enabled hardware type must be enabled here, or the ironic-conductor service will not start. Must not be an empty list. The default value is a recommended set of production-oriented console interfaces. A complete list of console interfaces present on your system may be found by enumerating the "ironic.hardware.interfaces.console" endpoint. When setting this value, please make sure that every enabled hardware type will have the same set of enabled console interfaces on every ironic-conductor service.
<b>enabled_deploy_interfaces = [u'iscsi', u'direct']</b>	list value	Specify the list of deploy interfaces to load during service initialization. Missing deploy interfaces, or deploy interfaces which fail to initialize, will prevent the ironic-conductor service from starting. At least one deploy interface that is supported by each enabled hardware type must be enabled here, or the ironic-conductor service will not start. Must not be an empty list. The default value is a recommended set of production-oriented deploy interfaces. A complete list of deploy interfaces present on your system may be found by enumerating the "ironic.hardware.interfaces.deploy" endpoint. When setting this value, please make sure that every enabled hardware type will have the same set of enabled deploy interfaces on every ironic-conductor service.
<b>enabled_drivers = [u'pxe_ipmitool']</b>	list value	Specify the list of drivers to load during service initialization. Missing drivers, or drivers which fail to initialize, will prevent the conductor service from starting. The option default is a recommended set of production-oriented drivers. A complete list of drivers present on your system may be found by enumerating the "ironic.drivers" endpoint. An example may be found in the developer documentation online.

Configuration option = Default value	Type	Description
<b>enabled_hardware_types = [u'ipmi']</b>	list value	Specify the list of hardware types to load during service initialization. Missing hardware types, or hardware types which fail to initialize, will prevent the conductor service from starting. This option defaults to a recommended set of production-oriented hardware types. A complete list of hardware types present on your system may be found by enumerating the "ironic.hardware.types" endpoint.
<b>enabled_inspect_interfaces = [u'no-inspect']</b>	list value	Specify the list of inspect interfaces to load during service initialization. Missing inspect interfaces, or inspect interfaces which fail to initialize, will prevent the ironic-conductor service from starting. At least one inspect interface that is supported by each enabled hardware type must be enabled here, or the ironic-conductor service will not start. Must not be an empty list. The default value is a recommended set of production-oriented inspect interfaces. A complete list of inspect interfaces present on your system may be found by enumerating the "ironic.hardware.interfaces.inspect" endpoint. When setting this value, please make sure that every enabled hardware type will have the same set of enabled inspect interfaces on every ironic-conductor service.
<b>enabled_management_interfaces = [u'ipmitool']</b>	list value	Specify the list of management interfaces to load during service initialization. Missing management interfaces, or management interfaces which fail to initialize, will prevent the ironic-conductor service from starting. At least one management interface that is supported by each enabled hardware type must be enabled here, or the ironic-conductor service will not start. Must not be an empty list. The default value is a recommended set of production-oriented management interfaces. A complete list of management interfaces present on your system may be found by enumerating the "ironic.hardware.interfaces.management" endpoint. When setting this value, please make sure that every enabled hardware type will have the same set of enabled management interfaces on every ironic-conductor service.



Configuration option = Default value	Type	Description
<b>enabled_network_interfaces = [u'flat', u'noop']</b>	list value	Specify the list of network interfaces to load during service initialization. Missing network interfaces, or network interfaces which fail to initialize, will prevent the ironic-conductor service from starting. At least one network interface that is supported by each enabled hardware type must be enabled here, or the ironic-conductor service will not start. Must not be an empty list. The default value is a recommended set of production-oriented network interfaces. A complete list of network interfaces present on your system may be found by enumerating the "ironic.hardware.interfaces.network" endpoint. When setting this value, please make sure that every enabled hardware type will have the same set of enabled network interfaces on every ironic-conductor service.
<b>enabled_power_interfaces = [u'ipmitool']</b>	list value	Specify the list of power interfaces to load during service initialization. Missing power interfaces, or power interfaces which fail to initialize, will prevent the ironic-conductor service from starting. At least one power interface that is supported by each enabled hardware type must be enabled here, or the ironic-conductor service will not start. Must not be an empty list. The default value is a recommended set of production-oriented power interfaces. A complete list of power interfaces present on your system may be found by enumerating the "ironic.hardware.interfaces.power" endpoint. When setting this value, please make sure that every enabled hardware type will have the same set of enabled power interfaces on every ironic-conductor service.

Configuration option = Default value	Type	Description
<b>enabled_raid_interfaces = [u'agent', u'no-raid']</b>	list value	Specify the list of raid interfaces to load during service initialization. Missing raid interfaces, or raid interfaces which fail to initialize, will prevent the ironic-conductor service from starting. At least one raid interface that is supported by each enabled hardware type must be enabled here, or the ironic-conductor service will not start. Must not be an empty list. The default value is a recommended set of production-oriented raid interfaces. A complete list of raid interfaces present on your system may be found by enumerating the "ironic.hardware.interfaces.raid" endpoint. When setting this value, please make sure that every enabled hardware type will have the same set of enabled raid interfaces on every ironic-conductor service.
<b>enabled_rescue_interfaces = [u'no-rescue']</b>	list value	Specify the list of rescue interfaces to load during service initialization. Missing rescue interfaces, or rescue interfaces which fail to initialize, will prevent the ironic-conductor service from starting. At least one rescue interface that is supported by each enabled hardware type must be enabled here, or the ironic-conductor service will not start. Must not be an empty list. The default value is a recommended set of production-oriented rescue interfaces. A complete list of rescue interfaces present on your system may be found by enumerating the "ironic.hardware.interfaces.rescue" endpoint. When setting this value, please make sure that every enabled hardware type will have the same set of enabled rescue interfaces on every ironic-conductor service.

Configuration option = Default value	Type	Description
<b>enabled_storage_interfaces = [u'cinder', u'noop']</b>	list value	Specify the list of storage interfaces to load during service initialization. Missing storage interfaces, or storage interfaces which fail to initialize, will prevent the ironic-conductor service from starting. At least one storage interface that is supported by each enabled hardware type must be enabled here, or the ironic-conductor service will not start. Must not be an empty list. The default value is a recommended set of production-oriented storage interfaces. A complete list of storage interfaces present on your system may be found by enumerating the "ironic.hardware.interfaces.storage" endpoint. When setting this value, please make sure that every enabled hardware type will have the same set of enabled storage interfaces on every ironic-conductor service.
<b>enabled_vendor_interfaces = [u'ipmitool', u'no-vendor']</b>	list value	Specify the list of vendor interfaces to load during service initialization. Missing vendor interfaces, or vendor interfaces which fail to initialize, will prevent the ironic-conductor service from starting. At least one vendor interface that is supported by each enabled hardware type must be enabled here, or the ironic-conductor service will not start. Must not be an empty list. The default value is a recommended set of production-oriented vendor interfaces. A complete list of vendor interfaces present on your system may be found by enumerating the "ironic.hardware.interfaces.vendor" endpoint. When setting this value, please make sure that every enabled hardware type will have the same set of enabled vendor interfaces on every ironic-conductor service.
<b>executor_thread_pool_size = 64</b>	integer value	Size of executor thread pool when executor is threading or eventlet.
<b>fatal_deprecations = False</b>	boolean value	Enables or disables fatal status of deprecations.
<b>fatal_exception_format_errors = False</b>	boolean value	Used if there is a formatting error when generating an exception message (a programming error). If True, raise an exception; if False, use the unformatted message.
<b>force_raw_images = True</b>	boolean value	If True, convert backing images to "raw" disk image format.

Configuration option = Default value	Type	Description
<b>graceful_shutdown_timeout = 60</b>	integer value	Specify a timeout after which a gracefully shutdown server will exit. Zero value means endless wait.
<b>grub_config_template = \$pybasedir/common/grub_conf.template</b>	string value	Template file for grub configuration file.
<b>hash_distribution_replicas = 1</b>	integer value	[Experimental Feature] Number of hosts to map onto each hash partition. Setting this to more than one will cause additional conductor services to prepare deployment environments and potentially allow the Ironic cluster to recover more quickly if a conductor instance is terminated.
<b>hash_partition_exponent = 5</b>	integer value	Exponent to determine number of hash partitions to use when distributing load across conductors. Larger values will result in more even distribution of load and less load when rebalancing the ring, but more memory usage. Number of partitions per conductor is (2^hash_partition_exponent). This determines the granularity of rebalancing: given 10 hosts, and an exponent of the 2, there are 40 partitions in the ring. A few thousand partitions should make rebalancing smooth in most cases. The default is suitable for up to a few hundred conductors. Configuring for too many partitions has a negative impact on CPU usage.
<b>hash_ring_reset_interval = 15</b>	integer value	Time (in seconds) after which the hash ring is considered outdated and is refreshed on the next access.
<b>host = ccsosp-undercloud.usersys.redhat.com</b>	string value	Name of this node. This can be an opaque identifier. It is not necessarily a hostname, FQDN, or IP address. However, the node name must be valid within an AMQP key, and if using ZeroMQ, a valid hostname, FQDN, or IP address.
<code>`instance_format = [instance: %(uuid)s] `</code>	string value	The format for an instance that is passed with the log message.
<code>`instance_uuid_format = [instance: %(uuid)s] `</code>	string value	The format for an instance UUID that is passed with the log message.
<b>isolinux_bin = /usr/lib/syslinux/isolinux.bin</b>	string value	Path to isolinux binary file.

Configuration option = Default value	Type	Description
<b>isolinux_config_template</b> = \$pybasedir/common/isolinux_config.template	string value	Template file for isolinux configuration file.
<b>ldlinux_c32</b> = None	string value	Path to ldlinux.c32 file. This file is required for syslinux 5.0 or later. If not specified, the file is looked for in "/usr/lib/syslinux/modules/bios/ldlinux.c32" and "/usr/share/syslinux/ldlinux.c32".
<b>log-config-append</b> = None	string value	The name of a logging configuration file. This file is appended to any existing logging configuration files. For details about logging configuration files, see the Python logging module documentation. Note that when logging configuration files are used then all logging configuration is set in the configuration file and other logging configuration options are ignored (for example, logging_context_format_string).
<b>log-date-format</b> = %Y-%m-%d %H:%M:%S	string value	Defines the format string for %(asctime)s in log records. Default: %(default)s . This option is ignored if log_config_append is set.
<b>log-dir</b> = None	string value	(Optional) The base directory used for relative log_file paths. This option is ignored if log_config_append is set.
<b>log-file</b> = None	string value	(Optional) Name of log file to send logging output to. If no default is set, logging will go to stderr as defined by use_stderr. This option is ignored if log_config_append is set.
<b>log_options</b> = True	boolean value	Enables or disables logging values of all registered options when starting a service (at DEBUG level).
<b>logging_context_format_string</b> = %(asctime)s.% (msecs)03d %(process)d %(levelname)s %(name)s [% (request_id)s % (user_identity)s] % (instance)s%(message)s	string value	Format string to use for log messages with context.
<b>logging_debug_format_s</b> <b>uffix</b> = %(funcName)s % <b>(pathname)s:%(lineno)d</b>	string value	Additional data to append to log message when logging level for the message is DEBUG.

Configuration option = Default value	Type	Description
<b>logging_default_format_string</b> = <code>%(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [-] %(instance)s%(message)s</code>	string value	Format string to use for log messages when context is undefined.
<b>logging_exception_prefix</b> = <code>%(asctime)s.%(msecs)03d %(process)d ERROR %(name)s %(instance)s</code>	string value	Prefix each line of exception output with this format.
<b>logging_user_identity_format</b> = <code>%(user)s %(tenant)s %(domain)s %(user_domain)s %(project_domain)s</code>	string value	Defines the format string for <code>%(user_identity)s</code> that is used in <code>logging_context_format_string</code> .
<b>my_ip</b> = <code>10.65.176.99</code>	string value	IP address of this host. If unset, will determine the IP programmatically. If unable to do so, will use "127.0.0.1".
<b>notification_level</b> = <code>None</code>	string value	Specifies the minimum level for which to send notifications. If not set, no notifications will be sent. The default is for this option to be unset.
<b>parallel_image_downloads</b> = <code>False</code>	boolean value	Run image downloads and raw format conversions in parallel.
<b>pecan_debug</b> = <code>False</code>	boolean value	Enable pecan debug mode. WARNING: this is insecure and should not be used in a production environment.
<b>pin_release_version</b> = <code>None</code>	string value	Used for rolling upgrades. Setting this option downgrades (or pins) the Bare Metal API, the internal ironic RPC communication, and the database objects to their respective versions, so they are compatible with older services. When doing a rolling upgrade from version N to version N+1, set (to pin) this to N. To unpin (default), leave it unset and the latest versions will be used.
<b>publish_errors</b> = <code>False</code>	boolean value	Enables or disables publication of error events.
<b>pybasedir</b> = <code>/usr/lib/python2.7/site-packages/ironic</code>	string value	Directory where the ironic python module is installed.

Configuration option = Default value	Type	Description
<b>rate_limit_burst = 0</b>	integer value	Maximum number of logged messages per rate_limit_interval.
<b>rate_limit_except_level = CRITICAL</b>	string value	Log level name used by rate limiting: CRITICAL, ERROR, INFO, WARNING, DEBUG or empty string. Logs with level greater or equal to rate_limit_except_level are not filtered. An empty string means that all levels are filtered.
<b>rate_limit_interval = 0</b>	integer value	Interval, number of seconds, of log rate limiting.
<b>rootwrap_config = /etc/ironic/rootwrap.conf</b>	string value	Path to the rootwrap configuration file to use for running commands as root.
<b>rpc_ack_timeout_base = 15</b>	integer value	Number of seconds to wait for an ack from a cast/call. After each retry attempt this timeout is multiplied by some specified multiplier.
<b>rpc_ack_timeout_multiplier = 2</b>	integer value	Number to multiply base ack timeout by after each retry attempt.
<b>rpc_backend = rabbit</b>	string value	The messaging driver to use, defaults to rabbit. Other drivers include amqp and zmq.
<b>rpc_conn_pool_size = 30</b>	integer value	Size of RPC connection pool.
<b>rpc_message_ttl = 300</b>	integer value	Expiration timeout in seconds of a sent/received message after which it is not tracked anymore by a client/server.
<b>rpc_poll_timeout = 1</b>	integer value	The default number of seconds that poll should wait. Poll raises timeout exception when timeout expired.
<b>rpc_response_timeout = 60</b>	integer value	Seconds to wait for a response from a call.
<b>rpc_retry_attempts = 3</b>	integer value	Default number of message sending attempts in case of any problems occurred: positive value N means at most N retries, 0 means no retries, None or -1 (or any other negative values) mean to retry forever. This option is used only if acknowledgments are enabled.

Configuration option = Default value	Type	Description
<b>rpc_thread_pool_size = 100</b>	integer value	Maximum number of (green) threads to work concurrently.
<b>rpc_use_acks = False</b>	boolean value	Wait for message acknowledgements from receivers. This mechanism works only via proxy without PUB/SUB.
<b>rpc_zmq_bind_address = *</b>	string value	ZeroMQ bind address. Should be a wildcard (*), an ethernet interface, or IP. The "host" option should point or resolve to this address.
<b>rpc_zmq_bind_port_retries = 100</b>	integer value	Number of retries to find free port number before fail with ZMQBindError.
<b>rpc_zmq_contexts = 1</b>	integer value	Number of ZeroMQ contexts, defaults to 1.
<b>rpc_zmq_host = ccosp-undercloud.usersys.redhat.com</b>	string value	Name of this node. Must be a valid hostname, FQDN, or IP address. Must match "host" option, if running Nova.
<b>rpc_zmq_ipc_dir = /var/run/openstack</b>	string value	Directory for holding IPC sockets.
<b>rpc_zmq_matchmaker = redis</b>	string value	MatchMaker driver.
<b>rpc_zmq_max_port = 65536</b>	integer value	Maximal port number for random ports range.
<b>rpc_zmq_min_port = 49153</b>	port value	Minimal port number for random ports range.
<b>rpc_zmq_serialization = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>rpc_zmq_topic_backlog = None</b>	integer value	Maximum number of ingress messages to locally buffer per topic. Default is unlimited.
<b>run_external_periodic_tasks = True</b>	boolean value	Some periodic tasks can be run in a separate process. Should we run them here?
<b>state_path = \$pybasedir</b>	string value	Top-level directory for maintaining ironic's state.
<b>subscribe_on = []</b>	list value	List of publisher hosts SubConsumer can subscribe on. This option has higher priority then the default publishers list taken from the matchmaker.



Configuration option = Default value	Type	Description
<b>syslog-log-facility = LOG_USER</b>	string value	Syslog facility to receive log lines. This option is ignored if log_config_append is set.
<b>tempdir = /tmp</b>	string value	Temporary working directory, default is Python temp dir.
<b>transport_url = None</b>	string value	<p>The network address and optional user credentials for connecting to the messaging backend, in URL format. The expected format is:</p> <pre>driver://[user:pass@]host:port[, [userN:passN@]hostN:portN]/virtual_host?query</pre> <p>Example: rabbit://rabbitmq:password@127.0.0.1:5672//</p> <p>For full details on the fields in the URL see the documentation of oslo_messaging.TransportURL at <a href="https://docs.openstack.org/oslo.messaging/latest/reference/transport.html">https://docs.openstack.org/oslo.messaging/latest/reference/transport.html</a></p>
<b>use-journal = False</b>	boolean value	Enable journald for logging. If running in a systemd environment you may wish to enable journal support. Doing so will use the journal native protocol which includes structured metadata in addition to log messages. This option is ignored if log_config_append is set.
<b>use-json = False</b>	boolean value	Use JSON formatting for logging. This option is ignored if log_config_append is set.
<b>use-syslog = False</b>	boolean value	Use syslog for logging. Existing syslog format is DEPRECATED and will be changed later to honor RFC5424. This option is ignored if log_config_append is set.
<b>use_dynamic_connections = False</b>	boolean value	This option makes direct connections dynamic or static. It makes sense only with use_router_proxy=False which means to use direct connections for direct message types (ignored otherwise).
<b>use_pub_sub = False</b>	boolean value	Use PUB/SUB pattern for fanout methods. PUB/SUB always uses proxy.

Configuration option = Default value	Type	Description
<b>use_router_proxy = False</b>	boolean value	Use ROUTER remote proxy.
<b>use_stderr = False</b>	boolean value	Log output to standard error. This option is ignored if log_config_append is set.
<b>watch-log-file = False</b>	boolean value	Uses logging handler designed to watch file system. When log file is moved or removed this handler will open a new log file with specified path instantaneously. It makes sense only if log_file option is specified and Linux platform is used. This option is ignored if log_config_append is set.
<b>zmq_failover_connections = 2</b>	integer value	How many additional connections to a host will be made for failover reasons. This option is actual only in dynamic connections mode.
<b>zmq_immediate = True</b>	boolean value	This option configures round-robin mode in zmq socket. True means not keeping a queue when server side disconnects. False means to keep queue and messages even if server is disconnected, when the server appears we send all accumulated messages to it.
<b>zmq_linger = -1</b>	integer value	Number of seconds to wait before all pending messages will be sent after closing a socket. The default value of -1 specifies an infinite linger period. The value of 0 specifies no linger period. Pending messages shall be discarded immediately when the socket is closed. Positive values specify an upper bound for the linger period.
<b>zmq_target_expire = 300</b>	integer value	Expiration timeout in seconds of a name service record about existing target (< 0 means no timeout).
<b>zmq_target_update = 180</b>	integer value	Update period in seconds of a name service record about existing target.
<b>zmq_tcp_keepalive = -1</b>	integer value	Enable/disable TCP keepalive (KA) mechanism. The default value of -1 (or any other negative value) means to skip any overrides and leave it to OS default; 0 and 1 (or any other positive value) mean to disable and enable the option respectively.

Configuration option = Default value	Type	Description
<b>zmq_tcp_keepalive_cnt = -1</b>	integer value	The number of retransmissions to be carried out before declaring that remote end is not available. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_idle = -1</b>	integer value	The duration between two keepalive transmissions in idle condition. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_intvl = -1</b>	integer value	The duration between two successive keepalive retransmissions, if acknowledgement to the previous keepalive transmission is not received. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.

### 7.1.2. agent

The following table outlines the options available under the **[agent]** group in the `/etc/ironic/ironic.conf` file.

Table 7.2. agent

Configuration option = Default value	Type	Description
<b>agent_api_version = v1</b>	string value	API version to use for communicating with the ramdisk agent.
<b>deploy_logs_collect = on_failure</b>	string value	Whether Ironic should collect the deployment logs on deployment failure (on_failure), always or never.
<b>deploy_logs_local_path = /var/log/ironic/deploy</b>	string value	The path to the directory where the logs should be stored, used when the <code>deploy_logs_storage_backend</code> is configured to "local".
<b>deploy_logs_storage_backend = local</b>	string value	The name of the storage backend where the logs will be stored.

Configuration option = Default value	Type	Description
<b>deploy_logs_swift_container = ironic_deploy_logs_container</b>	string value	The name of the Swift container to store the logs, used when the <code>deploy_logs_storage_backend</code> is configured to "swift".
<b>deploy_logs_swift_days_to_expire = 30</b>	integer value	Number of days before a log object is marked as expired in Swift. If None, the logs will be kept forever or until manually deleted. Used when the <code>deploy_logs_storage_backend</code> is configured to "swift".
<b>manage_agent_boot = True</b>	boolean value	Whether Ironic will manage booting of the agent ramdisk. If set to False, you will need to configure your mechanism to allow booting the agent ramdisk.
<b>memory_consumed_by_agent = 0</b>	integer value	The memory size in MiB consumed by agent when it is booted on a bare metal node. This is used for checking if the image can be downloaded and deployed on the bare metal node after booting agent ramdisk. This may be set according to the memory consumed by the agent ramdisk image.
<b>post_deploy_get_power_state_retries = 6</b>	integer value	Number of times to retry getting power state to check if bare metal node has been powered off after a soft power off.
<b>post_deploy_get_power_state_retry_interval = 5</b>	integer value	Amount of time (in seconds) to wait between polling power state after trigger soft poweroff.
<b>stream_raw_images = True</b>	boolean value	Whether the agent ramdisk should stream raw images directly onto the disk or not. By streaming raw images directly onto the disk the agent ramdisk will not spend time copying the image to a tmpfs partition (therefore consuming less memory) prior to writing it to the disk. Unless the disk where the image will be copied to is really slow, this option should be set to True. Defaults to True.

### 7.1.3. ansible

The following table outlines the options available under the **[ansible]** group in the `/etc/ironic/ironic.conf` file.

Table 7.3. ansible

Configuration option = Default value	Type	Description
<b>ansible_extra_args = None</b>	string value	Extra arguments to pass on every invocation of Ansible.
<b>ansible_playbook_script = ansible-playbook</b>	string value	Path to "ansible-playbook" script. Default will search the \$PATH configured for user running ironic-conductor process. Provide the full path when ansible-playbook is not in \$PATH or installed in not default location.
<b>config_file_path = \$pybasedir/drivers/modules/ansible/playbooks/ansible.cfg</b>	string value	Path to ansible configuration file. If set to empty, system default will be used.
<b>default_clean_playbook = clean.yaml</b>	string value	Path (relative to \$playbooks_path or absolute) to the default playbook used for node cleaning. It may be overridden by per-node <i>ansible_clean_playbook</i> option in node's <i>driver_info</i> field.
<b>default_clean_steps_config = clean_steps.yaml</b>	string value	Path (relative to \$playbooks_path or absolute) to the default auxiliary cleaning steps file used during the node cleaning. It may be overridden by per-node <i>ansible_clean_steps_config</i> option in node's <i>driver_info</i> field.
<b>default_deploy_playbook = deploy.yaml</b>	string value	Path (relative to \$playbooks_path or absolute) to the default playbook used for deployment. It may be overridden by per-node <i>ansible_deploy_playbook</i> option in node's <i>driver_info</i> field.
<b>default_key_file = None</b>	string value	Absolute path to the private SSH key file to use by Ansible by default when connecting to the ramdisk over SSH. Default is to use default SSH keys configured for the user running the ironic-conductor service. Private keys with password must be pre-loaded into <i>ssh-agent</i> . It may be overridden by per-node <i>ansible_key_file</i> option in node's <i>driver_info</i> field.
<b>default_shutdown_playbook = shutdown.yaml</b>	string value	Path (relative to \$playbooks_path or absolute) to the default playbook used for graceful in-band shutdown of the node. It may be overridden by per-node <i>ansible_shutdown_playbook</i> option in node's <i>driver_info</i> field.

Configuration option = Default value	Type	Description
<b>default_username = ansible</b>	string value	Name of the user to use for Ansible when connecting to the ramdisk over SSH. It may be overridden by per-node <i>ansible_username</i> option in node's <i>driver_info</i> field.
<b>extra_memory = 10</b>	integer value	Extra amount of memory in MiB expected to be consumed by Ansible-related processes on the node. Affects decision whether image will fit into RAM.
<b>image_store_cafile = None</b>	string value	Specific CA bundle to use for validating SSL connections to the image store. If not specified, CA available in the ramdisk will be used. Is not used by default playbooks included with the driver. Suitable for environments that use self-signed certificates.
<b>image_store_certfile = None</b>	string value	Client cert to use for SSL connections to image store. Is not used by default playbooks included with the driver.
<b>image_store_insecure = False</b>	boolean value	Skip verifying SSL connections to the image store when downloading the image. Setting it to "True" is only recommended for testing environments that use self-signed certificates.
<b>image_store_keyfile = None</b>	string value	Client key to use for SSL connections to image store. Is not used by default playbooks included with the driver.
<b>playbooks_path = \$pybasedir/drivers/modules/ansible/playbooks</b>	string value	Path to directory with playbooks, roles and local inventory.
<b>post_deploy_get_power_state_retries = 6</b>	integer value	Number of times to retry getting power state to check if bare metal node has been powered off after a soft power off. Value of 0 means do not retry on failure.
<b>post_deploy_get_power_state_retry_interval = 5</b>	integer value	Amount of time (in seconds) to wait between polling power state after trigger soft poweroff.
<b>verbosity = None</b>	integer value	Set ansible verbosity level requested when invoking "ansible-playbook" command. 4 includes detailed SSH session logging. Default is 4 when global debug is enabled and 0 otherwise.

#### 7.1.4. api

The following table outlines the options available under the **[api]** group in the **/etc/ironic/ironic.conf** file.

Table 7.4. api

Configuration option = Default value	Type	Description
<b>api_workers = None</b>	integer value	Number of workers for OpenStack Ironic API service. The default is equal to the number of CPUs available if that can be determined, else a default worker count of 1 is returned.
<b>enable_ssl_api = False</b>	boolean value	Enable the integrated stand-alone API to service requests via HTTPS instead of HTTP. If there is a front-end service performing HTTPS offloading from the service, this option should be False; note, you will want to change public API endpoint to represent SSL termination URL with <i>public_endpoint</i> option.
<b>host_ip = 0.0.0.0</b>	string value	The IP address on which ironic-api listens.
<b>max_limit = 1000</b>	integer value	The maximum number of items returned in a single response from a collection resource.
<b>port = 6385</b>	port value	The TCP port on which ironic-api listens.
<b>public_endpoint = None</b>	string value	Public URL to use when building the links to the API resources (for example, "https://ironic.rocks:6384"). If None the links will be built using the request's host URL. If the API is operating behind a proxy, you will want to change this to represent the proxy's URL. Defaults to None.
<b>ramdisk_heartbeat_timeout = 300</b>	integer value	Maximum interval (in seconds) for agent heartbeats.
<b>restrict_lookup = True</b>	boolean value	Whether to restrict the lookup API to only nodes in certain states.

### 7.1.5. audit

The following table outlines the options available under the **[audit]** group in the `/etc/ironic/ironic.conf` file.

Table 7.5. audit

Configuration option = Default value	Type	Description
<b>audit_map_file = /etc/ironic/api_audit_map.conf</b>	string value	Path to audit map file for ironic-api service. Used only when API audit is enabled.

Configuration option = Default value	Type	Description
<b>enabled = False</b>	boolean value	Enable auditing of API requests (for ironic-api service).
<code>ignore_req_list = ``</code>	string value	Comma separated list of Ironic REST API HTTP methods to be ignored during audit logging. For example: auditing will not be done on any GET or POST requests if this is set to "GET,POST". It is used only when API audit is enabled.

### 7.1.6. cimc

The following table outlines the options available under the **[cimc]** group in the `/etc/ironic/ironic.conf` file.

Table 7.6. cimc

Configuration option = Default value	Type	Description
<b>action_interval = 10</b>	integer value	Amount of time in seconds to wait in between power operations
<b>max_retry = 6</b>	integer value	Number of times a power operation needs to be retried

### 7.1.7. cinder

The following table outlines the options available under the **[cinder]** group in the `/etc/ironic/ironic.conf` file.

Table 7.7. cinder

Configuration option = Default value	Type	Description
<b>action_retries = 3</b>	integer value	Number of retries in the case of a failed action (currently only used when detaching volumes).
<b>action_retry_interval = 5</b>	integer value	Retry interval in seconds in the case of a failed action (only specific actions are retried).
<b>auth-url = None</b>	string value	Authentication URL
<b>auth_type = None</b>	string value	Authentication type to load



Configuration option = Default value	Type	Description
<b>cafile = None</b>	string value	PEM encoded Certificate Authority to use when verifying HTTPs connections.
<b>certfile = None</b>	string value	PEM encoded client certificate cert file
<b>default-domain-id = None</b>	string value	Optional domain ID to use with v3 and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.
<b>default-domain-name = None</b>	string value	Optional domain name to use with v3 API and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.
<b>domain-id = None</b>	string value	Domain ID to scope to
<b>domain-name = None</b>	string value	Domain name to scope to
<b>endpoint-override = None</b>	string value	Always use this endpoint URL for requests for this client. NOTE: The unversioned endpoint should be specified here; to request a particular API version, use the <b>version</b> , <b>min-version</b> , and/or <b>max-version</b> options.
<b>insecure = False</b>	boolean value	Verify HTTPS connections.
<b>keyfile = None</b>	string value	PEM encoded client certificate key file
<b>max-version = None</b>	string value	The maximum major version of a given API, intended to be used as the upper bound of a range with min_version. Mutually exclusive with version.
<b>min-version = None</b>	string value	The minimum major version of a given API, intended to be used as the lower bound of a range with max_version. Mutually exclusive with version. If min_version is given with no max_version it is as if max version is "latest".
<b>password = None</b>	string value	User's password
<b>project-domain-id = None</b>	string value	Domain ID containing project
<b>project-domain-name = None</b>	string value	Domain name containing project

Configuration option = Default value	Type	Description
<b>project-id = None</b>	string value	Project ID to scope to
<b>project-name = None</b>	string value	Project name to scope to
<b>region-name = None</b>	string value	The default region_name for endpoint URL discovery.
<b>retries = 3</b>	integer value	Client retries in the case of a failed request connection.
<b>service-name = None</b>	string value	The default service_name for endpoint URL discovery.
<b>service-type = volumev3</b>	string value	The default service_type for endpoint URL discovery.
<b>system-scope = None</b>	string value	Scope for system operations
<b>tenant-id = None</b>	string value	Tenant ID
<b>tenant-name = None</b>	string value	Tenant Name
<b>timeout = None</b>	integer value	Timeout value for http requests
<b>trust-id = None</b>	string value	Trust ID
<b>url = None</b>	uri value	URL for connecting to cinder. If set, the value must start with either http:// or https://.
<b>user-domain-id = None</b>	string value	User's domain id
<b>user-domain-name = None</b>	string value	User's domain name
<b>user-id = None</b>	string value	User id
<b>username = None</b>	string value	Username
<b>valid-interfaces = [u'internal', u'public']</b>	list value	List of interfaces, in order of preference, for endpoint URL.
<b>version = None</b>	string value	Minimum Major API version within a given Major API version for endpoint URL discovery. Mutually exclusive with min_version and max_version

### 7.1.8. cisco\_ucs

The following table outlines the options available under the **[cisco\_ucs]** group in the **/etc/ironic/ironic.conf** file.

Table 7.8. **cisco\_ucs**

Configuration option = Default value	Type	Description
<b>action_interval = 5</b>	integer value	Amount of time in seconds to wait in between power operations
<b>max_retry = 6</b>	integer value	Number of times a power operation needs to be retried

### 7.1.9. conductor

The following table outlines the options available under the **[conductor]** group in the **/etc/ironic/ironic.conf** file.

Table 7.9. **conductor**

Configuration option = Default value	Type	Description
<b>api_url = None</b>	uri value	URL of Ironic API service. If not set ironic can get the current value from the keystone service catalog. If set, the value must start with either <code>http://</code> or <code>https://</code> .
<b>automated_clean = True</b>	boolean value	Enables or disables automated cleaning. Automated cleaning is a configurable set of steps, such as erasing disk drives, that are performed on the node to ensure it is in a baseline state and ready to be deployed to. This is done after instance deletion as well as during the transition from a "manageable" to "available" state. When enabled, the particular steps performed to clean a node depend on which driver that node is managed by; see the individual driver's documentation for details. NOTE: The introduction of the cleaning operation causes instance deletion to take significantly longer. In an environment where all tenants are trusted (eg, because there is only one tenant), this option could be safely disabled.
<b>check_provision_state_interval = 60</b>	integer value	Interval between checks of provision timeouts, in seconds.
<b>check_rescue_state_interval = 60</b>	integer value	Interval (seconds) between checks of rescue timeouts.

Configuration option = Default value	Type	Description
<b>clean_callback_timeout = 1800</b>	integer value	Timeout (seconds) to wait for a callback from the ramdisk doing the cleaning. If the timeout is reached the node will be put in the "clean failed" provision state. Set to 0 to disable timeout.
<b>configdrive_swift_container = ironic_configdrive_container</b>	string value	Name of the Swift container to store config drive data. Used when configdrive_use_object_store is True.
<b>deploy_callback_timeout = 1800</b>	integer value	Timeout (seconds) to wait for a callback from a deploy ramdisk. Set to 0 to disable timeout.
<b>force_power_state_during_sync = True</b>	boolean value	During sync_power_state, should the hardware power state be set to the state recorded in the database (True) or should the database be updated based on the hardware state (False).
<b>heartbeat_interval = 10</b>	integer value	Seconds between conductor heart beats.
<b>heartbeat_timeout = 60</b>	integer value	Maximum time (in seconds) since the last check-in of a conductor. A conductor is considered inactive when this time has been exceeded.
<b>inspect_timeout = 1800</b>	integer value	Timeout (seconds) for waiting for node inspection. 0 - unlimited.
<b>node_locked_retry_attempts = 3</b>	integer value	Number of attempts to grab a node lock.
<b>node_locked_retry_interval = 1</b>	integer value	Seconds to sleep between node lock attempts.
<b>periodic_max_workers = 8</b>	integer value	Maximum number of worker threads that can be started simultaneously by a periodic task. Should be less than RPC thread pool size.
<b>power_state_change_timeout = 30</b>	integer value	Number of seconds to wait for power operations to complete, i.e., so that a baremetal node is in the desired power state. If timed out, the power operation is considered a failure.
<b>power_state_sync_max_retries = 3</b>	integer value	During sync_power_state failures, limit the number of times Ironic should try syncing the hardware node power state with the node power state in DB

Configuration option = Default value	Type	Description
<b>rescue_callback_timeout = 1800</b>	integer value	Timeout (seconds) to wait for a callback from the rescue ramdisk. If the timeout is reached the node will be put in the "rescue failed" provision state. Set to 0 to disable timeout.
<b>send_sensor_data = False</b>	boolean value	Enable sending sensor data message via the notification bus
<b>send_sensor_data_interval = 600</b>	integer value	Seconds between conductor sending sensor data message to ceilometer via the notification bus.
<b>send_sensor_data_types = [u'ALL']</b>	list value	List of comma separated meter types which need to be sent to Ceilometer. The default value, "ALL", is a special value meaning send all the sensor data.
<b>send_sensor_data_wait_timeout = 300</b>	integer value	The time in seconds to wait for send sensors data periodic task to be finished before allowing periodic call to happen again. Should be less than send_sensor_data_interval value.
<b>send_sensor_data_workers = 4</b>	integer value	The maximum number of workers that can be started simultaneously for send data from sensors periodic task.
<b>soft_power_off_timeout = 600</b>	integer value	Timeout (in seconds) of soft reboot and soft power off operation. This value always has to be positive.
<b>sync_local_state_interval = 180</b>	integer value	When conductors join or leave the cluster, existing conductors may need to update any persistent local state as nodes are moved around the cluster. This option controls how often, in seconds, each conductor will check for nodes that it should "take over". Set it to a negative value to disable the check entirely.
<b>sync_power_state_interval = 60</b>	integer value	Interval between syncing the node power state to the database, in seconds.
<b>workers_pool_size = 100</b>	integer value	The size of the workers greenthread pool. Note that 2 threads will be reserved by the conductor itself for handling heart beats and periodic tasks.

### 7.1.10. console

The following table outlines the options available under the **[console]** group in the **/etc/ironic/ironic.conf** file.

Table 7.10. console

Configuration option = Default value	Type	Description
<b>socat_address = \$my_ip</b>	IP address value	IP address of Socat service running on the host of ironic conductor. Used only by Socat console.
<b>subprocess_checking_interval = 1</b>	integer value	Time interval (in seconds) for checking the status of console subprocess.
<b>subprocess_timeout = 10</b>	integer value	Time (in seconds) to wait for the console subprocess to start.
<b>terminal = shellinaboxd</b>	string value	Path to serial console terminal program. Used only by Shell In A Box console.
<b>terminal_cert_dir = None</b>	string value	Directory containing the terminal SSL cert (PEM) for serial console access. Used only by Shell In A Box console.
<b>terminal_pid_dir = None</b>	string value	Directory for holding terminal pid files. If not specified, the temporary directory will be used.
<b>terminal_timeout = 600</b>	integer value	Timeout (in seconds) for the terminal session to be closed on inactivity. Set to 0 to disable timeout. Used only by Socat console.

### 7.1.11. cors

The following table outlines the options available under the **[cors]** group in the `/etc/ironic/ironic.conf` file.

Table 7.11. cors

Configuration option = Default value	Type	Description
<b>allow_credentials = True</b>	boolean value	Indicate that the actual request can include user credentials
<b>allow_headers = []</b>	list value	Indicate which header field names may be used during the actual request.
<b>allow_methods = [u'OPTIONS', u'GET', u'HEAD', u'POST', u'PUT', u'DELETE', u'TRACE', u'PATCH']</b>	list value	Indicate which methods can be used during the actual request.

Configuration option = Default value	Type	Description
<b>allowed_origin = None</b>	list value	Indicate whether this resource may be shared with the domain received in the requests "origin" header. Format: "<protocol>://<host>[:<port>]", no trailing slash. Example: <a href="https://horizon.example.com">https://horizon.example.com</a>
<b>expose_headers = []</b>	list value	Indicate which headers are safe to expose to the API. Defaults to HTTP Simple Headers.
<b>max_age = 3600</b>	integer value	Maximum cache age of CORS preflight requests.

### 7.1.12. database

The following table outlines the options available under the **[database]** group in the `/etc/ironic/ironic.conf` file.

Table 7.12. database

Configuration option = Default value	Type	Description
<b>backend = sqlalchemy</b>	string value	The back end to use for the database.
<b>connection = None</b>	string value	The SQLAlchemy connection string to use to connect to the database.
<b>connection_debug = 0</b>	integer value	Verbosity of SQL debugging information: 0=None, 100=Everything.
<b>connection_recycle_time = 3600</b>	integer value	Connections which have been present in the connection pool longer than this number of seconds will be replaced with a new one the next time they are checked out from the pool.
<b>connection_trace = False</b>	boolean value	Add Python stack traces to SQL as comment strings.
<b>db_inc_retry_interval = True</b>	boolean value	If True, increases the interval between retries of a database operation up to <code>db_max_retry_interval</code> .
<b>db_max_retries = 20</b>	integer value	Maximum retries in case of connection error or deadlock error before error is raised. Set to -1 to specify an infinite retry count.
<b>db_max_retry_interval = 10</b>	integer value	If <code>db_inc_retry_interval</code> is set, the maximum seconds between retries of a database operation.

Configuration option = Default value	Type	Description
<b>db_retry_interval = 1</b>	integer value	Seconds between retries of a database transaction.
<b>max_overflow = 50</b>	integer value	If set, use this value for max_overflow with SQLAlchemy.
<b>max_pool_size = 5</b>	integer value	Maximum number of SQL connections to keep open in a pool. Setting a value of 0 indicates no limit.
<b>max_retries = 10</b>	integer value	Maximum number of database connection retries during startup. Set to -1 to specify an infinite retry count.
<b>min_pool_size = 1</b>	integer value	Minimum number of SQL connections to keep open in a pool.
<b>mysql_enable_ndb = False</b>	boolean value	If True, transparently enables support for handling MySQL Cluster (NDB).
<b>mysql_engine = InnoDB</b>	string value	MySQL engine to use.
<b>mysql_sql_mode = TRADITIONAL</b>	string value	The SQL mode to be used for MySQL sessions. This option, including the default, overrides any server-set SQL mode. To use whatever SQL mode is set by the server configuration, set this to no value. Example: mysql_sql_mode=
<b>pool_timeout = None</b>	integer value	If set, use this value for pool_timeout with SQLAlchemy.
<b>retry_interval = 10</b>	integer value	Interval between retries of opening a SQL connection.
<b>slave_connection = None</b>	string value	The SQLAlchemy connection string to use to connect to the slave database.
<b>sqlite_synchronous = True</b>	boolean value	If True, SQLite uses synchronous mode.
<b>use_db_reconnect = False</b>	boolean value	Enable the experimental use of database reconnect on connection lost.

### 7.1.13. deploy

The following table outlines the options available under the **[deploy]** group in the `/etc/ironic/ironic.conf` file.

Table 7.13. deploy



Configuration option = Default value	Type	Description
<b>configdrive_use_object_store = False</b>	boolean value	Whether to upload the config drive to object store. Set this option to True to store config drive in swift or radosgw.
<b>continue_if_disk_secure_erase_fails = False</b>	boolean value	Defines what to do if an ATA secure erase operation fails during cleaning in the Ironic Python Agent. If False, the cleaning operation will fail and the node will be put in <b>clean failed</b> state. If True, shred will be invoked and cleaning will continue.
<b>default_boot_option = None</b>	string value	Default boot option to use when no boot option is requested in node's driver_info. Currently the default is "netboot", but it will be changed to "local" in the future. It is recommended to set an explicit value for this option.
<b>erase_devices_metadata_priority = None</b>	integer value	Priority to run in-band clean step that erases metadata from devices, via the Ironic Python Agent ramdisk. If unset, will use the priority set in the ramdisk (defaults to 99 for the GenericHardwareManager). If set to 0, will not run during cleaning.
<b>erase_devices_priority = None</b>	integer value	Priority to run in-band erase devices via the Ironic Python Agent ramdisk. If unset, will use the priority set in the ramdisk (defaults to 10 for the GenericHardwareManager). If set to 0, will not run during cleaning.
<b>http_root = /httpboot</b>	string value	ironic-conductor node's HTTP root path.
<b>http_url = None</b>	string value	ironic-conductor node's HTTP server URL. Example: <a href="http://192.1.2.3:8080">http://192.1.2.3:8080</a>
<b>object_store_endpoint_type = swift</b>	string value	Type of object store endpoint type to be used as a backend
<b>power_off_after_deploy_failure = True</b>	boolean value	Whether to power off a node after deploy failure. Defaults to True.
<b>shred_final_overwrite_with_zeros = True</b>	boolean value	Whether to write zeros to a node's block devices after writing random data. This will write zeros to the device even when <code>deploy.shred_random_overwrite_iterations</code> is 0. This option is only used if a device could not be ATA Secure Erased. Defaults to True.

Configuration option = Default value	Type	Description
<b>shred_random_overwrite_iterations = 1</b>	integer value	During shred, overwrite all block devices N times with random data. This is only used if a device could not be ATA Secure Erased. Defaults to 1.

#### 7.1.14. dhcp

The following table outlines the options available under the **[dhcp]** group in the `/etc/ironic/ironic.conf` file.

Table 7.14. dhcp

Configuration option = Default value	Type	Description
<b>dhcp_provider = neutron</b>	string value	DHCP provider to use. "neutron" uses Neutron, and "none" uses a no-op provider.

#### 7.1.15. disk\_partitioner

The following table outlines the options available under the **[disk\_partitioner]** group in the `/etc/ironic/ironic.conf` file.

Table 7.15. disk\_partitioner

Configuration option = Default value	Type	Description
<b>check_device_interval = 1</b>	integer value	After Ironic has completed creating the partition table, it continues to check for activity on the attached iSCSI device status at this interval prior to copying the image to the node, in seconds
<b>check_device_max_retries = 20</b>	integer value	The maximum number of times to check that the device is not accessed by another process. If the device is still busy after that, the disk partitioning will be treated as having failed.

#### 7.1.16. disk\_utils

The following table outlines the options available under the **[disk\_utils]** group in the `/etc/ironic/ironic.conf` file.

Table 7.16. disk\_utils

Configuration option = Default value	Type	Description
<b>bios_boot_partition_size = 1</b>	integer value	Size of BIOS Boot partition in MiB when configuring GPT partitioned systems for local boot in BIOS.
<b>dd_block_size = 1M</b>	string value	Block size to use when writing to the nodes disk.
<b>efi_system_partition_size = 200</b>	integer value	Size of EFI system partition in MiB when configuring UEFI systems for local boot.
<b>iscsi_verify_attempts = 3</b>	integer value	Maximum attempts to verify an iSCSI connection is active, sleeping 1 second between attempts.
<b>partprobe_attempts = 10</b>	integer value	Maximum number of attempts to try to read the partition.

### 7.1.17. drac

The following table outlines the options available under the **[drac]** group in the `/etc/ironic/ironic.conf` file.

Table 7.17. drac

Configuration option = Default value	Type	Description
<b>query_raid_config_job_status_interval = 120</b>	integer value	Interval (in seconds) between periodic RAID job status checks to determine whether the asynchronous RAID configuration was successfully finished or not.

### 7.1.18. glance

The following table outlines the options available under the **[glance]** group in the `/etc/ironic/ironic.conf` file.

Table 7.18. glance

Configuration option = Default value	Type	Description
<b>allowed_direct_url_schemes = []</b>	list value	A list of URL schemes that can be downloaded directly via the <code>direct_url</code> . Currently supported schemes: <code>[file]</code> .
<b>auth-url = None</b>	string value	Authentication URL

Configuration option = Default value	Type	Description
<b>auth_strategy = keystone</b>	string value	Authentication strategy to use when connecting to glance.
<b>auth_type = None</b>	string value	Authentication type to load
<b>cafile = None</b>	string value	PEM encoded Certificate Authority to use when verifying HTTPs connections.
<b>certfile = None</b>	string value	PEM encoded client certificate cert file
<b>default-domain-id = None</b>	string value	Optional domain ID to use with v3 and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.
<b>default-domain-name = None</b>	string value	Optional domain name to use with v3 API and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.
<b>domain-id = None</b>	string value	Domain ID to scope to
<b>domain-name = None</b>	string value	Domain name to scope to
<b>endpoint-override = None</b>	string value	Always use this endpoint URL for requests for this client. NOTE: The unversioned endpoint should be specified here; to request a particular API version, use the <b>version</b> , <b>min-version</b> , and/or <b>max-version</b> options.
<b>glance_api_insecure = False</b>	boolean value	Allow to perform insecure SSL (https) requests to glance.
<b>glance_api_servers = None</b>	list value	A list of the glance api servers available to ironic. Prefix with https:// for SSL-based glance API servers. Format is [hostname IP]:port.
<b>glance_api_version = 2</b>	integer value	Glance API version (1 or 2) to use.
<b>glance_cafile = None</b>	string value	Optional path to a CA certificate bundle to be used to validate the SSL certificate served by glance. It is used when glance_api_insecure is set to False.
<b>glance_num_retries = 0</b>	integer value	Number of retries when downloading an image from glance.
<b>insecure = False</b>	boolean value	Verify HTTPS connections.

Configuration option = Default value	Type	Description
<b>keyfile = None</b>	string value	PEM encoded client certificate key file
<b>max-version = None</b>	string value	The maximum major version of a given API, intended to be used as the upper bound of a range with <code>min_version</code> . Mutually exclusive with <code>version</code> .
<b>min-version = None</b>	string value	The minimum major version of a given API, intended to be used as the lower bound of a range with <code>max_version</code> . Mutually exclusive with <code>version</code> . If <code>min_version</code> is given with no <code>max_version</code> it is as if <code>max_version</code> is "latest".
<b>password = None</b>	string value	User's password
<b>project-domain-id = None</b>	string value	Domain ID containing project
<b>project-domain-name = None</b>	string value	Domain name containing project
<b>project-id = None</b>	string value	Project ID to scope to
<b>project-name = None</b>	string value	Project name to scope to
<b>region-name = None</b>	string value	The default <code>region_name</code> for endpoint URL discovery.
<b>service-name = None</b>	string value	The default <code>service_name</code> for endpoint URL discovery.
<b>service-type = image</b>	string value	The default <code>service_type</code> for endpoint URL discovery.
<b>swift_account = None</b>	string value	The account that Glance uses to communicate with Swift. The format is "AUTH_uuid". "uuid" is the UUID for the account configured in the <code>glance-api.conf</code> . For example: "AUTH_a422b2-91f3-2f46-74b7-d7c9e8958f5d30". If not set, the default value is calculated based on the ID of the project used to access Swift (as set in the <code>[swift]</code> section). Swift temporary URL format: "endpoint_url/api_version/[account/]container/object_id"
<b>swift_api_version = v1</b>	string value	The Swift API version to create a temporary URL for. Defaults to "v1". Swift temporary URL format: "endpoint_url/api_version/[account/]container/object_id"

Configuration option = Default value	Type	Description
<b>swift_container = glance</b>	string value	The Swift container Glance is configured to store its images in. Defaults to "glance", which is the default in glance-api.conf. Swift temporary URL format: "endpoint_url/api_version/[account/]container/object_id"
<b>swift_endpoint_url = None</b>	string value	The "endpoint" (scheme, hostname, optional port) for the Swift URL of the form "endpoint_url/api_version/[account/]container/object_id". Do not include trailing "/". For example, use "https://swift.example.com". If using RADOS Gateway, endpoint may also contain /swift path; if it does not, it will be appended. Used for temporary URLs, will be fetched from the service catalog, if not provided.
<b>swift_store_multiple_containers_seed = 0</b>	integer value	This should match a config by the same name in the Glance configuration file. When set to 0, a single-tenant store will only use one container to store all images. When set to an integer value between 1 and 32, a single-tenant store will use multiple containers to store images, and this value will determine how many containers are created.
<b>swift_temp_url_cache_enabled = False</b>	boolean value	Whether to cache generated Swift temporary URLs. Setting it to true is only useful when an image caching proxy is used. Defaults to False.
<b>swift_temp_url_duration = 1200</b>	integer value	The length of time in seconds that the temporary URL will be valid for. Defaults to 20 minutes. If some deploys get a 401 response code when trying to download from the temporary URL, try raising this duration. This value must be greater than or equal to the value for swift_temp_url_expected_download_start_delay
<b>swift_temp_url_expected_download_start_delay = 0</b>	integer value	This is the delay (in seconds) from the time of the deploy request (when the Swift temporary URL is generated) to when the IPA ramdisk starts up and URL is used for the image download. This value is used to check if the Swift temporary URL duration is large enough to let the image download begin. Also if temporary URL caching is enabled this will determine if a cached entry will still be valid when the download starts. swift_temp_url_duration value must be greater than or equal to this option's value. Defaults to 0.

Configuration option = Default value	Type	Description
<b>swift_temp_url_key = None</b>	string value	The secret token given to Swift to allow temporary URL downloads. Required for temporary URLs. For the Swift backend, the key on the service project (as set in the [swift] section) is used by default.
<b>system-scope = None</b>	string value	Scope for system operations
<b>tenant-id = None</b>	string value	Tenant ID
<b>tenant-name = None</b>	string value	Tenant Name
<b>timeout = None</b>	integer value	Timeout value for http requests
<b>trust-id = None</b>	string value	Trust ID
<b>user-domain-id = None</b>	string value	User's domain id
<b>user-domain-name = None</b>	string value	User's domain name
<b>user-id = None</b>	string value	User id
<b>username = None</b>	string value	Username
<b>valid-interfaces = [u'internal', u'public']</b>	list value	List of interfaces, in order of preference, for endpoint URL.
<b>version = None</b>	string value	Minimum Major API version within a given Major API version for endpoint URL discovery. Mutually exclusive with min_version and max_version

### 7.1.19. ilo

The following table outlines the options available under the **[ilo]** group in the **/etc/ironic/ironic.conf** file.

Table 7.19. ilo

Configuration option = Default value	Type	Description
<b>ca_file = None</b>	string value	CA certificate file to validate iLO.
<b>clean_priority_clear_secure_boot_keys = 0</b>	integer value	Priority for clear_secure_boot_keys clean step. This step is not enabled by default. It can be enabled to clear all secure boot keys enrolled with iLO.

Configuration option = Default value	Type	Description
<b>clean_priority_reset_bios_to_default = 10</b>	integer value	Priority for reset_bios_to_default clean step.
<b>clean_priority_reset_ilo = 0</b>	integer value	Priority for reset_ilo clean step.
<b>clean_priority_reset_ilo_credential = 30</b>	integer value	Priority for reset_ilo_credential clean step. This step requires "ilo_change_password" parameter to be updated in nodes's driver_info with the new password.
<b>clean_priority_reset_secure_boot_keys_to_default = 20</b>	integer value	Priority for reset_secure_boot_keys clean step. This step will reset the secure boot keys to manufacturing defaults.
<b>client_port = 443</b>	port value	Port to be used for iLO operations
<b>client_timeout = 60</b>	integer value	Timeout (in seconds) for iLO operations
<b>default_boot_mode = auto</b>	string value	Default boot mode to be used in provisioning when "boot_mode" capability is not provided in the "properties/capabilities" of the node. The default is "auto" for backward compatibility. When "auto" is specified, default boot mode will be selected based on boot mode settings on the system.
<b>power_retry = 6</b>	integer value	Number of times a power operation needs to be retried
<b>power_wait = 2</b>	integer value	Amount of time in seconds to wait in between power operations
<b>swift_ilo_container = ironic_ilo_container</b>	string value	The Swift iLO container to store data.
<b>swift_object_expiry_timeout = 900</b>	integer value	Amount of time in seconds for Swift objects to auto-expire.
<b>use_web_server_for_images = False</b>	boolean value	Set this to True to use http web server to host floppy images and generated boot ISO. This requires http_root and http_url to be configured in the [deploy] section of the config file. If this is set to False, then Ironic will use Swift to host the floppy images and generated boot_iso.

### 7.1.20. inspector



The following table outlines the options available under the **[inspector]** group in the `/etc/ironic/ironic.conf` file.

Table 7.20. inspector

Configuration option = Default value	Type	Description
<b>auth-url = None</b>	string value	Authentication URL
<b>auth_type = None</b>	string value	Authentication type to load
<b>cafile = None</b>	string value	PEM encoded Certificate Authority to use when verifying HTTPs connections.
<b>certfile = None</b>	string value	PEM encoded client certificate cert file
<b>default-domain-id = None</b>	string value	Optional domain ID to use with v3 and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.
<b>default-domain-name = None</b>	string value	Optional domain name to use with v3 API and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.
<b>domain-id = None</b>	string value	Domain ID to scope to
<b>domain-name = None</b>	string value	Domain name to scope to
<b>enabled = False</b>	boolean value	whether to enable inspection using ironic-inspector. This option does not affect new-style dynamic drivers and the fake_inspector driver.
<b>endpoint-override = None</b>	string value	Always use this endpoint URL for requests for this client. NOTE: The unversioned endpoint should be specified here; to request a particular API version, use the <b>version</b> , <b>min-version</b> , and/or <b>max-version</b> options.
<b>insecure = False</b>	boolean value	Verify HTTPS connections.
<b>keyfile = None</b>	string value	PEM encoded client certificate key file
<b>max-version = None</b>	string value	The maximum major version of a given API, intended to be used as the upper bound of a range with min_version. Mutually exclusive with version.

Configuration option = Default value	Type	Description
<b>min-version = None</b>	string value	The minimum major version of a given API, intended to be used as the lower bound of a range with <code>max_version</code> . Mutually exclusive with <code>version</code> . If <code>min_version</code> is given with no <code>max_version</code> it is as if <code>max_version</code> is "latest".
<b>password = None</b>	string value	User's password
<b>project-domain-id = None</b>	string value	Domain ID containing project
<b>project-domain-name = None</b>	string value	Domain name containing project
<b>project-id = None</b>	string value	Project ID to scope to
<b>project-name = None</b>	string value	Project name to scope to
<b>region-name = None</b>	string value	The default <code>region_name</code> for endpoint URL discovery.
<b>service-name = None</b>	string value	The default <code>service_name</code> for endpoint URL discovery.
<b>service-type = baremetal-introspection</b>	string value	The default <code>service_type</code> for endpoint URL discovery.
<b>service_url = None</b>	string value	ironic-inspector HTTP endpoint. If this is not set, the service catalog will be used.
<b>status_check_period = 60</b>	integer value	period (in seconds) to check status of nodes on inspection
<b>system-scope = None</b>	string value	Scope for system operations
<b>tenant-id = None</b>	string value	Tenant ID
<b>tenant-name = None</b>	string value	Tenant Name
<b>timeout = None</b>	integer value	Timeout value for http requests
<b>trust-id = None</b>	string value	Trust ID
<b>user-domain-id = None</b>	string value	User's domain id

Configuration option = Default value	Type	Description
<b>user-domain-name = None</b>	string value	User's domain name
<b>user-id = None</b>	string value	User id
<b>username = None</b>	string value	Username
<b>valid-interfaces = [u'internal', u'public']</b>	list value	List of interfaces, in order of preference, for endpoint URL.
<b>version = None</b>	string value	Minimum Major API version within a given Major API version for endpoint URL discovery. Mutually exclusive with min_version and max_version

### 7.1.21. ipmi

The following table outlines the options available under the **[ipmi]** group in the `/etc/ironic/ironic.conf` file.

Table 7.21. ipmi

Configuration option = Default value	Type	Description
<b>command_retry_timeout = 60</b>	integer value	Maximum time in seconds to retry retryable IPMI operations. (An operation is retryable, for example, if the requested operation fails because the BMC is busy.) Setting this too high can cause the sync power state periodic task to hang when there are slow or unresponsive BMCs.
<b>min_command_interval = 5</b>	integer value	Minimum time, in seconds, between IPMI operations sent to a server. There is a risk with some hardware that setting this too low may cause the BMC to crash. Recommended setting is 5 seconds.
<b>retry_timeout = None</b>	integer value	Maximum time in seconds to retry IPMI operations. Setting this too high can cause the sync power state periodic task to hang when there are slow or unresponsive BMCs.

### 7.1.22. irmc

The following table outlines the options available under the **[irmc]** group in the `/etc/ironic/ironic.conf` file.

Table 7.22. irmc

Configuration option = Default value	Type	Description
<b>auth_method = basic</b>	string value	Authentication method to be used for iRMC operations
<b>clean_priority_restore_irmc_bios_config = 0</b>	integer value	Priority for restore_irmc_bios_config clean step.
<b>client_timeout = 60</b>	integer value	Timeout (in seconds) for iRMC operations
<b>gpu_ids = []</b>	list value	List of vendor IDs and device IDs for GPU device to inspect. List items are in format vendorID/deviceID and separated by commas. GPU inspection will use this value to count the number of GPU device in a node. If this option is not defined, then leave out pci_gpu_devices in capabilities property. Sample gpu_ids value: 0x1000/0x0079,0x2100/0x0080
<b>port = 443</b>	port value	Port to be used for iRMC operations
<b>remote_image_server = None</b>	string value	IP of remote image server
<b>remote_image_share_name = share</b>	string value	share name of remote_image_server
<b>remote_image_share_root = /remote_image_share_root</b>	string value	Ironic conductor node's "NFS" or "CIFS" root path
<b>remote_image_share_type = CIFS</b>	string value	Share type of virtual media
<b>remote_image_user_domain = `</b>	string value	Domain name of remote_image_user_name
<b>remote_image_user_name = None</b>	string value	User name of remote_image_server
<b>remote_image_user_password = None</b>	string value	Password of remote_image_user_name
<b>sensor_method = ipmitool</b>	string value	Sensor data retrieval method.
<b>snmp_community = public</b>	string value	SNMP community. Required for versions "v1" and "v2c"

Configuration option = Default value	Type	Description
<b>snmp_polling_interval = 10</b>	integer value	SNMP polling interval in seconds
<b>snmp_port = 161</b>	port value	SNMP port
<b>snmp_security = None</b>	string value	SNMP security name. Required for version "v3"
<b>snmp_version = v2c</b>	string value	SNMP protocol version

### 7.1.23. ironic\_lib

The following table outlines the options available under the **[ironic\_lib]** group in the `/etc/ironic/ironic.conf` file.

Table 7.23. ironic\_lib

Configuration option = Default value	Type	Description
<b>root_helper = sudo ironic-rootwrap</b> <code>/etc/ironic/rootwrap.conf</code>	string value	Command that is prefixed to commands that are run as root. If not specified, no commands are run as root.

### 7.1.24. iscsi

The following table outlines the options available under the **[iscsi]** group in the `/etc/ironic/ironic.conf` file.

Table 7.24. iscsi

Configuration option = Default value	Type	Description
<b>portal_port = 3260</b>	port value	The port number on which the iSCSI portal listens for incoming connections.

### 7.1.25. keystone

The following table outlines the options available under the **[keystone]** group in the `/etc/ironic/ironic.conf` file.

Table 7.25. keystone

Configuration option = Default value	Type	Description
<b>region_name = None</b>	string value	The region used for getting endpoints of OpenStack services.

### 7.1.26. keystone\_authtoken

The following table outlines the options available under the **[keystone\_authtoken]** group in the **/etc/ironic/ironic.conf** file.

Table 7.26. keystone\_authtoken

Configuration option = Default value	Type	Description
<b>auth_section = None</b>	string value	Config Section from which to load plugin specific options
<b>auth_type = None</b>	string value	Authentication type to load
<b>auth_uri = None</b>	string value	Complete "public" Identity API endpoint. This endpoint should not be an "admin" endpoint, as it should be accessible by all end users. Unauthenticated clients are redirected to this endpoint to authenticate. Although this endpoint should ideally be unversioned, client support in the wild varies. If you're using a versioned v2 endpoint here, then this should <b>not</b> be the same endpoint the service user utilizes for validating tokens, because normal end users may not be able to reach that endpoint. This option is deprecated in favor of <code>www_authenticate_uri</code> and will be removed in the S release.
<b>auth_version = None</b>	string value	API version of the admin Identity API endpoint.
<b>cache = None</b>	string value	Request environment key where the Swift cache object is stored. When <code>auth_token</code> middleware is deployed with a Swift cache, use this option to have the middleware share a caching backend with swift. Otherwise, use the <b>memcached_servers</b> option instead.
<b>cafile = None</b>	string value	A PEM encoded Certificate Authority to use when verifying HTTPs connections. Defaults to system CAs.
<b>certfile = None</b>	string value	Required if identity server requires client certificate

Configuration option = Default value	Type	Description
<b>check_revocations_for_cached = False</b>	boolean value	If true, the revocation list will be checked for cached tokens. This requires that PKI tokens are configured on the identity server.
<b>delay_auth_decision = False</b>	boolean value	Do not handle authorization requests within the middleware, but delegate the authorization decision to downstream WSGI components.
<b>enforce_token_bind = permissive</b>	string value	Used to control the use and type of token binding. Can be set to: "disabled" to not check token binding. "permissive" (default) to validate binding information if the bind type is of a form known to the server and ignore it if not. "strict" like "permissive" but if the bind type is unknown the token will be rejected. "required" any form of token binding is needed to be allowed. Finally the name of a binding method that must be present in tokens.
<b>hash_algorithms = [u'md5']</b>	list value	Hash algorithms to use for hashing PKI tokens. This may be a single algorithm or multiple. The algorithms are those supported by Python standard <code>hashlib.new()</code> . The hashes will be tried in the order given, so put the preferred one first for performance. The result of the first hash will be stored in the cache. This will typically be set to multiple values only while migrating from a less secure algorithm to a more secure one. Once all the old tokens are expired this option should be set to a single value for better performance.
<b>http_connect_timeout = None</b>	integer value	Request timeout value for communicating with Identity API server.
<b>http_request_max_retries = 3</b>	integer value	How many times are we trying to reconnect when communicating with Identity API Server.
<b>include_service_catalog = True</b>	boolean value	(Optional) Indicate whether to set the X-Service-Catalog header. If False, middleware will not ask for service catalog on token validation and will not set the X-Service-Catalog header.
<b>insecure = False</b>	boolean value	Verify HTTPS connections.
<b>keyfile = None</b>	string value	Required if identity server requires client certificate

Configuration option = Default value	Type	Description
<b>memcache_pool_conn_get_timeout = 10</b>	integer value	(Optional) Number of seconds that an operation will wait to get a memcached client connection from the pool.
<b>memcache_pool_dead_retry = 300</b>	integer value	(Optional) Number of seconds memcached server is considered dead before it is tried again.
<b>memcache_pool_maxsize = 10</b>	integer value	(Optional) Maximum total number of open connections to every memcached server.
<b>memcache_pool_socket_timeout = 3</b>	integer value	(Optional) Socket timeout in seconds for communicating with a memcached server.
<b>memcache_pool_unused_timeout = 60</b>	integer value	(Optional) Number of seconds a connection to memcached is held unused in the pool before it is closed.
<b>memcache_secret_key = None</b>	string value	(Optional, mandatory if memcache_security_strategy is defined) This string is used for key derivation.
<b>memcache_security_strategy = None</b>	string value	(Optional) If defined, indicate whether token data should be authenticated or authenticated and encrypted. If MAC, token data is authenticated (with HMAC) in the cache. If ENCRYPT, token data is encrypted and authenticated in the cache. If the value is not one of these options or empty, auth_token will raise an exception on initialization.
<b>memcache_use_advanced_pool = False</b>	boolean value	(Optional) Use the advanced (eventlet safe) memcached client pool. The advanced pool will only work under python 2.x.
<b>memcached_servers = None</b>	list value	Optionally specify a list of memcached server(s) to use for caching. If left undefined, tokens will instead be cached in-process.
<b>region_name = None</b>	string value	The region in which the identity server can be found.
<b>revocation_cache_time = 10</b>	integer value	Determines the frequency at which the list of revoked tokens is retrieved from the Identity service (in seconds). A high number of revocation events combined with a low cache duration may significantly reduce performance. Only valid for PKI tokens. This option has been deprecated in the Ocata release and will be removed in the P release.



Configuration option = Default value	Type	Description
<b>service_token_roles = [u'service']</b>	list value	A choice of roles that must be present in a service token. Service tokens are allowed to request that an expired token can be used and so this check should tightly control that only actual services should be sending this token. Roles here are applied as an ANY check so any role in this list must be present. For backwards compatibility reasons this currently only affects the allow_expired check.
<b>service_token_roles_required = False</b>	boolean value	For backwards compatibility reasons we must let valid service tokens pass that don't pass the service_token_roles check as valid. Setting this true will become the default in a future release and should be enabled if possible.
<b>signing_dir = None</b>	string value	Directory used to cache files related to PKI tokens. This option has been deprecated in the Ocata release and will be removed in the P release.
<b>token_cache_time = 300</b>	integer value	In order to prevent excessive effort spent validating tokens, the middleware caches previously-seen tokens for a configurable duration (in seconds). Set to -1 to disable caching completely.
<b>www_authenticate_uri = None</b>	string value	Complete "public" Identity API endpoint. This endpoint should not be an "admin" endpoint, as it should be accessible by all end users. Unauthenticated clients are redirected to this endpoint to authenticate. Although this endpoint should ideally be unversioned, client support in the wild varies. If you're using a versioned v2 endpoint here, then this should <b>not</b> be the same endpoint the service user utilizes for validating tokens, because normal end users may not be able to reach that endpoint.

### 7.1.27. matchmaker\_redis

The following table outlines the options available under the **[matchmaker\_redis]** group in the **/etc/ironic/ironic.conf** file.

Table 7.27. matchmaker\_redis

Configuration option = Default value	Type	Description
<b>check_timeout = 20000</b>	integer value	Time in ms to wait before the transaction is killed.

Configuration option = Default value	Type	Description
<b>host = 127.0.0.1</b>	string value	Host to locate redis.
<code>`password = `</code>	string value	Password for Redis server (optional).
<b>port = 6379</b>	port value	Use this port to connect to redis host.
<b>sentinel_group_name = oslo-messaging-zeromq</b>	string value	Redis replica set name.
<b>sentinel_hosts = []</b>	list value	List of Redis Sentinel hosts (fault tolerance mode), e.g., [host:port, host!:port ... ]
<b>socket_timeout = 10000</b>	integer value	Timeout in ms on blocking socket operations.
<b>wait_timeout = 2000</b>	integer value	Time in ms to wait between connection attempts.

### 7.1.28. metrics

The following table outlines the options available under the **[metrics]** group in the `/etc/ironic/ironic.conf` file.

Table 7.28. metrics

Configuration option = Default value	Type	Description
<b>agent_backend = noop</b>	string value	Backend for the agent ramdisk to use for metrics. Default possible backends are "noop" and "statsd".
<b>agent_global_prefix = None</b>	string value	Prefix all metric names sent by the agent ramdisk with this value. The format of metric names is [global_prefix.][uuid.] [host_name.]prefix.metric_name.
<b>agent_prepend_host = False</b>	boolean value	Prepend the hostname to all metric names sent by the agent ramdisk. The format of metric names is [global_prefix.][uuid.] [host_name.]prefix.metric_name.
<b>agent_prepend_host_reverse = True</b>	boolean value	Split the prepended host value by "." and reverse it for metrics sent by the agent ramdisk (to better match the reverse hierarchical form of domain names).

Configuration option = Default value	Type	Description
<b>agent_prepend_uuid = False</b>	boolean value	Prepend the node's Ironic uuid to all metric names sent by the agent ramdisk. The format of metric names is [global_prefix.][uuid.] [host_name.]prefix.metric_name.
<b>backend = noop</b>	string value	Backend to use for the metrics system.
<b>global_prefix = None</b>	string value	Prefix all metric names with this value. By default, there is no global prefix. The format of metric names is [global_prefix.][host_name.]prefix.metric_name.
<b>prepend_host = False</b>	boolean value	Prepend the hostname to all metric names. The format of metric names is [global_prefix.] [host_name.]prefix.metric_name.
<b>prepend_host_reverse = True</b>	boolean value	Split the prepended host value by "." and reverse it (to better match the reverse hierarchical form of domain names).

### 7.1.29. metrics\_statsd

The following table outlines the options available under the **[metrics\_statsd]** group in the `/etc/ironic/ironic.conf` file.

Table 7.29. metrics\_statsd

Configuration option = Default value	Type	Description
<b>agent_statsd_host = localhost</b>	string value	Host for the agent ramdisk to use with the statsd backend. This must be accessible from networks the agent is booted on.
<b>agent_statsd_port = 8125</b>	port value	Port for the agent ramdisk to use with the statsd backend.
<b>statsd_host = localhost</b>	string value	Host for use with the statsd backend.
<b>statsd_port = 8125</b>	port value	Port to use with the statsd backend.

### 7.1.30. neutron

The following table outlines the options available under the **[neutron]** group in the `/etc/ironic/ironic.conf` file.

Table 7.30. neutron

Configuration option = Default value	Type	Description
<b>auth-url = None</b>	string value	Authentication URL
<b>auth_strategy = keystone</b>	string value	Authentication strategy to use when connecting to neutron. Running neutron in noauth mode (related to but not affected by this setting) is insecure and should only be used for testing.
<b>auth_type = None</b>	string value	Authentication type to load
<b>cafile = None</b>	string value	PEM encoded Certificate Authority to use when verifying HTTPs connections.
<b>certfile = None</b>	string value	PEM encoded client certificate cert file
<b>cleaning_network = None</b>	string value	Neutron network UUID or name for the ramdisk to be booted into for cleaning nodes. Required for "neutron" network interface. It is also required if cleaning nodes when using "flat" network interface or "neutron" DHCP provider. If a name is provided, it must be unique among all networks or cleaning will fail.
<b>cleaning_network_security_groups = []</b>	list value	List of Neutron Security Group UUIDs to be applied during cleaning of the nodes. Optional for the "neutron" network interface and not used for the "flat" or "noop" network interfaces. If not specified, default security group is used.
<b>default-domain-id = None</b>	string value	Optional domain ID to use with v3 and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.
<b>default-domain-name = None</b>	string value	Optional domain name to use with v3 API and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.
<b>domain-id = None</b>	string value	Domain ID to scope to
<b>domain-name = None</b>	string value	Domain name to scope to
<b>endpoint-override = None</b>	string value	Always use this endpoint URL for requests for this client. NOTE: The unversioned endpoint should be specified here; to request a particular API version, use the <b>version</b> , <b>min-version</b> , and/or <b>max-version</b> options.

Configuration option = Default value	Type	Description
<b>insecure = False</b>	boolean value	Verify HTTPS connections.
<b>keyfile = None</b>	string value	PEM encoded client certificate key file
<b>max-version = None</b>	string value	The maximum major version of a given API, intended to be used as the upper bound of a range with <code>min_version</code> . Mutually exclusive with <code>version</code> .
<b>min-version = None</b>	string value	The minimum major version of a given API, intended to be used as the lower bound of a range with <code>max_version</code> . Mutually exclusive with <code>version</code> . If <code>min_version</code> is given with no <code>max_version</code> it is as if <code>max_version</code> is "latest".
<b>password = None</b>	string value	User's password
<b>port_setup_delay = 0</b>	integer value	Delay value to wait for Neutron agents to setup sufficient DHCP configuration for port.
<b>project-domain-id = None</b>	string value	Domain ID containing project
<b>project-domain-name = None</b>	string value	Domain name containing project
<b>project-id = None</b>	string value	Project ID to scope to
<b>project-name = None</b>	string value	Project name to scope to
<b>provisioning_network = None</b>	string value	Neutron network UUID or name for the ramdisk to be booted into for provisioning nodes. Required for "neutron" network interface. If a name is provided, it must be unique among all networks or deploy will fail.
<b>provisioning_network_security_groups = []</b>	list value	List of Neutron Security Group UUIDs to be applied during provisioning of the nodes. Optional for the "neutron" network interface and not used for the "flat" or "noop" network interfaces. If not specified, default security group is used.
<b>region-name = None</b>	string value	The default <code>region_name</code> for endpoint URL discovery.

Configuration option = Default value	Type	Description
<b>rescuing_network = None</b>	string value	Neutron network UUID or name for booting the ramdisk for rescue mode. This is not the network that the rescue ramdisk will use post-boot – the tenant network is used for that. Required for "neutron" network interface, if rescue mode will be used. It is not used for the "flat" or "noop" network interfaces. If a name is provided, it must be unique among all networks or rescue will fail.
<b>rescuing_network_security_groups = []</b>	list value	List of Neutron Security Group UUIDs to be applied during the node rescue process. Optional for the "neutron" network interface and not used for the "flat" or "noop" network interfaces. If not specified, the default security group is used.
<b>retries = 3</b>	integer value	Client retries in the case of a failed request.
<b>service-name = None</b>	string value	The default service_name for endpoint URL discovery.
<b>service-type = network</b>	string value	The default service_type for endpoint URL discovery.
<b>system-scope = None</b>	string value	Scope for system operations
<b>tenant-id = None</b>	string value	Tenant ID
<b>tenant-name = None</b>	string value	Tenant Name
<b>timeout = None</b>	integer value	Timeout value for http requests
<b>trust-id = None</b>	string value	Trust ID
<b>url = None</b>	string value	URL for connecting to neutron. Default value translates to <a href="http://\$my_ip:9696">http://\$my_ip:9696</a> when auth_strategy is <i>noauth</i> , and to discovery from Keystone catalog when auth_strategy is <i>keystone</i> .
<b>url_timeout = 30</b>	integer value	Timeout value for connecting to neutron in seconds.
<b>user-domain-id = None</b>	string value	User's domain id
<b>user-domain-name = None</b>	string value	User's domain name

Configuration option = Default value	Type	Description
<b>user-id = None</b>	string value	User id
<b>username = None</b>	string value	Username
<b>valid-interfaces = [u'internal', u'public']</b>	list value	List of interfaces, in order of preference, for endpoint URL.
<b>version = None</b>	string value	Minimum Major API version within a given Major API version for endpoint URL discovery. Mutually exclusive with min_version and max_version

### 7.1.31. oneflow

The following table outlines the options available under the **[oneflow]** group in the **/etc/ironic/ironic.conf** file.

Table 7.31. oneflow

Configuration option = Default value	Type	Description
<b>allow_insecure_connections = False</b>	boolean value	Option to allow insecure connection with OneView.
<b>enable_periodic_tasks = True</b>	boolean value	Whether to enable the periodic tasks for OneView driver be aware when OneView hardware resources are taken and released by Ironic or OneView users and proactively manage nodes in clean fail state according to Dynamic Allocation model of hardware resources allocation in OneView.
<b>manager_url = None</b>	string value	URL where OneView is available.
<b>password = None</b>	string value	OneView password to be used.
<b>periodic_check_interval = 300</b>	integer value	Period (in seconds) for periodic tasks to be executed when enable_periodic_tasks=True.
<b>tls_cacert_file = None</b>	string value	Path to CA certificate.
<b>username = None</b>	string value	OneView username to be used.

### 7.1.32. oslo\_concurrency

The following table outlines the options available under the **[oslo\_concurrency]** group in the **/etc/ironic/ironic.conf** file.

Table 7.32. oslo\_concurrency

Configuration option = Default value	Type	Description
<b>disable_process_locking</b> = <b>False</b>	boolean value	Enables or disables inter-process locks.
<b>lock_path</b> = <b>None</b>	string value	Directory to use for lock files. For security, the specified directory should only be writable by the user running the processes that need locking. Defaults to environment variable OSLO_LOCK_PATH. If external locks are used, a lock path must be set.

### 7.1.33. oslo\_messaging\_amqp

The following table outlines the options available under the **[oslo\_messaging\_amqp]** group in the **/etc/ironic/ironic.conf** file.

Table 7.33. oslo\_messaging\_amqp

Configuration option = Default value	Type	Description
<b>addressing_mode</b> = <b>dynamic</b>	string value	Indicates the addressing mode used by the driver. Permitted values: <i>legacy</i> - use legacy non-routable addressing <i>routable</i> - use routable addresses <i>dynamic</i> - use legacy addresses if the message bus does not support routing otherwise use routable addressing
<b>allow_insecure_clients</b> = <b>False</b>	boolean value	Accept clients using either SSL or plain TCP
<b>anycast_address</b> = <b>anycast</b>	string value	Appended to the address prefix when sending to a group of consumers. Used by the message bus to identify messages that should be delivered in a round-robin fashion across consumers.
<b>broadcast_prefix</b> = <b>broadcast</b>	string value	address prefix used when broadcasting to all servers
<b>connection_retry_backoff</b> = <b>2</b>	integer value	Increase the connection_retry_interval by this many seconds after each unsuccessful failover attempt.
<b>connection_retry_interval</b> = <b>1</b>	integer value	Seconds to pause before attempting to re-connect.
<b>connection_retry_interval_max</b> = <b>30</b>	integer value	Maximum limit for connection_retry_interval + connection_retry_backoff



Configuration option = Default value	Type	Description
<b>container_name = None</b>	string value	Name for the AMQP container. must be globally unique. Defaults to a generated UUID
<b>default_notification_exchange = None</b>	string value	Exchange name used in notification addresses. Exchange name resolution precedence: Target.exchange if set else default_notification_exchange if set else control_exchange if set else <i>notify</i>
<b>default_notify_timeout = 30</b>	integer value	The deadline for a sent notification message delivery. Only used when caller does not provide a timeout expiry.
<b>default_reply_retry = 0</b>	integer value	The maximum number of attempts to re-send a reply message which failed due to a recoverable error.
<b>default_reply_timeout = 30</b>	integer value	The deadline for an rpc reply message delivery.
<b>default_rpc_exchange = None</b>	string value	Exchange name used in RPC addresses. Exchange name resolution precedence: Target.exchange if set else default_rpc_exchange if set else control_exchange if set else <i>rpc</i>
<b>default_send_timeout = 30</b>	integer value	The deadline for an rpc cast or call message delivery. Only used when caller does not provide a timeout expiry.
<b>default_sender_link_timeout = 600</b>	integer value	The duration to schedule a purge of idle sender links. Detach link after expiry.
<b>group_request_prefix = unicast</b>	string value	address prefix when sending to any server in group
<b>idle_timeout = 0</b>	integer value	Timeout for inactive connections (in seconds)
<b>link_retry_delay = 10</b>	integer value	Time to pause between re-connecting an AMQP 1.0 link that failed due to a recoverable error.
<b>multicast_address = multicast</b>	string value	Appended to the address prefix when sending a fanout message. Used by the message bus to identify fanout messages.
<b>notify_address_prefix = openstack.org/om/notify</b>	string value	Address prefix for all generated Notification addresses
<b>notify_server_credit = 100</b>	integer value	Window size for incoming Notification messages

Configuration option = Default value	Type	Description
<code>`password = `</code>	string value	Password for message broker authentication
<b>pre_settled = [u'rpc-cast', u'rpc-reply']</b>	multi valued	Send messages of this type pre-settled. Pre-settled messages will not receive acknowledgement from the peer. Note well: pre-settled messages may be silently discarded if the delivery fails. Permitted values: <i>rpc-call</i> - send RPC Calls pre-settled <i>rpc-reply</i> - send RPC Replies pre-settled <i>rpc-cast</i> - Send RPC Casts pre-settled <i>notify</i> - Send Notifications pre-settled
<b>pseudo_vhost = True</b>	boolean value	Enable virtual host support for those message buses that do not natively support virtual hosting (such as <i>qpidd</i> ). When set to true the virtual host name will be added to all message bus addresses, effectively creating a private <i>subnet</i> per virtual host. Set to False if the message bus supports virtual hosting using the <i>hostname</i> field in the AMQP 1.0 Open performative as the name of the virtual host.
<b>reply_link_credit = 200</b>	integer value	Window size for incoming RPC Reply messages.
<b>rpc_address_prefix = openstack.org/om/rpc</b>	string value	Address prefix for all generated RPC addresses
<b>rpc_server_credit = 100</b>	integer value	Window size for incoming RPC Request messages
<code>`sasl_config_dir = `</code>	string value	Path to directory that contains the SASL configuration
<code>`sasl_config_name = `</code>	string value	Name of configuration file (without <i>.conf</i> suffix)
<code>`sasl_default_realm = `</code>	string value	SASL realm to use if no realm present in username
<code>`sasl_mechanisms = `</code>	string value	Space separated list of acceptable SASL mechanisms
<b>server_request_prefix = exclusive</b>	string value	address prefix used when sending to a specific server
<b>ssl = False</b>	boolean value	Attempt to connect via SSL. If no other <i>ssl</i> -related parameters are given, it will use the system's CA-bundle to verify the server's certificate.
<code>`ssl_ca_file = `</code>	string value	CA certificate PEM file used to verify the server's certificate

Configuration option = Default value	Type	Description
<code>`ssl_cert_file = `</code>	string value	Self-identifying certificate PEM file for client authentication
<code>`ssl_key_file = `</code>	string value	Private key PEM file used to sign <code>ssl_cert_file</code> certificate (optional)
<b>ssl_key_password = None</b>	string value	Password for decrypting <code>ssl_key_file</code> (if encrypted)
<b>ssl_verify_vhost = False</b>	boolean value	By default SSL checks that the name in the server's certificate matches the hostname in the <code>transport_url</code> . In some configurations it may be preferable to use the virtual hostname instead, for example if the server uses the Server Name Indication TLS extension (rfc6066) to provide a certificate per virtual host. Set <code>ssl_verify_vhost</code> to <code>True</code> if the server's SSL certificate uses the virtual host name instead of the DNS name.
<b>trace = False</b>	boolean value	Debug: dump AMQP frames to stdout
<b>unicast_address = unicast</b>	string value	Appended to the address prefix when sending to a particular RPC/Notification server. Used by the message bus to identify messages sent to a single destination.
<code>`username = `</code>	string value	User name for message broker authentication

### 7.1.34. oslo\_messaging\_kafka

The following table outlines the options available under the **[oslo\_messaging\_kafka]** group in the `/etc/ironic/ironic.conf` file.

Table 7.34. oslo\_messaging\_kafka

Configuration option = Default value	Type	Description
<b>conn_pool_min_size = 2</b>	integer value	The pool size limit for connections expiration policy
<b>conn_pool_ttl = 1200</b>	integer value	The time-to-live in sec of idle connections in the pool
<b>consumer_group = oslo_messaging_consumer</b>	string value	Group id for Kafka consumer. Consumers in one group will coordinate message consumption

Configuration option = Default value	Type	Description
<b>kafka_consumer_timeout = 1.0</b>	floating point value	Default timeout(s) for Kafka consumers
<b>kafka_default_host = localhost</b>	string value	Default Kafka broker Host
<b>kafka_default_port = 9092</b>	port value	Default Kafka broker Port
<b>kafka_max_fetch_bytes = 1048576</b>	integer value	Max fetch bytes of Kafka consumer
<b>pool_size = 10</b>	integer value	Pool Size for Kafka Consumers
<b>producer_batch_size = 16384</b>	integer value	Size of batch for the producer async send
<b>producer_batch_timeout = 0.0</b>	floating point value	Upper bound on the delay for KafkaProducer batching in seconds

### 7.1.35. oslo\_messaging\_notifications

The following table outlines the options available under the **[oslo\_messaging\_notifications]** group in the **/etc/ironic/ironic.conf** file.

Table 7.35. oslo\_messaging\_notifications

Configuration option = Default value	Type	Description
<b>driver = []</b>	multi valued	The Drivers(s) to handle sending notifications. Possible values are messaging, messagingv2, routing, log, test, noop
<b>retry = -1</b>	integer value	The maximum number of attempts to re-send a notification message which failed to be delivered due to a recoverable error. 0 - No retry, -1 - indefinite
<b>topics = [u'notifications']</b>	list value	AMQP topic used for OpenStack notifications.
<b>transport_url = None</b>	string value	A URL representing the messaging driver to use for notifications. If not set, we fall back to the same configuration used for RPC.

### 7.1.36. oslo\_messaging\_rabbit

The following table outlines the options available under the `[oslo_messaging_rabbit]` group in the `/etc/ironic/ironic.conf` file.

Table 7.36. oslo\_messaging\_rabbit

Configuration option = Default value	Type	Description
<b>amqp_auto_delete = False</b>	boolean value	Auto-delete queues in AMQP.
<b>amqp_durable_queues = False</b>	boolean value	Use durable queues in AMQP.
<b>channel_max = None</b>	integer value	Maximum number of channels to allow
<b>connection_factory = single</b>	string value	Connection factory implementation
<b>default_notification_exchange = \${control_exchange}_notification</b>	string value	Exchange name for sending notifications
<b>default_notification_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending notification, -1 means infinite retry.
<b>default_rpc_exchange = \${control_exchange}_rpc</b>	string value	Exchange name for sending RPC messages
<b>default_rpc_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending RPC message, -1 means infinite retry. If actual retry attempts in not 0 the rpc request could be processed more than one time
<b>default_serializer_type = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>fake_rabbit = False</b>	boolean value	Deprecated, use <code>rpc_backend=kombu+memory</code> or <code>rpc_backend=fake</code>
<b>frame_max = None</b>	integer value	The maximum byte size for an AMQP frame
<b>heartbeat_interval = 3</b>	integer value	How often to send heartbeats for consumer's connections
<b>heartbeat_rate = 2</b>	integer value	How often times during the <code>heartbeat_timeout_threshold</code> we check the heartbeat.

Configuration option = Default value	Type	Description
<b>heartbeat_timeout_thresh old = 60</b>	integer value	Number of seconds after which the Rabbit broker is considered down if heartbeat's keep-alive fails (0 disable the heartbeat). EXPERIMENTAL
<b>host_connection_reconn ect_delay = 0.25</b>	floating point value	Set delay for reconnection to some host which has connection error
<b>kombu_compression = None</b>	string value	EXPERIMENTAL: Possible values are: gzip, bz2. If not set compression will not be used. This option may not be available in future versions.
<b>kombu_failover_strategy = round-robin</b>	string value	Determines how the next RabbitMQ node is chosen in case the one we are currently connected to becomes unavailable. Takes effect only if more than one RabbitMQ node is provided in config.
<b>kombu_missing_consum er_retry_timeout = 60</b>	integer value	How long to wait a missing client before abandoning to send it its replies. This value should not be longer than <code>rpc_response_timeout</code> .
<b>kombu_reconnect_delay = 1.0</b>	floating point value	How long to wait before reconnecting in response to an AMQP consumer cancel notification.
<b>notification_listener_pref etch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to notification listener.
<b>notification_persistence = False</b>	boolean value	Persist notification messages.
<b>notification_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending notification message
<b>pool_max_overflow = 0</b>	integer value	Maximum number of connections to create above <b>pool_max_size</b> .
<b>pool_max_size = 30</b>	integer value	Maximum number of connections to keep queued.
<b>pool_recycle = 600</b>	integer value	Lifetime of a connection (since creation) in seconds or None for no recycling. Expired connections are closed on acquire.
<b>pool_stale = 60</b>	integer value	Threshold at which inactive (since release) connections are considered stale in seconds or None for no staleness. Stale connections are closed on acquire.

Configuration option = Default value	Type	Description
<b>pool_timeout = 30</b>	integer value	Default number of seconds to wait for a connections to available
<b>rabbit_ha_queues = False</b>	boolean value	Try to use HA queues in RabbitMQ (x-ha-policy: all). If you change this option, you must wipe the RabbitMQ database. In RabbitMQ 3.0, queue mirroring is no longer controlled by the x-ha-policy argument when declaring a queue. If you just want to make sure that all queues (except those with auto-generated names) are mirrored across all nodes, run: <code>"rabbitmqctl set_policy HA ^(?!amq\).* {"ha-mode": "all"}"</code>
<b>rabbit_host = localhost</b>	string value	The RabbitMQ broker address where a single node is used.
<b>rabbit_hosts = [u'\$rabbit_host:\$rabbit_port']</b>	list value	RabbitMQ HA cluster host:port pairs.
<b>rabbit_interval_max = 30</b>	integer value	Maximum interval of RabbitMQ connection retries. Default is 30 seconds.
<b>rabbit_login_method = AMQPLAIN</b>	string value	The RabbitMQ login method.
<b>rabbit_max_retries = 0</b>	integer value	Maximum number of RabbitMQ connection retries. Default is 0 (infinite retry count).
<b>rabbit_password = guest</b>	string value	The RabbitMQ password.
<b>rabbit_port = 5672</b>	port value	The RabbitMQ broker port where a single node is used.
<b>rabbit_qos_prefetch_count = 0</b>	integer value	Specifies the number of messages to prefetch. Setting to zero allows unlimited messages.
<b>rabbit_retry_backoff = 2</b>	integer value	How long to backoff for between retries when connecting to RabbitMQ.
<b>rabbit_retry_interval = 1</b>	integer value	How frequently to retry connecting with RabbitMQ.
<b>rabbit_transient_queues_ttl = 1800</b>	integer value	Positive integer representing duration in seconds for queue TTL (x-expires). Queues which are unused for the duration of the TTL are automatically deleted. The parameter affects only reply and fanout queues.

Configuration option = Default value	Type	Description
<b>rabbit_userid = guest</b>	string value	The RabbitMQ userid.
<b>rabbit_virtual_host = /</b>	string value	The RabbitMQ virtual host.
<b>rpc_listener_prefetch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to rpc listener.
<b>rpc_queue_expiration = 60</b>	integer value	Time to live for rpc queues without consumers in seconds.
<b>rpc_reply_exchange = \${control_exchange}_rpc_reply</b>	string value	Exchange name for receiving RPC replies
<b>rpc_reply_listener_prefetch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to rpc reply listener.
<b>rpc_reply_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending reply. -1 means infinite retry during rpc_timeout
<b>rpc_reply_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending reply.
<b>rpc_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending RPC message
<b>socket_timeout = 0.25</b>	floating point value	Set socket timeout in seconds for connection's socket
<b>ssl = None</b>	boolean value	Enable SSL
<code>`ssl_ca_file = `</code>	string value	SSL certification authority file (valid only if SSL enabled).
<code>`ssl_cert_file = `</code>	string value	SSL cert file (valid only if SSL enabled).
<code>`ssl_key_file = `</code>	string value	SSL key file (valid only if SSL enabled).
<b>ssl_options = None</b>	dict value	Arguments passed to ssl.wrap_socket
<code>`ssl_version = `</code>	string value	SSL version to use (valid only if SSL enabled). Valid values are TLSv1 and SSLv23. SSLv2, SSLv3, TLSv1_1, and TLSv1_2 may be available on some distributions.



Configuration option = Default value	Type	Description
<b>tcp_user_timeout = 0.25</b>	floating point value	Set TCP_USER_TIMEOUT in seconds for connection's socket

### 7.1.37. oslo\_messaging\_zmq

The following table outlines the options available under the **[oslo\_messaging\_zmq]** group in the **/etc/ironic/ironic.conf** file.

Table 7.37. oslo\_messaging\_zmq

Configuration option = Default value	Type	Description
<b>rpc_ack_timeout_base = 15</b>	integer value	Number of seconds to wait for an ack from a cast/call. After each retry attempt this timeout is multiplied by some specified multiplier.
<b>rpc_ack_timeout_multiplier = 2</b>	integer value	Number to multiply base ack timeout by after each retry attempt.
<b>rpc_message_ttl = 300</b>	integer value	Expiration timeout in seconds of a sent/received message after which it is not tracked anymore by a client/server.
<b>rpc_poll_timeout = 1</b>	integer value	The default number of seconds that poll should wait. Poll raises timeout exception when timeout expired.
<b>rpc_retry_attempts = 3</b>	integer value	Default number of message sending attempts in case of any problems occurred: positive value N means at most N retries, 0 means no retries, None or -1 (or any other negative values) mean to retry forever. This option is used only if acknowledgments are enabled.
<b>rpc_thread_pool_size = 100</b>	integer value	Maximum number of (green) threads to work concurrently.
<b>rpc_use_acks = False</b>	boolean value	Wait for message acknowledgements from receivers. This mechanism works only via proxy without PUB/SUB.
<b>rpc_zmq_bind_address = *</b>	string value	ZeroMQ bind address. Should be a wildcard (*), an ethernet interface, or IP. The "host" option should point or resolve to this address.
<b>rpc_zmq_bind_port_retries = 100</b>	integer value	Number of retries to find free port number before fail with ZMQBindError.

Configuration option = Default value	Type	Description
<b>rpc_zmq_contexts = 1</b>	integer value	Number of ZeroMQ contexts, defaults to 1.
<b>rpc_zmq_host = ccsosp-undercloud.usersys.redhat.com</b>	string value	Name of this node. Must be a valid hostname, FQDN, or IP address. Must match "host" option, if running Nova.
<b>rpc_zmq_ipc_dir = /var/run/openstack</b>	string value	Directory for holding IPC sockets.
<b>rpc_zmq_matchmaker = redis</b>	string value	MatchMaker driver.
<b>rpc_zmq_max_port = 65536</b>	integer value	Maximal port number for random ports range.
<b>rpc_zmq_min_port = 49153</b>	port value	Minimal port number for random ports range.
<b>rpc_zmq_serialization = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>rpc_zmq_topic_backlog = None</b>	integer value	Maximum number of ingress messages to locally buffer per topic. Default is unlimited.
<b>subscribe_on = []</b>	list value	List of publisher hosts SubConsumer can subscribe on. This option has higher priority then the default publishers list taken from the matchmaker.
<b>use_dynamic_connections = False</b>	boolean value	This option makes direct connections dynamic or static. It makes sense only with <code>use_router_proxy=False</code> which means to use direct connections for direct message types (ignored otherwise).
<b>use_pub_sub = False</b>	boolean value	Use PUB/SUB pattern for fanout methods. PUB/SUB always uses proxy.
<b>use_router_proxy = False</b>	boolean value	Use ROUTER remote proxy.
<b>zmq_failover_connections = 2</b>	integer value	How many additional connections to a host will be made for failover reasons. This option is actual only in dynamic connections mode.

Configuration option = Default value	Type	Description
<b>zmq_immediate = True</b>	boolean value	This option configures round-robin mode in zmq socket. True means not keeping a queue when server side disconnects. False means to keep queue and messages even if server is disconnected, when the server appears we send all accumulated messages to it.
<b>zmq_linger = -1</b>	integer value	Number of seconds to wait before all pending messages will be sent after closing a socket. The default value of -1 specifies an infinite linger period. The value of 0 specifies no linger period. Pending messages shall be discarded immediately when the socket is closed. Positive values specify an upper bound for the linger period.
<b>zmq_target_expire = 300</b>	integer value	Expiration timeout in seconds of a name service record about existing target (< 0 means no timeout).
<b>zmq_target_update = 180</b>	integer value	Update period in seconds of a name service record about existing target.
<b>zmq_tcp_keepalive = -1</b>	integer value	Enable/disable TCP keepalive (KA) mechanism. The default value of -1 (or any other negative value) means to skip any overrides and leave it to OS default; 0 and 1 (or any other positive value) mean to disable and enable the option respectively.
<b>zmq_tcp_keepalive_cnt = -1</b>	integer value	The number of retransmissions to be carried out before declaring that remote end is not available. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_idle = -1</b>	integer value	The duration between two keepalive transmissions in idle condition. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_intvl = -1</b>	integer value	The duration between two successive keepalive retransmissions, if acknowledgement to the previous keepalive transmission is not received. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.

### 7.1.38. oslo\_policy

The following table outlines the options available under the **[oslo\_policy]** group in the `/etc/ironic/ironic.conf` file.

Table 7.38. oslo\_policy

Configuration option = Default value	Type	Description
<b>enforce_scope = False</b>	boolean value	This option controls whether or not to enforce scope when evaluating policies. If <b>True</b> , the scope of the token used in the request is compared to the <b>scope_types</b> of the policy being enforced. If the scopes do not match, an <b>InvalidScope</b> exception will be raised. If <b>False</b> , a message will be logged informing operators that policies are being invoked with mismatching scope.
<b>policy_default_rule = default</b>	string value	Default rule. Enforced when a requested rule is not found.
<b>policy_dirs = [u'policy.d']</b>	multi valued	Directories where policy configuration files are stored. They can be relative to any directory in the search path defined by the <code>config_dir</code> option, or absolute paths. The file defined by <code>policy_file</code> must exist for these directories to be searched. Missing or empty directories are ignored.
<b>policy_file = policy.json</b>	string value	The file that defines policies.
<b>remote_content_type = application/x-www-form-urlencoded</b>	string value	Content Type to send and receive data for REST based policy check
<b>remote_ssl_ca_cert_file = None</b>	string value	Absolute path to ca cert file for REST based policy check
<b>remote_ssl_client_cert_file = None</b>	string value	Absolute path to client cert for REST based policy check
<b>remote_ssl_client_key_file = None</b>	string value	Absolute path client key file REST based policy check
<b>remote_ssl_verify_server_cert = False</b>	boolean value	server identity verification for REST based policy check

### 7.1.39. profiler

The following table outlines the options available under the **[profiler]** group in the `/etc/ironic/ironic.conf` file.

Table 7.39. profiler

Configuration option = Default value	Type	Description
<b>connection_string = messaging://</b>	string value	<p>Connection string for a notifier backend. Default value is messaging:// which sets the notifier to oslo_messaging.</p> <p>Examples of possible values:</p> <ul style="list-style-type: none"> <li>● messaging://: use oslo_messaging driver for sending notifications.</li> <li>● mongodb://127.0.0.1:27017 : use mongodb driver for sending notifications.</li> <li>● elasticsearch://127.0.0.1:9200 : use elasticsearch driver for sending notifications.</li> </ul>
<b>enabled = False</b>	boolean value	<p>Enables the profiling for all services on this node. Default value is False (fully disable the profiling feature).</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True: Enables the feature</li> <li>● False: Disables the feature. The profiling cannot be started via this project operations. If the profiling is triggered by another project, this project part will be empty.</li> </ul>
<b>es_doc_type = notification</b>	string value	Document type for notification indexing in elasticsearch.
<b>es_scroll_size = 10000</b>	integer value	Elasticsearch splits large requests in batches. This parameter defines maximum size of each batch (for example: es_scroll_size=10000).
<b>es_scroll_time = 2m</b>	string value	This parameter is a time value parameter (for example: es_scroll_time=2m), indicating for how long the nodes that participate in the search will maintain relevant resources in order to continue and support it.

Configuration option = Default value	Type	Description
<b>hmac_keys = SECRET_KEY</b>	string value	<p>Secret key(s) to use for encrypting context data for performance profiling. This string value should have the following format: &lt;key1&gt;[,&lt;key2&gt;,...&lt;keyn&gt;], where each key is some random string. A user who triggers the profiling via the REST API has to set one of these keys in the headers of the REST API call to include profiling results of this node for this particular project.</p> <p>Both "enabled" flag and "hmac_keys" config options should be set to enable profiling. Also, to generate correct profiling information across all services at least one key needs to be consistent between OpenStack projects. This ensures it can be used from client side to generate the trace, containing information from all possible resources.</p>
<b>sentinel_service_name = mymaster</b>	string value	<p>Redis sentinel uses a service name to identify a master redis service. This parameter defines the name (for example: sentinel_service_name=mymaster).</p>
<b>socket_timeout = 0.1</b>	floating point value	<p>Redis sentinel provides a timeout option on the connections. This parameter defines that timeout (for example: socket_timeout=0.1).</p>
<b>trace_sqlalchemy = False</b>	boolean value	<p>Enables SQL requests profiling in services. Default value is False (SQL requests won't be traced).</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True: Enables SQL requests profiling. Each SQL query will be part of the trace and can be analyzed by how much time was spent for that.</li> <li>● False: Disables SQL requests profiling. The spent time is only shown on a higher level of operations. Single SQL queries cannot be analyzed this way.</li> </ul>

### 7.1.40. pxe

The following table outlines the options available under the **[pxe]** group in the **/etc/ironic/ironic.conf** file.

Table 7.40. pxe

Configuration option = Default value	Type	Description
<b>default_ephemeral_format = ext4</b>	string value	Default file system format for ephemeral partition, if one is created.
<b>dir_permission = None</b>	integer value	The permission that will be applied to the TFTP folders upon creation. This should be set to the permission such that the tftpsrv has access to read the contents of the configured TFTP folder. This setting is only required when the operating system's umask is restrictive such that ironic-conductor is creating files that cannot be read by the TFTP server. Setting to <None> will result in the operating system's umask to be utilized for the creation of new tftp folders. It is recommended that an octal representation is specified. For example: 0o755
<b>image_cache_size = 20480</b>	integer value	Maximum size (in MiB) of cache for master images, including those in use.
<b>image_cache_ttl = 10080</b>	integer value	Maximum TTL (in minutes) for old master images in cache.
<b>images_path = /var/lib/ironic/images/</b>	string value	On the ironic-conductor node, directory where images are stored on disk.
<b>instance_master_path = /var/lib/ironic/master_images</b>	string value	On the ironic-conductor node, directory where master instance images are stored on disk. Setting to the empty string disables image caching.
<b>ip_version = 4</b>	string value	The IP version that will be used for PXE booting. Defaults to 4. EXPERIMENTAL
<b>ipxe_boot_script = \$pybasedir/drivers/modules/boot.ipxe</b>	string value	On ironic-conductor node, the path to the main iPXE script file.
<b>ipxe_enabled = False</b>	boolean value	Enable iPXE boot.
<b>ipxe_timeout = 0</b>	integer value	Timeout value (in seconds) for downloading an image via iPXE. Defaults to 0 (no timeout)
<b>ipxe_use_swift = False</b>	boolean value	Download deploy and rescue images directly from swift using temporary URLs. If set to false (default), images are downloaded to the ironic-conductor node and served over its local HTTP server. Applicable only when <i>ipxe_enabled</i> option is set to true.

Configuration option = Default value	Type	Description
<b>pxe_append_params = nofb nomodeset vga=normal</b>	string value	Additional append parameters for baremetal PXE boot.
<b>pxe_bootfile_name = pxelinux.0</b>	string value	Bootfile DHCP parameter.
<b>pxe_bootfile_name_by_arch = ch = {}</b>	dict value	Bootfile DHCP parameter per node architecture. For example: aarch64:grubaa64.efi
<b>pxe_config_template = \$pybasedir/drivers/modules/pxe_config.template</b>	string value	On ironic-conductor node, template file for PXE configuration.
<b>pxe_config_template_by_arch = arch = {}</b>	dict value	On ironic-conductor node, template file for PXE configuration per node architecture. For example: aarch64:/opt/share/grubaa64_pxe_config.template
<b>tftp_master_path = /tftpboot/master_images</b>	string value	On ironic-conductor node, directory where master TFTP images are stored on disk. Setting to <None> disables image caching.
<b>tftp_root = /tftpboot</b>	string value	ironic-conductor node's TFTP root path. The ironic-conductor must have read/write access to this path.
<b>tftp_server = \$my_ip</b>	string value	IP address of ironic-conductor node's TFTP server.
<b>uefi_pxe_bootfile_name = bootx64.efi</b>	string value	Bootfile DHCP parameter for UEFI boot mode.
<b>uefi_pxe_config_template = \$pybasedir/drivers/modules/pxe_grub_config.template</b>	string value	On ironic-conductor node, template file for PXE configuration for UEFI boot loader.

### 7.1.41. service\_catalog

The following table outlines the options available under the **[service\_catalog]** group in the **/etc/ironic/ironic.conf** file.

Table 7.41. service\_catalog

Configuration option = Default value	Type	Description
<b>auth-url = None</b>	string value	Authentication URL



Configuration option = Default value	Type	Description
<b>auth_type = None</b>	string value	Authentication type to load
<b>cafile = None</b>	string value	PEM encoded Certificate Authority to use when verifying HTTPs connections.
<b>certfile = None</b>	string value	PEM encoded client certificate cert file
<b>default-domain-id = None</b>	string value	Optional domain ID to use with v3 and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.
<b>default-domain-name = None</b>	string value	Optional domain name to use with v3 API and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.
<b>domain-id = None</b>	string value	Domain ID to scope to
<b>domain-name = None</b>	string value	Domain name to scope to
<b>endpoint-override = None</b>	string value	Always use this endpoint URL for requests for this client. NOTE: The unversioned endpoint should be specified here; to request a particular API version, use the <b>version</b> , <b>min-version</b> , and/or <b>max-version</b> options.
<b>insecure = False</b>	boolean value	Verify HTTPS connections.
<b>keyfile = None</b>	string value	PEM encoded client certificate key file
<b>max-version = None</b>	string value	The maximum major version of a given API, intended to be used as the upper bound of a range with min_version. Mutually exclusive with version.
<b>min-version = None</b>	string value	The minimum major version of a given API, intended to be used as the lower bound of a range with max_version. Mutually exclusive with version. If min_version is given with no max_version it is as if max version is "latest".
<b>password = None</b>	string value	User's password
<b>project-domain-id = None</b>	string value	Domain ID containing project
<b>project-domain-name = None</b>	string value	Domain name containing project

Configuration option = Default value	Type	Description
<b>project-id = None</b>	string value	Project ID to scope to
<b>project-name = None</b>	string value	Project name to scope to
<b>region-name = None</b>	string value	The default region_name for endpoint URL discovery.
<b>service-name = None</b>	string value	The default service_name for endpoint URL discovery.
<b>service-type = baremetal</b>	string value	The default service_type for endpoint URL discovery.
<b>system-scope = None</b>	string value	Scope for system operations
<b>tenant-id = None</b>	string value	Tenant ID
<b>tenant-name = None</b>	string value	Tenant Name
<b>timeout = None</b>	integer value	Timeout value for http requests
<b>trust-id = None</b>	string value	Trust ID
<b>user-domain-id = None</b>	string value	User's domain id
<b>user-domain-name = None</b>	string value	User's domain name
<b>user-id = None</b>	string value	User id
<b>username = None</b>	string value	Username
<b>valid-interfaces = [u'internal', u'public']</b>	list value	List of interfaces, in order of preference, for endpoint URL.
<b>version = None</b>	string value	Minimum Major API version within a given Major API version for endpoint URL discovery. Mutually exclusive with min_version and max_version

### 7.1.42. snmp

The following table outlines the options available under the **[snmp]** group in the `/etc/ironic/ironic.conf` file.

Table 7.42. snmp

Configuration option = Default value	Type	Description
<b>power_timeout = 10</b>	integer value	Seconds to wait for power action to be completed
<b>reboot_delay = 0</b>	integer value	Time (in seconds) to sleep between when rebooting (powering off and on again)
<b>udp_transport_retries = 5</b>	integer value	Maximum number of UDP request retries, 0 means no retries.
<b>udp_transport_timeout = 1.0</b>	floating point value	Response timeout in seconds used for UDP transport. Timeout should be a multiple of 0.5 seconds and is applicable to each retry.

### 7.1.43. ssl

The following table outlines the options available under the **[ssl]** group in the `/etc/ironic/ironic.conf` file.

Table 7.43. ssl

Configuration option = Default value	Type	Description
<b>ca_file = None</b>	string value	CA certificate file to use to verify connecting clients.
<b>cert_file = None</b>	string value	Certificate file to use when starting the server securely.
<b>ciphers = None</b>	string value	Sets the list of available ciphers. value should be a string in the OpenSSL cipher list format.
<b>key_file = None</b>	string value	Private key file to use when starting the server securely.
<b>version = None</b>	string value	SSL version to use (valid only if SSL enabled). Valid values are TLSv1 and SSLv23. SSLv2, SSLv3, TLSv1_1, and TLSv1_2 may be available on some distributions.

### 7.1.44. swift

The following table outlines the options available under the **[swift]** group in the `/etc/ironic/ironic.conf` file.

Table 7.44. swift

Configuration option = Default value	Type	Description
<b>auth-url = None</b>	string value	Authentication URL
<b>auth_type = None</b>	string value	Authentication type to load
<b>cafile = None</b>	string value	PEM encoded Certificate Authority to use when verifying HTTPs connections.
<b>certfile = None</b>	string value	PEM encoded client certificate cert file
<b>default-domain-id = None</b>	string value	Optional domain ID to use with v3 and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.
<b>default-domain-name = None</b>	string value	Optional domain name to use with v3 API and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.
<b>domain-id = None</b>	string value	Domain ID to scope to
<b>domain-name = None</b>	string value	Domain name to scope to
<b>endpoint-override = None</b>	string value	Always use this endpoint URL for requests for this client. NOTE: The unversioned endpoint should be specified here; to request a particular API version, use the <b>version</b> , <b>min-version</b> , and/or <b>max-version</b> options.
<b>insecure = False</b>	boolean value	Verify HTTPS connections.
<b>keyfile = None</b>	string value	PEM encoded client certificate key file
<b>max-version = None</b>	string value	The maximum major version of a given API, intended to be used as the upper bound of a range with min_version. Mutually exclusive with version.
<b>min-version = None</b>	string value	The minimum major version of a given API, intended to be used as the lower bound of a range with max_version. Mutually exclusive with version. If min_version is given with no max_version it is as if max version is "latest".
<b>password = None</b>	string value	User's password
<b>project-domain-id = None</b>	string value	Domain ID containing project

Configuration option = Default value	Type	Description
<b>project-domain-name = None</b>	string value	Domain name containing project
<b>project-id = None</b>	string value	Project ID to scope to
<b>project-name = None</b>	string value	Project name to scope to
<b>region-name = None</b>	string value	The default region_name for endpoint URL discovery.
<b>service-name = None</b>	string value	The default service_name for endpoint URL discovery.
<b>service-type = object-store</b>	string value	The default service_type for endpoint URL discovery.
<b>swift_max_retries = 2</b>	integer value	Maximum number of times to retry a Swift request, before failing.
<b>system-scope = None</b>	string value	Scope for system operations
<b>tenant-id = None</b>	string value	Tenant ID
<b>tenant-name = None</b>	string value	Tenant Name
<b>timeout = None</b>	integer value	Timeout value for http requests
<b>trust-id = None</b>	string value	Trust ID
<b>user-domain-id = None</b>	string value	User's domain id
<b>user-domain-name = None</b>	string value	User's domain name
<b>user-id = None</b>	string value	User id
<b>username = None</b>	string value	Username
<b>valid-interfaces = [u'internal', u'public']</b>	list value	List of interfaces, in order of preference, for endpoint URL.
<b>version = None</b>	string value	Minimum Major API version within a given Major API version for endpoint URL discovery. Mutually exclusive with min_version and max_version

## 7.1.45. xclarity

The following table outlines the options available under the **[xclarity]** group in the **/etc/ironic/ironic.conf** file.

Table 7.45. xclarity

Configuration option = Default value	Type	Description
<b>manager_ip = None</b>	string value	IP address of XClarity controller.
<b>password = None</b>	string value	Password for XClarity controller username.
<b>port = 443</b>	port value	Port to be used for XClarity operations.
<b>username = None</b>	string value	Username to access the XClarity controller.

## CHAPTER 8. IRONIC-INSPECTOR

The following chapter contains information about the configuration options in the **ironic-inspector** service.

### 8.1. INSPECTOR.CONF

This section contains options for the `/etc/ironic-inspector/inspector.conf` file.

#### 8.1.1. DEFAULT

The following table outlines the options available under the **[DEFAULT]** group in the `/etc/ironic-inspector/inspector.conf` file.

Table 8.1. DEFAULT

Configuration option = Default value	Type	Description
<b>api_max_limit = 1000</b>	integer value	Limit the number of elements an API list-call returns
<b>auth_strategy = keystone</b>	string value	Authentication method used on the ironic-inspector API. Either "noauth" or "keystone" are currently valid options. "noauth" will disable all authentication.
<b>clean_up_period = 60</b>	integer value	Amount of time in seconds, after which repeat clean up of timed out nodes and old nodes status information.
<b>debug = False</b>	boolean value	If set to true, the logging level will be set to DEBUG instead of the default INFO level.
<b>default_log_levels =</b> [u'sqlalchemy=WARNING', u'iso8601=WARNING', u'requests=WARNING', u'urllib3.connectionpool= WARNING', u'keystonemiddleware=W ARNING', u'swiftclient=WARNING', u'keystoneauth=WARNIN G', u'ironicclient=WARNING']	list value	List of package logging levels in logger=LEVEL pairs. This option is ignored if log_config_append is set.
<b>fatal_deprecations =</b> <b>False</b>	boolean value	Enables or disables fatal status of deprecations.
<code>`instance_format = [instance: %(uuid)s] `</code>	string value	The format for an instance that is passed with the log message.

Configuration option = Default value	Type	Description
<code>`instance_uuid_format = [instance:%(uuid)s]`</code>	string value	The format for an instance UUID that is passed with the log message.
<b>introspection_delay = 5</b>	integer value	Delay (in seconds) between two introspections.
<b>ipmi_address_fields = [u'ilo_address', u'drac_host', u'drac_address', u'cimc_address']</b>	list value	Ironic driver_info fields that are equivalent to ipmi_address.
<b>listen_address = 0.0.0.0</b>	string value	IP to listen on.
<b>listen_port = 5050</b>	port value	Port to listen on.
<b>log-config-append = None</b>	string value	The name of a logging configuration file. This file is appended to any existing logging configuration files. For details about logging configuration files, see the Python logging module documentation. Note that when logging configuration files are used then all logging configuration is set in the configuration file and other logging configuration options are ignored (for example, logging_context_format_string).
<b>log-date-format = %Y-%m-%d %H:%M:%S</b>	string value	Defines the format string for <code>%(asctime)s</code> in log records. Default: <code>%(default)s</code> . This option is ignored if <code>log_config_append</code> is set.
<b>log-dir = None</b>	string value	(Optional) The base directory used for relative <code>log_file</code> paths. This option is ignored if <code>log_config_append</code> is set.
<b>log-file = None</b>	string value	(Optional) Name of log file to send logging output to. If no default is set, logging will go to <code>stderr</code> as defined by <code>use_stderr</code> . This option is ignored if <code>log_config_append</code> is set.
<b>logging_context_format_string = %(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [%(request_id)s %(user_identity)s] %(instance)s%(message)s</b>	string value	Format string to use for log messages with context.
<b>logging_debug_format_s uffix = %(funcName)s % (pathname)s:%(lineno)d</b>	string value	Additional data to append to log message when logging level for the message is <code>DEBUG</code> .



Configuration option = Default value	Type	Description
<b>logging_default_format_string =</b> <code>%(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [-] %(instance)s%(message)s</code>	string value	Format string to use for log messages when context is undefined.
<b>logging_exception_prefix =</b> <code>%(asctime)s.%(msecs)03d %(process)d ERROR %(name)s %(instance)s</code>	string value	Prefix each line of exception output with this format.
<b>logging_user_identity_format =</b> <code>%(user)s %(tenant)s %(domain)s %(user_domain)s %(project_domain)s</code>	string value	Defines the format string for <code>%(user_identity)s</code> that is used in <code>logging_context_format_string</code> .
<b>max_concurrency = 1000</b>	integer value	The green thread pool size.
<b>node_status_keep_time = 0</b>	integer value	For how much time (in seconds) to keep status information about nodes after introspection was finished for them. Set to 0 (the default) to disable the timeout.
<b>publish_errors = False</b>	boolean value	Enables or disables publication of error events.
<b>rate_limit_burst = 0</b>	integer value	Maximum number of logged messages per <code>rate_limit_interval</code> .
<b>rate_limit_except_level = CRITICAL</b>	string value	Log level name used by rate limiting: CRITICAL, ERROR, INFO, WARNING, DEBUG or empty string. Logs with level greater or equal to <code>rate_limit_except_level</code> are not filtered. An empty string means that all levels are filtered.
<b>rate_limit_interval = 0</b>	integer value	Interval, number of seconds, of log rate limiting.
<b>rootwrap_config =</b> <code>/etc/ironic-inspector/rootwrap.conf</code>	string value	Path to the rootwrap configuration file to use for running commands as root
<code>`ssl_cert_path = `</code>	string value	Path to SSL certificate
<code>`ssl_key_path = `</code>	string value	Path to SSL key

Configuration option = Default value	Type	Description
<b>syslog-log-facility = LOG_USER</b>	string value	Syslog facility to receive log lines. This option is ignored if log_config_append is set.
<b>timeout = 3600</b>	integer value	Timeout after which introspection is considered failed, set to 0 to disable.
<b>use-journal = False</b>	boolean value	Enable journald for logging. If running in a systemd environment you may wish to enable journal support. Doing so will use the journal native protocol which includes structured metadata in addition to log messages. This option is ignored if log_config_append is set.
<b>use-json = False</b>	boolean value	Use JSON formatting for logging. This option is ignored if log_config_append is set.
<b>use-syslog = False</b>	boolean value	Use syslog for logging. Existing syslog format is DEPRECATED and will be changed later to honor RFC5424. This option is ignored if log_config_append is set.
<b>use_ssl = False</b>	boolean value	SSL Enabled/Disabled
<b>use_stderr = False</b>	boolean value	Log output to standard error. This option is ignored if log_config_append is set.
<b>watch-log-file = False</b>	boolean value	Uses logging handler designed to watch file system. When log file is moved or removed this handler will open a new log file with specified path instantaneously. It makes sense only if log_file option is specified and Linux platform is used. This option is ignored if log_config_append is set.

### 8.1.2. capabilities

The following table outlines the options available under the **[capabilities]** group in the **/etc/ironic-inspector/inspector.conf** file.

Table 8.2. capabilities

Configuration option = Default value	Type	Description
<b>boot_mode = False</b>	boolean value	Whether to store the boot mode (BIOS or UEFI).

Configuration option = Default value	Type	Description
<b>cpu_flags = {u'aes': u'cpu_aes', u'svm': u'cpu_vt', u'pdpe1gb': u'cpu_hugepages_1g', u'smx': u'cpu_txt', u'pse': u'cpu_hugepages', u'vmx': u'cpu_vt'}</b>	dict value	Mapping between a CPU flag and a capability to set if this flag is present.

### 8.1.3. cors

The following table outlines the options available under the **[cors]** group in the `/etc/ironic-inspector/inspector.conf` file.

Table 8.3. cors

Configuration option = Default value	Type	Description
<b>allow_credentials = True</b>	boolean value	Indicate that the actual request can include user credentials
<b>allow_headers = [u'X-Auth-Token', u'X-OpenStack-Ironic-Inspector-API-Minimum-Version', u'X-OpenStack-Ironic-Inspector-API-Maximum-Version', u'X-OpenStack-Ironic-Inspector-API-Version']</b>	list value	Indicate which header field names may be used during the actual request.
<b>allow_methods = [u'GET', u'POST', u'PUT', u'HEAD', u'PATCH', u'DELETE', u'OPTIONS']</b>	list value	Indicate which methods can be used during the actual request.
<b>allowed_origin = None</b>	list value	Indicate whether this resource may be shared with the domain received in the requests "origin" header. Format: "<protocol>://<host>[:<port>]", no trailing slash. Example: <a href="https://horizon.example.com">https://horizon.example.com</a>
<b>expose_headers = []</b>	list value	Indicate which headers are safe to expose to the API. Defaults to HTTP Simple Headers.
<b>max_age = 3600</b>	integer value	Maximum cache age of CORS preflight requests.

### 8.1.4. database

The following table outlines the options available under the **[database]** group in the `/etc/ironic-inspector/inspector.conf` file.

Table 8.4. database

Configuration option = Default value	Type	Description
<b>backend = sqlalchemy</b>	string value	The back end to use for the database.
<b>connection = None</b>	string value	The SQLAlchemy connection string to use to connect to the database.
<b>connection_debug = 0</b>	integer value	Verbosity of SQL debugging information: 0=None, 100=Everything.
<b>connection_recycle_time = 3600</b>	integer value	Connections which have been present in the connection pool longer than this number of seconds will be replaced with a new one the next time they are checked out from the pool.
<b>connection_trace = False</b>	boolean value	Add Python stack traces to SQL as comment strings.
<b>db_inc_retry_interval = True</b>	boolean value	If True, increases the interval between retries of a database operation up to <code>db_max_retry_interval</code> .
<b>db_max_retries = 20</b>	integer value	Maximum retries in case of connection error or deadlock error before error is raised. Set to -1 to specify an infinite retry count.
<b>db_max_retry_interval = 10</b>	integer value	If <code>db_inc_retry_interval</code> is set, the maximum seconds between retries of a database operation.
<b>db_retry_interval = 1</b>	integer value	Seconds between retries of a database transaction.
<b>max_overflow = 50</b>	integer value	If set, use this value for <code>max_overflow</code> with SQLAlchemy.
<b>max_pool_size = 5</b>	integer value	Maximum number of SQL connections to keep open in a pool. Setting a value of 0 indicates no limit.
<b>max_retries = 10</b>	integer value	Maximum number of database connection retries during startup. Set to -1 to specify an infinite retry count.
<b>min_pool_size = 1</b>	integer value	Minimum number of SQL connections to keep open in a pool.

Configuration option = Default value	Type	Description
<b>mysql_enable_ndb = False</b>	boolean value	If True, transparently enables support for handling MySQL Cluster (NDB).
<b>mysql_sql_mode = TRADITIONAL</b>	string value	The SQL mode to be used for MySQL sessions. This option, including the default, overrides any server-set SQL mode. To use whatever SQL mode is set by the server configuration, set this to no value. Example: <code>mysql_sql_mode=</code>
<b>pool_timeout = None</b>	integer value	If set, use this value for <code>pool_timeout</code> with SQLAlchemy.
<b>retry_interval = 10</b>	integer value	Interval between retries of opening a SQL connection.
<b>slave_connection = None</b>	string value	The SQLAlchemy connection string to use to connect to the slave database.
<b>sqlite_synchronous = True</b>	boolean value	If True, SQLite uses synchronous mode.
<b>use_db_reconnect = False</b>	boolean value	Enable the experimental use of database reconnect on connection lost.

### 8.1.5. discovery

The following table outlines the options available under the **[discovery]** group in the `/etc/ironic-inspector/inspector.conf` file.

Table 8.5. discovery

Configuration option = Default value	Type	Description
<b>enroll_node_driver = fake</b>	string value	The name of the Ironic driver used by the enroll hook when creating a new node in Ironic.

### 8.1.6. dnsmasq\_pxe\_filter

The following table outlines the options available under the **[dnsmasq\_pxe\_filter]** group in the `/etc/ironic-inspector/inspector.conf` file.

Table 8.6. dnsmasq\_pxe\_filter

Configuration option = Default value	Type	Description
<b>dhcp_hostsdir = /var/lib/ironic- inspector/dhcp-hostsdir</b>	string value	The MAC address cache directory, exposed to dnsmasq. This directory is expected to be in exclusive control of the driver.
<code>^dnsmasq_start_command = ^`</code>	string value	A (shell) command line to start the dnsmasq service upon filter initialization. Default: don't start.
<code>^dnsmasq_stop_command = ^`</code>	string value	A (shell) command line to stop the dnsmasq service upon inspector (error) exit. Default: don't stop.
<b>purge_dhcp_hostsdir = True</b>	boolean value	Purge the hostsdir upon driver initialization. Setting to false should only be performed when the deployment of inspector is such that there are multiple processes executing inside of the same host and namespace. In this case, the Operator is responsible for setting up a custom cleaning facility.

### 8.1.7. iptables

The following table outlines the options available under the **[iptables]** group in the `/etc/ironic-inspector/inspector.conf` file.

Table 8.7. iptables

Configuration option = Default value	Type	Description
<b>dnsmasq_interface = br- ctlplane</b>	string value	Interface on which dnsmasq listens, the default is for VM's.
<b>ethoib_interfaces = []</b>	list value	List of Ethernet Over InfiniBand interfaces on the Inspector host which are used for physical access to the DHCP network. Multiple interfaces would be attached to a bond or bridge specified in <code>dnsmasq_interface</code> . The MACs of the InfiniBand nodes which are not in desired state are going to be blacklisted based on the list of neighbor MACs on these interfaces.
<b>firewall_chain = ironic- inspector</b>	string value	iptables chain name to use.
<b>manage_firewall = True</b>	boolean value	Whether to manage firewall rules for PXE port. This configuration option was deprecated in favor of the <b>driver</b> option in the <b>pxe_filter</b> section. Please, use the <b>noop</b> filter driver to disable the firewall filtering or the <b>iptables</b> filter driver to enable it.

### 8.1.8. ironic

The following table outlines the options available under the **[ironic]** group in the `/etc/ironic-inspector/inspector.conf` file.

Table 8.8. ironic

Configuration option = Default value	Type	Description
<b>auth-url = None</b>	string value	Authentication URL
<b>auth_strategy = keystone</b>	string value	Method to use for authentication: noauth or keystone.
<b>auth_type = None</b>	string value	Authentication type to load
<b>cafile = None</b>	string value	PEM encoded Certificate Authority to use when verifying HTTPs connections.
<b>certfile = None</b>	string value	PEM encoded client certificate cert file
<b>default-domain-id = None</b>	string value	Optional domain ID to use with v3 and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.
<b>default-domain-name = None</b>	string value	Optional domain name to use with v3 API and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.
<b>domain-id = None</b>	string value	Domain ID to scope to
<b>domain-name = None</b>	string value	Domain name to scope to
<b>endpoint-override = None</b>	string value	Always use this endpoint URL for requests for this client. NOTE: The unversioned endpoint should be specified here; to request a particular API version, use the <b>version</b> , <b>min-version</b> , and/or <b>max-version</b> options.
<b>insecure = False</b>	boolean value	Verify HTTPS connections.
<b>ironic_url = http://localhost:6385/</b>	string value	Ironic API URL, used to set Ironic API URL when <code>auth_strategy</code> option is <code>noauth</code> or <code>auth_type</code> is "none" to work with standalone Ironic without keystone.
<b>keyfile = None</b>	string value	PEM encoded client certificate key file

Configuration option = Default value	Type	Description
<b>max-version = None</b>	string value	The maximum major version of a given API, intended to be used as the upper bound of a range with <code>min_version</code> . Mutually exclusive with <code>version</code> .
<b>max_retries = 30</b>	integer value	Maximum number of retries in case of conflict error (HTTP 409).
<b>min-version = None</b>	string value	The minimum major version of a given API, intended to be used as the lower bound of a range with <code>max_version</code> . Mutually exclusive with <code>version</code> . If <code>min_version</code> is given with no <code>max_version</code> it is as if <code>max version</code> is "latest".
<b>os_endpoint_type = internalURL</b>	string value	Ironic endpoint type.
<b>os_region = None</b>	string value	Keystone region used to get Ironic endpoints.
<b>os_service_type = baremetal</b>	string value	Ironic service type.
<b>password = None</b>	string value	User's password
<b>project-domain-id = None</b>	string value	Domain ID containing project
<b>project-domain-name = None</b>	string value	Domain name containing project
<b>project-id = None</b>	string value	Project ID to scope to
<b>project-name = None</b>	string value	Project name to scope to
<b>region-name = None</b>	string value	The default <code>region_name</code> for endpoint URL discovery.
<b>retry_interval = 2</b>	integer value	Interval between retries in case of conflict error (HTTP 409).
<b>service-name = None</b>	string value	The default <code>service_name</code> for endpoint URL discovery.
<b>service-type = baremetal</b>	string value	The default <code>service_type</code> for endpoint URL discovery.
<b>system-scope = None</b>	string value	Scope for system operations



Configuration option = Default value	Type	Description
<b>tenant-id = None</b>	string value	Tenant ID
<b>tenant-name = None</b>	string value	Tenant Name
<b>timeout = None</b>	integer value	Timeout value for http requests
<b>trust-id = None</b>	string value	Trust ID
<b>user-domain-id = None</b>	string value	User's domain id
<b>user-domain-name = None</b>	string value	User's domain name
<b>user-id = None</b>	string value	User id
<b>username = None</b>	string value	Username
<b>valid-interfaces = [u'internal', u'public']</b>	list value	List of interfaces, in order of preference, for endpoint URL.
<b>version = None</b>	string value	Minimum Major API version within a given Major API version for endpoint URL discovery. Mutually exclusive with min_version and max_version

### 8.1.9. keystone\_authtoken

The following table outlines the options available under the **[keystone\_authtoken]** group in the **/etc/ironic-inspector/inspector.conf** file.

Table 8.9. keystone\_authtoken

Configuration option = Default value	Type	Description
<b>auth_section = None</b>	string value	Config Section from which to load plugin specific options
<b>auth_type = None</b>	string value	Authentication type to load

Configuration option = Default value	Type	Description
<b>auth_uri = None</b>	string value	Complete "public" Identity API endpoint. This endpoint should not be an "admin" endpoint, as it should be accessible by all end users. Unauthenticated clients are redirected to this endpoint to authenticate. Although this endpoint should ideally be unversioned, client support in the wild varies. If you're using a versioned v2 endpoint here, then this should <b>not</b> be the same endpoint the service user utilizes for validating tokens, because normal end users may not be able to reach that endpoint. This option is deprecated in favor of <code>www_authenticate_uri</code> and will be removed in the S release.
<b>auth_version = None</b>	string value	API version of the admin Identity API endpoint.
<b>cache = None</b>	string value	Request environment key where the Swift cache object is stored. When <code>auth_token</code> middleware is deployed with a Swift cache, use this option to have the middleware share a caching backend with swift. Otherwise, use the <b>memcached_servers</b> option instead.
<b>cafile = None</b>	string value	A PEM encoded Certificate Authority to use when verifying HTTPs connections. Defaults to system CAs.
<b>certfile = None</b>	string value	Required if identity server requires client certificate
<b>check_revocations_for_cached = False</b>	boolean value	If true, the revocation list will be checked for cached tokens. This requires that PKI tokens are configured on the identity server.
<b>delay_auth_decision = False</b>	boolean value	Do not handle authorization requests within the middleware, but delegate the authorization decision to downstream WSGI components.
<b>enforce_token_bind = permissive</b>	string value	Used to control the use and type of token binding. Can be set to: "disabled" to not check token binding. "permissive" (default) to validate binding information if the bind type is of a form known to the server and ignore it if not. "strict" like "permissive" but if the bind type is unknown the token will be rejected. "required" any form of token binding is needed to be allowed. Finally the name of a binding method that must be present in tokens.

Configuration option = Default value	Type	Description
<b>hash_algorithms = [u'md5']</b>	list value	Hash algorithms to use for hashing PKI tokens. This may be a single algorithm or multiple. The algorithms are those supported by Python standard <code>hashlib.new()</code> . The hashes will be tried in the order given, so put the preferred one first for performance. The result of the first hash will be stored in the cache. This will typically be set to multiple values only while migrating from a less secure algorithm to a more secure one. Once all the old tokens are expired this option should be set to a single value for better performance.
<b>http_connect_timeout = None</b>	integer value	Request timeout value for communicating with Identity API server.
<b>http_request_max_retries = 3</b>	integer value	How many times are we trying to reconnect when communicating with Identity API Server.
<b>include_service_catalog = True</b>	boolean value	(Optional) Indicate whether to set the X-Service-Catalog header. If False, middleware will not ask for service catalog on token validation and will not set the X-Service-Catalog header.
<b>insecure = False</b>	boolean value	Verify HTTPS connections.
<b>keyfile = None</b>	string value	Required if identity server requires client certificate
<b>memcache_pool_conn_g et_timeout = 10</b>	integer value	(Optional) Number of seconds that an operation will wait to get a memcached client connection from the pool.
<b>memcache_pool_dead_re try = 300</b>	integer value	(Optional) Number of seconds memcached server is considered dead before it is tried again.
<b>memcache_pool_maxsize = 10</b>	integer value	(Optional) Maximum total number of open connections to every memcached server.
<b>memcache_pool_socket_ timeout = 3</b>	integer value	(Optional) Socket timeout in seconds for communicating with a memcached server.
<b>memcache_pool_unused_ timeout = 60</b>	integer value	(Optional) Number of seconds a connection to memcached is held unused in the pool before it is closed.
<b>memcache_secret_key = None</b>	string value	(Optional, mandatory if <code>memcache_security_strategy</code> is defined) This string is used for key derivation.

Configuration option = Default value	Type	Description
<b>memcache_security_strategy = None</b>	string value	(Optional) If defined, indicate whether token data should be authenticated or authenticated and encrypted. If MAC, token data is authenticated (with HMAC) in the cache. If ENCRYPT, token data is encrypted and authenticated in the cache. If the value is not one of these options or empty, auth_token will raise an exception on initialization.
<b>memcache_use_advanced_pool = False</b>	boolean value	(Optional) Use the advanced (eventlet safe) memcached client pool. The advanced pool will only work under python 2.x.
<b>memcached_servers = None</b>	list value	Optionally specify a list of memcached server(s) to use for caching. If left undefined, tokens will instead be cached in-process.
<b>region_name = None</b>	string value	The region in which the identity server can be found.
<b>revocation_cache_time = 10</b>	integer value	Determines the frequency at which the list of revoked tokens is retrieved from the Identity service (in seconds). A high number of revocation events combined with a low cache duration may significantly reduce performance. Only valid for PKI tokens. This option has been deprecated in the Ocata release and will be removed in the P release.
<b>service_token_roles = [u'service']</b>	list value	A choice of roles that must be present in a service token. Service tokens are allowed to request that an expired token can be used and so this check should tightly control that only actual services should be sending this token. Roles here are applied as an ANY check so any role in this list must be present. For backwards compatibility reasons this currently only affects the allow_expired check.
<b>service_token_roles_required = False</b>	boolean value	For backwards compatibility reasons we must let valid service tokens pass that don't pass the service_token_roles check as valid. Setting this true will become the default in a future release and should be enabled if possible.
<b>signing_dir = None</b>	string value	Directory used to cache files related to PKI tokens. This option has been deprecated in the Ocata release and will be removed in the P release.

Configuration option = Default value	Type	Description
<b>token_cache_time = 300</b>	integer value	In order to prevent excessive effort spent validating tokens, the middleware caches previously-seen tokens for a configurable duration (in seconds). Set to -1 to disable caching completely.
<b>www_authenticate_uri = None</b>	string value	Complete "public" Identity API endpoint. This endpoint should not be an "admin" endpoint, as it should be accessible by all end users. Unauthenticated clients are redirected to this endpoint to authenticate. Although this endpoint should ideally be unversioned, client support in the wild varies. If you're using a versioned v2 endpoint here, then this should <b>not</b> be the same endpoint the service user utilizes for validating tokens, because normal end users may not be able to reach that endpoint.

### 8.1.10. oslo\_policy

The following table outlines the options available under the **[oslo\_policy]** group in the `/etc/ironic-inspector/inspector.conf` file.

Table 8.10. oslo\_policy

Configuration option = Default value	Type	Description
<b>enforce_scope = False</b>	boolean value	This option controls whether or not to enforce scope when evaluating policies. If <b>True</b> , the scope of the token used in the request is compared to the <b>scope_types</b> of the policy being enforced. If the scopes do not match, an <b>InvalidScope</b> exception will be raised. If <b>False</b> , a message will be logged informing operators that policies are being invoked with mismatching scope.
<b>policy_default_rule = default</b>	string value	Default rule. Enforced when a requested rule is not found.
<b>policy_dirs = [u'policy.d']</b>	multi valued	Directories where policy configuration files are stored. They can be relative to any directory in the search path defined by the <code>config_dir</code> option, or absolute paths. The file defined by <code>policy_file</code> must exist for these directories to be searched. Missing or empty directories are ignored.
<b>policy_file = policy.json</b>	string value	The file that defines policies.

Configuration option = Default value	Type	Description
<b>remote_content_type = application/x-www-form-urlencoded</b>	string value	Content Type to send and receive data for REST based policy check
<b>remote_ssl_ca_cert_file = None</b>	string value	Absolute path to ca cert file for REST based policy check
<b>remote_ssl_client_cert_file = None</b>	string value	Absolute path to client cert for REST based policy check
<b>remote_ssl_client_key_file = None</b>	string value	Absolute path client key file REST based policy check
<b>remote_ssl_verify_server_cert = False</b>	boolean value	server identity verification for REST based policy check

### 8.1.11. pci\_devices

The following table outlines the options available under the **[pci\_devices]** group in the `/etc/ironic-inspector/inspector.conf` file.

Table 8.11. pci\_devices

Configuration option = Default value	Type	Description
<b>alias = []</b>	multi valued	An alias for PCI device identified by <i>vendor_id</i> and <i>product_id</i> fields. Format: {"vendor_id": "1234", "product_id": "5678", "name": "pci_dev1"}

### 8.1.12. processing

The following table outlines the options available under the **[processing]** group in the `/etc/ironic-inspector/inspector.conf` file.

Table 8.12. processing

Configuration option = Default value	Type	Description
---	------	-------------

Configuration option = Default value	Type	Description
<b>add_ports = pxe</b>	string value	Which MAC addresses to add as ports during introspection. Possible values: all (all MAC addresses), active (MAC addresses of NIC with IP addresses), pxe (only MAC address of NIC node PXE booted from, falls back to "active" if PXE MAC is not supplied by the ramdisk).
<b>always_store_ramdisk_logs = False</b>	boolean value	Whether to store ramdisk logs even if it did not return an error message (dependent upon "ramdisk_logs_dir" option being set).
<b>default_processing_hooks = ramdisk_error,root_disk_selection,scheduler,validate_interfaces,capabilities,pci_devices</b>	string value	Comma-separated list of default hooks for processing pipeline. Hook <i>scheduler</i> updates the node with the minimum properties required by the Nova scheduler. Hook <i>validate_interfaces</i> ensures that valid NIC data was provided by the ramdisk. Do not exclude these two unless you really know what you're doing.
<b>disk_partitioning_spacing = True</b>	boolean value	Whether to leave 1 GiB of disk size untouched for partitioning. Only has effect when used with the IPA as a ramdisk, for older ramdisk local_gb is calculated on the ramdisk side.
<b>keep_ports = all</b>	string value	Which ports (already present on a node) to keep after introspection. Possible values: all (do not delete anything), present (keep ports which MACs were present in introspection data), added (keep only MACs that we added during introspection).
<b>node_not_found_hook = None</b>	string value	The name of the hook to run when inspector receives inspection information from a node it isn't already aware of. This hook is ignored by default.
<b>overwrite_existing = True</b>	boolean value	Whether to overwrite existing values in node database. Disable this option to make introspection a non-destructive operation.
<b>power_off = True</b>	boolean value	Whether to power off a node after introspection.
<b>processing_hooks = \$default_processing_hooks</b>	string value	Comma-separated list of enabled hooks for processing pipeline. The default for this is \$default_processing_hooks, hooks can be added before or after the defaults like this: "prehook,\$default_processing_hooks,posthook".
<b>ramdisk_logs_dir = None</b>	string value	If set, logs from ramdisk will be stored in this directory.

Configuration option = Default value	Type	Description
<b>ramdisk_logs_filename_format =</b> <b>{uuid}_{dt:%Y%m%d-%H%M%S.%f}.tar.gz</b>	string value	File name template for storing ramdisk logs. The following replacements can be used: {uuid} - node UUID or "unknown", {bmc} - node BMC address or "unknown", {dt} - current UTC date and time, {mac} - PXE booting MAC or "unknown".
<b>store_data = none</b>	string value	Method for storing introspection data. If set to <i>none</i> , introspection data will not be stored.
<b>store_data_location = None</b>	string value	Name of the key to store the location of stored data in the extra column of the Ironic database.

### 8.1.13. pxe\_filter

The following table outlines the options available under the **[pxe\_filter]** group in the **/etc/ironic-inspector/inspector.conf** file.

Table 8.13. pxe\_filter

Configuration option = Default value	Type	Description
<b>driver = iptables</b>	string value	PXE boot filter driver to use, such as iptables
<b>sync_period = 15</b>	integer value	Amount of time in seconds, after which repeat periodic update of the filter.

### 8.1.14. swift

The following table outlines the options available under the **[swift]** group in the **/etc/ironic-inspector/inspector.conf** file.

Table 8.14. swift

Configuration option = Default value	Type	Description
<b>auth-url = None</b>	string value	Authentication URL
<b>auth_type = None</b>	string value	Authentication type to load
<b>cafile = None</b>	string value	PEM encoded Certificate Authority to use when verifying HTTPs connections.
<b>certfile = None</b>	string value	PEM encoded client certificate cert file



Configuration option = Default value	Type	Description
<b>container = ironic-inspector</b>	string value	Default Swift container to use when creating objects.
<b>default-domain-id = None</b>	string value	Optional domain ID to use with v3 and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.
<b>default-domain-name = None</b>	string value	Optional domain name to use with v3 API and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.
<b>delete_after = 0</b>	integer value	Number of seconds that the Swift object will last before being deleted. (set to 0 to never delete the object).
<b>domain-id = None</b>	string value	Domain ID to scope to
<b>domain-name = None</b>	string value	Domain name to scope to
<b>endpoint-override = None</b>	string value	Always use this endpoint URL for requests for this client. NOTE: The unversioned endpoint should be specified here; to request a particular API version, use the <b>version</b> , <b>min-version</b> , and/or <b>max-version</b> options.
<b>insecure = False</b>	boolean value	Verify HTTPS connections.
<b>keyfile = None</b>	string value	PEM encoded client certificate key file
<b>max-version = None</b>	string value	The maximum major version of a given API, intended to be used as the upper bound of a range with min_version. Mutually exclusive with version.
<b>max_retries = 2</b>	integer value	Maximum number of times to retry a Swift request, before failing.
<b>min-version = None</b>	string value	The minimum major version of a given API, intended to be used as the lower bound of a range with max_version. Mutually exclusive with version. If min_version is given with no max_version it is as if max version is "latest".
<b>os_endpoint_type = internalURL</b>	string value	Swift endpoint type.

Configuration option = Default value	Type	Description
<b>os_region = None</b>	string value	Keystone region to get endpoint for.
<b>os_service_type = object-store</b>	string value	Swift service type.
<b>password = None</b>	string value	User's password
<b>project-domain-id = None</b>	string value	Domain ID containing project
<b>project-domain-name = None</b>	string value	Domain name containing project
<b>project-id = None</b>	string value	Project ID to scope to
<b>project-name = None</b>	string value	Project name to scope to
<b>region-name = None</b>	string value	The default region_name for endpoint URL discovery.
<b>service-name = None</b>	string value	The default service_name for endpoint URL discovery.
<b>service-type = object-store</b>	string value	The default service_type for endpoint URL discovery.
<b>system-scope = None</b>	string value	Scope for system operations
<b>tenant-id = None</b>	string value	Tenant ID
<b>tenant-name = None</b>	string value	Tenant Name
<b>timeout = None</b>	integer value	Timeout value for http requests
<b>trust-id = None</b>	string value	Trust ID
<b>user-domain-id = None</b>	string value	User's domain id
<b>user-domain-name = None</b>	string value	User's domain name
<b>user-id = None</b>	string value	User id
<b>username = None</b>	string value	Username
<b>valid-interfaces = [u'internal', u'public']</b>	list value	List of interfaces, in order of preference, for endpoint URL.

Configuration option = Default value	Type	Description
<b>version = None</b>	string value	Minimum Major API version within a given Major API version for endpoint URL discovery. Mutually exclusive with min_version and max_version

## CHAPTER 9. KEYSTONE

The following chapter contains information about the configuration options in the **keystone** service.

### 9.1. KEYSTONE.CONF

This section contains options for the `/etc/keystone/keystone.conf` file.

#### 9.1.1. DEFAULT

The following table outlines the options available under the **[DEFAULT]** group in the `/etc/keystone/keystone.conf` file.

Table 9.1. DEFAULT

Configuration option = Default value	Type	Description
<b>admin_endpoint = None</b>	uri value	The base admin endpoint URL for Keystone that is advertised to clients (NOTE: this does NOT affect how Keystone listens for connections). Defaults to the base host URL of the request. For example, if keystone receives a request to <b>http://server:35357/v3/users</b> , then this will option will be automatically treated as <b>http://server:35357</b> . You should only need to set option if either the value of the base URL contains a path that keystone does not automatically infer ( <b>/prefix/v3</b> ), or if the endpoint should be found on a different host.
<b>admin_token = None</b>	string value	Using this feature is <b>NOT</b> recommended. Instead, use the <b>keystone-manage bootstrap</b> command. The value of this option is treated as a "shared secret" that can be used to bootstrap Keystone through the API. This "token" does not represent a user (it has no identity), and carries no explicit authorization (it effectively bypasses most authorization checks). If set to <b>None</b> , the value is ignored and the <b>admin_token</b> middleware is effectively disabled.
<b>conn_pool_min_size = 2</b>	integer value	The pool size limit for connections expiration policy
<b>conn_pool_ttl = 1200</b>	integer value	The time-to-live in sec of idle connections in the pool
<b>control_exchange = keystone</b>	string value	The default exchange under which topics are scoped. May be overridden by an exchange name specified in the <code>transport_url</code> option.

Configuration option = Default value	Type	Description
<b>crypt_strength = 10000</b>	integer value	The value passed as the keyword "rounds" to passlib's encrypt method. This option represents a trade off between security and performance. Higher values lead to slower performance, but higher security. Changing this option will only affect newly created passwords as existing password hashes already have a fixed number of rounds applied, so it is safe to tune this option in a running cluster. For more information, see <a href="https://pythonhosted.org/passlib/password_hash_api.html#choosing-the-right-rounds-value">https://pythonhosted.org/passlib/password_hash_api.html#choosing-the-right-rounds-value</a>
<b>debug = False</b>	boolean value	If set to true, the logging level will be set to DEBUG instead of the default INFO level.
<b>default_log_levels =</b> [u'amqp=WARN', u'amqpplib=WARN', u'boto=WARN', u'qpidd=WARN', u'sqlalchemy=WARN', u'suds=INFO', u'oslo.messaging=INFO', u'oslo_messaging=INFO', u'iso8601=WARN', u'requests.packages.urllib3.connectionpool=WARN', u'urllib3.connectionpool=WARN', u'websocket=WARN', u'requests.packages.urllib3.util.retry=WARN', u'urllib3.util.retry=WARN', u'keystonemiddleware=WARN', u'routes.middleware=WARN', u'stevedore=WARN', u'taskflow=WARN', u'keystoneauth=WARN', u'oslo.cache=INFO', u'dogpile.core.dogpile=INFO']	list value	List of package logging levels in logger=LEVEL pairs. This option is ignored if log_config_append is set.
<b>default_publisher_id = None</b>	string value	Default <b>publisher_id</b> for outgoing notifications. If left undefined, Keystone will default to using the server's host name.
<b>executor_thread_pool_size = 64</b>	integer value	Size of executor thread pool when executor is threading or eventlet.

Configuration option = Default value	Type	Description
<b>fatal_deprecations = False</b>	boolean value	Enables or disables fatal status of deprecations.
<b>insecure_debug = False</b>	boolean value	If set to true, then the server will return information in HTTP responses that may allow an unauthenticated or authenticated user to get more information than normal, such as additional details about why authentication failed. This may be useful for debugging but is insecure.
<code>`instance_format = [instance: %(uuid)s] `</code>	string value	The format for an instance that is passed with the log message.
<code>`instance_uuid_format = [instance:%(uuid)s] `</code>	string value	The format for an instance UUID that is passed with the log message.
<b>list_limit = None</b>	integer value	The maximum number of entities that will be returned in a collection. This global limit may be then overridden for a specific driver, by specifying a <code>list_limit</code> in the appropriate section (for example, <b>[assignment]</b> ). No limit is set by default. In larger deployments, it is recommended that you set this to a reasonable number to prevent operations like listing all users and projects from placing an unnecessary load on the system.
<b>log-config-append = None</b>	string value	The name of a logging configuration file. This file is appended to any existing logging configuration files. For details about logging configuration files, see the Python logging module documentation. Note that when logging configuration files are used then all logging configuration is set in the configuration file and other logging configuration options are ignored (for example, <code>logging_context_format_string</code> ).
<b>log-date-format = %Y-%m-%d %H:%M:%S</b>	string value	Defines the format string for <code>%(asctime)s</code> in log records. Default: <code>%(default)s</code> . This option is ignored if <code>log_config_append</code> is set.
<b>log-dir = None</b>	string value	(Optional) The base directory used for relative <code>log_file</code> paths. This option is ignored if <code>log_config_append</code> is set.

Configuration option = Default value	Type	Description
<b>log-file = None</b>	string value	(Optional) Name of log file to send logging output to. If no default is set, logging will go to stderr as defined by use_stderr. This option is ignored if log_config_append is set.
<b>logging_context_format_string = %(asctime)s.% (msecs)03d %(process)d %(levelname)s %(name)s [% (request_id)s % (user_identity)s] % (instance)s%(message)s</b>	string value	Format string to use for log messages with context.
<b>logging_debug_format_s uffix = %(funcName)s % (pathname)s:%(lineno)d</b>	string value	Additional data to append to log message when logging level for the message is DEBUG.
<b>logging_default_format_s tring = %(asctime)s.% (msecs)03d %(process)d %(levelname)s %(name)s [-] %(instance)s% (message)s</b>	string value	Format string to use for log messages when context is undefined.
<b>logging_exception_prefix = %(asctime)s.% (msecs)03d %(process)d ERROR %(name)s % (instance)s</b>	string value	Prefix each line of exception output with this format.
<b>logging_user_identity_for mat = %(user)s % (tenant)s %(domain)s % (user_domain)s % (project_domain)s</b>	string value	Defines the format string for %(user_identity)s that is used in logging_context_format_string.
<b>max_param_size = 64</b>	integer value	Limit the sizes of user & project ID/names.
<b>max_project_tree_depth = 5</b>	integer value	Maximum depth of the project hierarchy, excluding the project acting as a domain at the top of the hierarchy. WARNING: Setting it to a large value may adversely impact performance.
<b>max_token_size = 255</b>	integer value	Similar to <b>[DEFAULT] max_param_size</b> , but provides an exception for token values. With Fernet tokens, this can be set as low as 255. With UUID tokens, this should be set to 32).

Configuration option = Default value	Type	Description
<b>member_role_id = 9fe2ff9ee4384b1894a9087 8d3e92bab</b>	string value	Similar to the <b>[DEFAULT] member_role_name</b> option, this represents the default role ID used to associate users with their default projects in the v2 API. This will be used as the explicit role where one is not specified by the v2 API. You do not need to set this value unless you want keystone to use an existing role with a different ID, other than the arbitrarily defined <b>_member_role</b> (in which case, you should set <b>[DEFAULT] member_role_name</b> as well).
<b>member_role_name = _member_</b>	string value	This is the role name used in combination with the <b>[DEFAULT] member_role_id</b> option; see that option for more detail. You do not need to set this option unless you want keystone to use an existing role (in which case, you should set <b>[DEFAULT] member_role_id</b> as well).
<b>notification_format = cadf</b>	string value	Define the notification format for identity service events. A <b>basic</b> notification only has information about the resource being operated on. A <b>cadf</b> notification has the same information, as well as information about the initiator of the event. The <b>cadf</b> option is entirely backwards compatible with the <b>basic</b> option, but is fully CADF-compliant, and is recommended for auditing use cases.
<b>notification_opt_out = [u'identity.authenticate.su ccess', u'identity.authenticate.pe nding', u'identity.authenticate.fail ed']</b>	multi valued	You can reduce the number of notifications keystone emits by explicitly opting out. Keystone will not emit notifications that match the patterns expressed in this list. Values are expected to be in the form of <b>identity.&lt;resource_type&gt;.&lt;operation&gt;</b> . By default, all notifications related to authentication are automatically suppressed. This field can be set multiple times in order to opt-out of multiple notification topics. For example, the following suppresses notifications describing user creation or successful authentication events: notification_opt_out=identity.user.create notification_opt_out=identity.authenticate.success



Configuration option = Default value	Type	Description
<b>public_endpoint = None</b>	uri value	The base public endpoint URL for Keystone that is advertised to clients (NOTE: this does NOT affect how Keystone listens for connections). Defaults to the base host URL of the request. For example, if keystone receives a request to <b>http://server:5000/v3/users</b> , then this will option will be automatically treated as <b>http://server:5000</b> . You should only need to set option if either the value of the base URL contains a path that keystone does not automatically infer ( <b>/prefix/v3</b> ), or if the endpoint should be found on a different host.
<b>publish_errors = False</b>	boolean value	Enables or disables publication of error events.
<b>rate_limit_burst = 0</b>	integer value	Maximum number of logged messages per <code>rate_limit_interval</code> .
<b>rate_limit_except_level = CRITICAL</b>	string value	Log level name used by rate limiting: CRITICAL, ERROR, INFO, WARNING, DEBUG or empty string. Logs with level greater or equal to <code>rate_limit_except_level</code> are not filtered. An empty string means that all levels are filtered.
<b>rate_limit_interval = 0</b>	integer value	Interval, number of seconds, of log rate limiting.
<b>rpc_ack_timeout_base = 15</b>	integer value	Number of seconds to wait for an ack from a cast/call. After each retry attempt this timeout is multiplied by some specified multiplier.
<b>rpc_ack_timeout_multiplier = 2</b>	integer value	Number to multiply base ack timeout by after each retry attempt.
<b>rpc_backend = rabbit</b>	string value	The messaging driver to use, defaults to rabbit. Other drivers include amqp and zmq.
<b>rpc_conn_pool_size = 30</b>	integer value	Size of RPC connection pool.
<b>rpc_message_ttl = 300</b>	integer value	Expiration timeout in seconds of a sent/received message after which it is not tracked anymore by a client/server.
<b>rpc_poll_timeout = 1</b>	integer value	The default number of seconds that poll should wait. Poll raises timeout exception when timeout expired.
<b>rpc_response_timeout = 60</b>	integer value	Seconds to wait for a response from a call.

Configuration option = Default value	Type	Description
<b>rpc_retry_attempts = 3</b>	integer value	Default number of message sending attempts in case of any problems occurred: positive value N means at most N retries, 0 means no retries, None or -1 (or any other negative values) mean to retry forever. This option is used only if acknowledgments are enabled.
<b>rpc_thread_pool_size = 100</b>	integer value	Maximum number of (green) threads to work concurrently.
<b>rpc_use_acks = False</b>	boolean value	Wait for message acknowledgements from receivers. This mechanism works only via proxy without PUB/SUB.
<b>rpc_zmq_bind_address = *</b>	string value	ZeroMQ bind address. Should be a wildcard (*), an ethernet interface, or IP. The "host" option should point or resolve to this address.
<b>rpc_zmq_bind_port_retries = 100</b>	integer value	Number of retries to find free port number before fail with ZMQBindError.
<b>rpc_zmq_contexts = 1</b>	integer value	Number of ZeroMQ contexts, defaults to 1.
<b>rpc_zmq_host = ccsosp-undercloud.usersys.redhat.com</b>	string value	Name of this node. Must be a valid hostname, FQDN, or IP address. Must match "host" option, if running Nova.
<b>rpc_zmq_ipc_dir = /var/run/openstack</b>	string value	Directory for holding IPC sockets.
<b>rpc_zmq_matchmaker = redis</b>	string value	MatchMaker driver.
<b>rpc_zmq_max_port = 65536</b>	integer value	Maximal port number for random ports range.
<b>rpc_zmq_min_port = 49153</b>	port value	Minimal port number for random ports range.
<b>rpc_zmq_serialization = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>rpc_zmq_topic_backlog = None</b>	integer value	Maximum number of ingress messages to locally buffer per topic. Default is unlimited.

Configuration option = Default value	Type	Description
<b>secure_proxy_ssl_header = HTTP_X_FORWARDED_PROTO</b>	string value	The HTTP header used to determine the scheme for the original request, even if it was removed by an SSL terminating proxy.
<b>strict_password_check = False</b>	boolean value	If set to true, strict password length checking is performed for password manipulation. If a password exceeds the maximum length, the operation will fail with an HTTP 403 Forbidden error. If set to false, passwords are automatically truncated to the maximum length.
<b>subscribe_on = []</b>	list value	List of publisher hosts SubConsumer can subscribe on. This option has higher priority then the default publishers list taken from the matchmaker.
<b>syslog-log-facility = LOG_USER</b>	string value	Syslog facility to receive log lines. This option is ignored if log_config_append is set.
<b>transport_url = None</b>	string value	The network address and optional user credentials for connecting to the messaging backend, in URL format. The expected format is:  driver://[user:pass@]host:port[, [userN:passN@]hostN:portN]/virtual_host?query  Example: rabbit://rabbitmq:password@127.0.0.1:5672//  For full details on the fields in the URL see the documentation of oslo_messaging.TransportURL at <a href="https://docs.openstack.org/oslo.messaging/latest/reference/transport.html">https://docs.openstack.org/oslo.messaging/latest/reference/transport.html</a>
<b>use-journal = False</b>	boolean value	Enable journald for logging. If running in a systemd environment you may wish to enable journal support. Doing so will use the journal native protocol which includes structured metadata in addition to log messages. This option is ignored if log_config_append is set.
<b>use-json = False</b>	boolean value	Use JSON formatting for logging. This option is ignored if log_config_append is set.
<b>use-syslog = False</b>	boolean value	Use syslog for logging. Existing syslog format is DEPRECATED and will be changed later to honor RFC5424. This option is ignored if log_config_append is set.

Configuration option = Default value	Type	Description
<b>use_dynamic_connections = False</b>	boolean value	This option makes direct connections dynamic or static. It makes sense only with <code>use_router_proxy=False</code> which means to use direct connections for direct message types (ignored otherwise).
<b>use_pub_sub = False</b>	boolean value	Use PUB/SUB pattern for fanout methods. PUB/SUB always uses proxy.
<b>use_router_proxy = False</b>	boolean value	Use ROUTER remote proxy.
<b>use_stderr = False</b>	boolean value	Log output to standard error. This option is ignored if <code>log_config_append</code> is set.
<b>watch-log-file = False</b>	boolean value	Uses logging handler designed to watch file system. When log file is moved or removed this handler will open a new log file with specified path instantaneously. It makes sense only if <code>log_file</code> option is specified and Linux platform is used. This option is ignored if <code>log_config_append</code> is set.
<b>zmq_failover_connections = 2</b>	integer value	How many additional connections to a host will be made for failover reasons. This option is actual only in dynamic connections mode.
<b>zmq_immediate = True</b>	boolean value	This option configures round-robin mode in zmq socket. True means not keeping a queue when server side disconnects. False means to keep queue and messages even if server is disconnected, when the server appears we send all accumulated messages to it.
<b>zmq_linger = -1</b>	integer value	Number of seconds to wait before all pending messages will be sent after closing a socket. The default value of -1 specifies an infinite linger period. The value of 0 specifies no linger period. Pending messages shall be discarded immediately when the socket is closed. Positive values specify an upper bound for the linger period.
<b>zmq_target_expire = 300</b>	integer value	Expiration timeout in seconds of a name service record about existing target (< 0 means no timeout).
<b>zmq_target_update = 180</b>	integer value	Update period in seconds of a name service record about existing target.

Configuration option = Default value	Type	Description
<b>zmq_tcp_keepalive = -1</b>	integer value	Enable/disable TCP keepalive (KA) mechanism. The default value of -1 (or any other negative value) means to skip any overrides and leave it to OS default; 0 and 1 (or any other positive value) mean to disable and enable the option respectively.
<b>zmq_tcp_keepalive_cnt = -1</b>	integer value	The number of retransmissions to be carried out before declaring that remote end is not available. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_idle = -1</b>	integer value	The duration between two keepalive transmissions in idle condition. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_intvl = -1</b>	integer value	The duration between two successive keepalive retransmissions, if acknowledgement to the previous keepalive transmission is not received. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.

### 9.1.2. application\_credential

The following table outlines the options available under the **[application\_credential]** group in the `/etc/keystone/keystone.conf` file.

Table 9.2. application\_credential

Configuration option = Default value	Type	Description
<b>cache_time = None</b>	integer value	Time to cache application credential data in seconds. This has no effect unless global caching is enabled.
<b>caching = True</b>	boolean value	Toggle for application credential caching. This has no effect unless global caching is enabled.

Configuration option = Default value	Type	Description
<b>driver = sql</b>	string value	Entry point for the application credential backend driver in the <b>keystone.application_credential</b> namespace. Keystone only provides a <b>sql</b> driver, so there is no reason to change this unless you are providing a custom entry point.
<b>user_limit = -1</b>	integer value	Maximum number of application credentials a user is permitted to create. A value of -1 means unlimited. If a limit is not set, users are permitted to create application credentials at will, which could lead to bloat in the keystone database or open keystone to a DoS attack.

### 9.1.3. assignment

The following table outlines the options available under the **[assignment]** group in the `/etc/keystone/keystone.conf` file.

Table 9.3. assignment

Configuration option = Default value	Type	Description
<b>driver = sql</b>	string value	Entry point for the assignment backend driver (where role assignments are stored) in the <b>keystone.assignment</b> namespace. Only a SQL driver is supplied by keystone itself. Unless you are writing proprietary drivers for keystone, you do not need to set this option.
<b>prohibited_implied_role = [u'admin']</b>	list value	A list of role names which are prohibited from being an implied role.

### 9.1.4. auth

The following table outlines the options available under the **[auth]** group in the `/etc/keystone/keystone.conf` file.

Table 9.4. auth

Configuration option = Default value	Type	Description
---	------	-------------

Configuration option = Default value	Type	Description
<b>application_credential = None</b>	string value	Entry point for the <code>application_credential</code> auth plugin module in the <b><code>keystone.auth.application_credential</code></b> namespace. You do not need to set this unless you are overriding keystone's own <b><code>application_credential</code></b> authentication plugin.
<b>external = None</b>	string value	Entry point for the external ( <b><code>REMOTE_USER</code></b> ) auth plugin module in the <b><code>keystone.auth.external</code></b> namespace. Supplied drivers are <b><code>DefaultDomain</code></b> and <b><code>Domain</code></b> . The default driver is <b><code>DefaultDomain</code></b> , which assumes that all users identified by the username specified to keystone in the <b><code>REMOTE_USER</code></b> variable exist within the context of the default domain. The <b><code>Domain</code></b> option expects an additional environment variable be presented to keystone, <b><code>REMOTE_DOMAIN</code></b> , containing the domain name of the <b><code>REMOTE_USER</code></b> (if <b><code>REMOTE_DOMAIN</code></b> is not set, then the default domain will be used instead). You do not need to set this unless you are taking advantage of "external authentication", where the application server (such as Apache) is handling authentication instead of keystone.
<b>mapped = None</b>	string value	Entry point for the mapped auth plugin module in the <b><code>keystone.auth.mapped</code></b> namespace. You do not need to set this unless you are overriding keystone's own <b><code>mapped</code></b> authentication plugin.
<b>methods = [u'external', u'password', u'token', u'oauth1', u'mapped', u'application_credential']</b>	list value	Allowed authentication methods. Note: You should disable the <b><code>external</code></b> auth method if you are currently using federation. External auth and federation both use the <code>REMOTE_USER</code> variable. Since both the mapped and external plugin are being invoked to validate attributes in the request environment, it can cause conflicts.
<b>oauth1 = None</b>	string value	Entry point for the OAuth 1.0a auth plugin module in the <b><code>keystone.auth.oauth1</code></b> namespace. You do not need to set this unless you are overriding keystone's own <b><code>oauth1</code></b> authentication plugin.
<b>password = None</b>	string value	Entry point for the password auth plugin module in the <b><code>keystone.auth.password</code></b> namespace. You do not need to set this unless you are overriding keystone's own password authentication plugin.

Configuration option = Default value	Type	Description
<b>token = None</b>	string value	Entry point for the token auth plugin module in the <b>keystone.auth.token</b> namespace. You do not need to set this unless you are overriding keystone's own token authentication plugin.

### 9.1.5. cache

The following table outlines the options available under the **[cache]** group in the **/etc/keystone/keystone.conf** file.

Table 9.5. cache

Configuration option = Default value	Type	Description
<b>backend = dogpile.cache.null</b>	string value	Cache backend module. For eventlet-based or environments with hundreds of threaded servers, Memcache with pooling ( <code>oslo_cache.memcache_pool</code> ) is recommended. For environments with less than 100 threaded servers, Memcached ( <code>dogpile.cache.memcached</code> ) or Redis ( <code>dogpile.cache.redis</code> ) is recommended. Test environments with a single instance of the server can use the <code>dogpile.cache.memory</code> backend.
<b>backend_argument = []</b>	multi valued	Arguments supplied to the backend module. Specify this option once per argument to be passed to the <code>dogpile.cache</code> backend. Example format: " <code>&lt;argname&gt;: &lt;value&gt;</code> ".
<b>config_prefix = cache.oslo</b>	string value	Prefix for building the configuration dictionary for the cache region. This should not need to be changed unless there is another <code>dogpile.cache</code> region with the same configuration name.
<b>debug_cache_backend = False</b>	boolean value	Extra debugging from the cache backend (cache keys, <code>get/set/delete/etc</code> calls). This is only really useful if you need to see the specific cache-backend <code>get/set/delete</code> calls with the keys/values. Typically this should be left set to false.
<b>enabled = True</b>	boolean value	Global toggle for caching.
<b>expiration_time = 600</b>	integer value	Default TTL, in seconds, for any cached item in the <code>dogpile.cache</code> region. This applies to any cached method that doesn't have an explicit cache expiration time defined for it.



Configuration option = Default value	Type	Description
<b>memcache_dead_retry = 300</b>	integer value	Number of seconds memcached server is considered dead before it is tried again. (dogpile.cache.memcache and oslo_cache.memcache_pool backends only).
<b>memcache_pool_connection_get_timeout = 10</b>	integer value	Number of seconds that an operation will wait to get a memcache client connection.
<b>memcache_pool_maxsize = 10</b>	integer value	Max total number of open connections to every memcached server. (oslo_cache.memcache_pool backend only).
<b>memcache_pool_unused_timeout = 60</b>	integer value	Number of seconds a connection to memcached is held unused in the pool before it is closed. (oslo_cache.memcache_pool backend only).
<b>memcache_servers = [u'localhost:11211']</b>	list value	Memcache servers in the format of "host:port". (dogpile.cache.memcache and oslo_cache.memcache_pool backends only).
<b>memcache_socket_timeout = 3</b>	integer value	Timeout in seconds for every call to a server. (dogpile.cache.memcache and oslo_cache.memcache_pool backends only).
<b>proxies = []</b>	list value	Proxy classes to import that will affect the way the dogpile.cache backend functions. See the dogpile.cache documentation on changing-backend-behavior.

### 9.1.6. catalog

The following table outlines the options available under the **[catalog]** group in the `/etc/keystone/keystone.conf` file.

Table 9.6. catalog

Configuration option = Default value	Type	Description
<b>cache_time = None</b>	integer value	Time to cache catalog data (in seconds). This has no effect unless global and catalog caching are both enabled. Catalog data (services, endpoints, etc.) typically does not change frequently, and so a longer duration than the global default may be desirable.

Configuration option = Default value	Type	Description
<b>caching = True</b>	boolean value	Toggle for catalog caching. This has no effect unless global caching is enabled. In a typical deployment, there is no reason to disable this.
<b>driver = sql</b>	string value	Entry point for the catalog driver in the <b>keystone.catalog</b> namespace. Keystone provides a <b>sql</b> option (which supports basic CRUD operations through SQL), a <b>templated</b> option (which loads the catalog from a templated catalog file on disk), and a <b>endpoint_filter.sql</b> option (which supports arbitrary service catalogs per project).
<b>list_limit = None</b>	integer value	Maximum number of entities that will be returned in a catalog collection. There is typically no reason to set this, as it would be unusual for a deployment to have enough services or endpoints to exceed a reasonable limit.
<b>template_file = default_catalog.templates</b>	string value	Absolute path to the file used for the templated catalog backend. This option is only used if the <b>[catalog] driver</b> is set to <b>templated</b> .

### 9.1.7. cors

The following table outlines the options available under the **[cors]** group in the `/etc/keystone/keystone.conf` file.

Table 9.7. cors

Configuration option = Default value	Type	Description
<b>allow_credentials = True</b>	boolean value	Indicate that the actual request can include user credentials
<b>allow_headers = [u'X-Auth-Token', u'X-Openstack-Request-Id', u'X-Subject-Token', u'X-Project-Id', u'X-Project-Name', u'X-Project-Domain-Id', u'X-Project-Domain-Name', u'X-Domain-Id', u'X-Domain-Name']</b>	list value	Indicate which header field names may be used during the actual request.

Configuration option = Default value	Type	Description
<b>allow_methods = [u'GET', u'PUT', u'POST', u'DELETE', u'PATCH']</b>	list value	Indicate which methods can be used during the actual request.
<b>allowed_origin = None</b>	list value	Indicate whether this resource may be shared with the domain received in the requests "origin" header. Format: "<protocol>://<host>[:<port>]", no trailing slash. Example: <a href="https://horizon.example.com">https://horizon.example.com</a>
<b>expose_headers = [u'X-Auth-Token', u'X-Openstack-Request-Id', u'X-Subject-Token']</b>	list value	Indicate which headers are safe to expose to the API. Defaults to HTTP Simple Headers.
<b>max_age = 3600</b>	integer value	Maximum cache age of CORS preflight requests.

### 9.1.8. credential

The following table outlines the options available under the **[credential]** group in the `/etc/keystone/keystone.conf` file.

Table 9.8. credential

Configuration option = Default value	Type	Description
<b>driver = sql</b>	string value	Entry point for the credential backend driver in the <b>keystone.credential</b> namespace. Keystone only provides a <b>sql</b> driver, so there's no reason to change this unless you are providing a custom entry point.
<b>key_repository = /etc/keystone/credential-keys/</b>	string value	Directory containing Fernet keys used to encrypt and decrypt credentials stored in the credential backend. Fernet keys used to encrypt credentials have no relationship to Fernet keys used to encrypt Fernet tokens. Both sets of keys should be managed separately and require different rotation policies. Do not share this repository with the repository used to manage keys for Fernet tokens.
<b>provider = fernet</b>	string value	Entry point for credential encryption and decryption operations in the <b>keystone.credential.provider</b> namespace. Keystone only provides a <b>fernet</b> driver, so there's no reason to change this unless you are providing a custom entry point to encrypt and decrypt credentials.

## 9.1.9. database

The following table outlines the options available under the **[database]** group in the `/etc/keystone/keystone.conf` file.

Table 9.9. database

Configuration option = Default value	Type	Description
<b>backend = sqlalchemy</b>	string value	The back end to use for the database.
<b>connection = None</b>	string value	The SQLAlchemy connection string to use to connect to the database.
<b>connection_debug = 0</b>	integer value	Verbosity of SQL debugging information: 0=None, 100=Everything.
<b>connection_recycle_time = 3600</b>	integer value	Connections which have been present in the connection pool longer than this number of seconds will be replaced with a new one the next time they are checked out from the pool.
<b>connection_trace = False</b>	boolean value	Add Python stack traces to SQL as comment strings.
<b>db_inc_retry_interval = True</b>	boolean value	If True, increases the interval between retries of a database operation up to <code>db_max_retry_interval</code> .
<b>db_max_retries = 20</b>	integer value	Maximum retries in case of connection error or deadlock error before error is raised. Set to -1 to specify an infinite retry count.
<b>db_max_retry_interval = 10</b>	integer value	If <code>db_inc_retry_interval</code> is set, the maximum seconds between retries of a database operation.
<b>db_retry_interval = 1</b>	integer value	Seconds between retries of a database transaction.
<b>max_overflow = 50</b>	integer value	If set, use this value for <code>max_overflow</code> with SQLAlchemy.
<b>max_pool_size = 5</b>	integer value	Maximum number of SQL connections to keep open in a pool. Setting a value of 0 indicates no limit.
<b>max_retries = 10</b>	integer value	Maximum number of database connection retries during startup. Set to -1 to specify an infinite retry count.
<b>min_pool_size = 1</b>	integer value	Minimum number of SQL connections to keep open in a pool.

Configuration option = Default value	Type	Description
<b>mysql_enable_ndb = False</b>	boolean value	If True, transparently enables support for handling MySQL Cluster (NDB).
<b>mysql_sql_mode = TRADITIONAL</b>	string value	The SQL mode to be used for MySQL sessions. This option, including the default, overrides any server-set SQL mode. To use whatever SQL mode is set by the server configuration, set this to no value. Example: <code>mysql_sql_mode=</code>
<b>pool_timeout = None</b>	integer value	If set, use this value for <code>pool_timeout</code> with SQLAlchemy.
<b>retry_interval = 10</b>	integer value	Interval between retries of opening a SQL connection.
<b>slave_connection = None</b>	string value	The SQLAlchemy connection string to use to connect to the slave database.
<b>sqlite_synchronous = True</b>	boolean value	If True, SQLite uses synchronous mode.
<b>use_db_reconnect = False</b>	boolean value	Enable the experimental use of database reconnect on connection lost.

### 9.1.10. domain\_config

The following table outlines the options available under the **[domain\_config]** group in the `/etc/keystone/keystone.conf` file.

Table 9.10. domain\_config

Configuration option = Default value	Type	Description
<b>cache_time = 300</b>	integer value	Time-to-live (TTL, in seconds) to cache domain-specific configuration data. This has no effect unless <b>[domain_config] caching</b> is enabled.
<b>caching = True</b>	boolean value	Toggle for caching of the domain-specific configuration backend. This has no effect unless global caching is enabled. There is normally no reason to disable this.

Configuration option = Default value	Type	Description
<b>driver = sql</b>	string value	Entry point for the domain-specific configuration driver in the <b>keystone.resource.domain_config</b> namespace. Only a <b>sql</b> option is provided by keystone, so there is no reason to set this unless you are providing a custom entry point.

### 9.1.11. endpoint\_filter

The following table outlines the options available under the **[endpoint\_filter]** group in the **/etc/keystone/keystone.conf** file.

Table 9.11. endpoint\_filter

Configuration option = Default value	Type	Description
<b>driver = sql</b>	string value	Entry point for the endpoint filter driver in the <b>keystone.endpoint_filter</b> namespace. Only a <b>sql</b> option is provided by keystone, so there is no reason to set this unless you are providing a custom entry point.
<b>return_all_endpoints_if_no_filter = True</b>	boolean value	This controls keystone's behavior if the configured endpoint filters do not result in any endpoints for a user + project pair (and therefore a potentially empty service catalog). If set to true, keystone will return the entire service catalog. If set to false, keystone will return an empty service catalog.

### 9.1.12. endpoint\_policy

The following table outlines the options available under the **[endpoint\_policy]** group in the **/etc/keystone/keystone.conf** file.

Table 9.12. endpoint\_policy

Configuration option = Default value	Type	Description
<b>driver = sql</b>	string value	Entry point for the endpoint policy driver in the <b>keystone.endpoint_policy</b> namespace. Only a <b>sql</b> driver is provided by keystone, so there is no reason to set this unless you are providing a custom entry point.

### 9.1.13. eventlet\_server

The following table outlines the options available under the **[eventlet\_server]** group in the `/etc/keystone/keystone.conf` file.

Table 9.13. eventlet\_server

Configuration option = Default value	Type	Description
<b>admin_bind_host = 0.0.0.0</b>	unknown value	The IP address of the network interface for the admin service to listen on.
<b>admin_port = 35357</b>	port value	The port number for the admin service to listen on.
<b>public_bind_host = 0.0.0.0</b>	unknown value	The IP address of the network interface for the public service to listen on.
<b>public_port = 5000</b>	port value	The port number for the public service to listen on.

#### 9.1.14. federation

The following table outlines the options available under the **[federation]** group in the `/etc/keystone/keystone.conf` file.

Table 9.14. federation

Configuration option = Default value	Type	Description
<code>`assertion_prefix = `</code>	string value	Prefix to use when filtering environment variable names for federated assertions. Matched variables are passed into the federated mapping engine.
<b>cached = True</b>	boolean value	Toggle for federation caching. This has no effect unless global caching is enabled. There is typically no reason to disable this.
<b>driver = sql</b>	string value	Entry point for the federation backend driver in the <b>keystone.federation</b> namespace. Keystone only provides a <b>sql</b> driver, so there is no reason to set this option unless you are providing a custom entry point.
<b>federated_domain_name = Federated</b>	string value	An arbitrary domain name that is reserved to allow federated ephemeral users to have a domain concept. Note that an admin will not be able to create a domain with this name or update an existing domain to this name. You are not advised to change this value unless you really have to.

Configuration option = Default value	Type	Description
<b>remote_id_attribute = None</b>	string value	Value to be used to obtain the entity ID of the Identity Provider from the environment. For <b>mod_shib</b> , this would be <b>Shib-Identity-Provider</b> . For <b>mod_auth_openidc</b> , this could be <b>HTTP_OIDC_ISS</b> . For <b>mod_auth_mellon</b> , this could be <b>MELLON_IDP</b> .
<b>sso_callback_template = /etc/keystone/sso_callback_template.html</b>	string value	Absolute path to an HTML file used as a Single Sign-On callback handler. This page is expected to redirect the user from keystone back to a trusted dashboard host, by form encoding a token in a POST request. Keystone's default value should be sufficient for most deployments.
<b>trusted_dashboard = []</b>	multi valued	A list of trusted dashboard hosts. Before accepting a Single Sign-On request to return a token, the origin host must be a member of this list. This configuration option may be repeated for multiple values. You must set this in order to use web-based SSO flows. For example: trusted_dashboard=https://acme.example.com/auth/webssso trusted_dashboard=https://beta.example.com/auth/webssso

### 9.1.15. fernet\_tokens

The following table outlines the options available under the **[fernet\_tokens]** group in the **/etc/keystone/keystone.conf** file.

Table 9.15. fernet\_tokens

Configuration option = Default value	Type	Description
---	------	-------------



Configuration option = Default value	Type	Description
<b>key_repository =</b> <b>/etc/keystone/fernet-keys/</b>	string value	Directory containing Fernet token keys. This directory must exist before using <b>keystone-manage fernet_setup</b> for the first time, must be writable by the user running <b>keystone-manage fernet_setup</b> or <b>keystone-manage fernet_rotate</b> , and of course must be readable by keystone's server process. The repository may contain keys in one of three states: a single staged key (always index 0) used for token validation, a single primary key (always the highest index) used for token creation and validation, and any number of secondary keys (all other index values) used for token validation. With multiple keystone nodes, each node must share the same key repository contents, with the exception of the staged key (index 0). It is safe to run <b>keystone-manage fernet_rotate</b> once on any one node to promote a staged key (index 0) to be the new primary (incremented from the previous highest index), and produce a new staged key (a new key with index 0); the resulting repository can then be atomically replicated to other nodes without any risk of race conditions (for example, it is safe to run <b>keystone-manage fernet_rotate</b> on host A, wait any amount of time, create a tarball of the directory on host A, unpack it on host B to a temporary location, and atomically move ( <b>mv</b> ) the directory into place on host B). Running <b>keystone-manage fernet_rotate twice</b> on a key repository without syncing other nodes will result in tokens that can not be validated by all nodes.
<b>max_active_keys = 3</b>	integer value	This controls how many keys are held in rotation by <b>keystone-manage fernet_rotate</b> before they are discarded. The default value of 3 means that keystone will maintain one staged key (always index 0), one primary key (the highest numerical index), and one secondary key (every other index). Increasing this value means that additional secondary keys will be kept in the rotation.

### 9.1.16. healthcheck

The following table outlines the options available under the **[healthcheck]** group in the **/etc/keystone/keystone.conf** file.

Table 9.16. healthcheck

Configuration option = Default value	Type	Description
<b>backends = []</b>	list value	Additional backends that can perform health checks and report that information back as part of a request.
<b>detailed = False</b>	boolean value	Show more detailed information as part of the response
<b>disable_by_file_path = None</b>	string value	Check the presence of a file to determine if an application is running on a port. Used by DisableByFileHealthcheck plugin.
<b>disable_by_file_paths = []</b>	list value	Check the presence of a file based on a port to determine if an application is running on a port. Expects a "port:path" list of strings. Used by DisableByFilesPortsHealthcheck plugin.
<b>path = /healthcheck</b>	string value	The path to respond to healthcheck requests on.

### 9.1.17. identity

The following table outlines the options available under the **[identity]** group in the `/etc/keystone/keystone.conf` file.

Table 9.17. identity

Configuration option = Default value	Type	Description
<b>cache_time = 600</b>	integer value	Time to cache identity data (in seconds). This has no effect unless global and identity caching are enabled.
<b>caching = True</b>	boolean value	Toggle for identity caching. This has no effect unless global caching is enabled. There is typically no reason to disable this.
<b>default_domain_id = default</b>	string value	This references the domain to use for all Identity API v2 requests (which are not aware of domains). A domain with this ID can optionally be created for you by <b>keystone-manage bootstrap</b> . The domain referenced by this ID cannot be deleted on the v3 API, to prevent accidentally breaking the v2 API. There is nothing special about this domain, other than the fact that it must exist to order to maintain support for your v2 clients. There is typically no reason to change this value.

Configuration option = Default value	Type	Description
<b>domain_config_dir = /etc/keystone/domains</b>	string value	Absolute path where keystone should locate domain-specific <b>[identity]</b> configuration files. This option has no effect unless <b>[identity] domain_specific_drivers_enabled</b> is set to true. There is typically no reason to change this value.
<b>domain_configurations_from_database = False</b>	boolean value	By default, domain-specific configuration data is read from files in the directory identified by <b>[identity] domain_config_dir</b> . Enabling this configuration option allows you to instead manage domain-specific configurations through the API, which are then persisted in the backend (typically, a SQL database), rather than using configuration files on disk.
<b>domain_specific_drivers_enabled = False</b>	boolean value	A subset (or all) of domains can have their own identity driver, each with their own partial configuration options, stored in either the resource backend or in a file in a domain configuration directory (depending on the setting of <b>[identity] domain_configurations_from_database</b> ). Only values specific to the domain need to be specified in this manner. This feature is disabled by default, but may be enabled by default in a future release; set to true to enable.
<b>driver = sql</b>	string value	Entry point for the identity backend driver in the <b>keystone.identity</b> namespace. Keystone provides a <b>sql</b> and <b>ldap</b> driver. This option is also used as the default driver selection (along with the other configuration variables in this section) in the event that <b>[identity] domain_specific_drivers_enabled</b> is enabled, but no applicable domain-specific configuration is defined for the domain in question. Unless your deployment primarily relies on <b>ldap</b> AND is not using domain-specific configuration, you should typically leave this set to <b>sql</b> .
<b>list_limit = None</b>	integer value	Maximum number of entities that will be returned in an identity collection.
<b>max_password_length = 4096</b>	integer value	Maximum allowed length for user passwords. Decrease this value to improve performance. Changing this value does not effect existing passwords.
<b>password_hash_algorithm = bcrypt</b>	string value	The password hashing algorithm to use for passwords stored within keystone.

Configuration option = Default value	Type	Description
<b>password_hash_rounds = None</b>	integer value	This option represents a trade off between security and performance. Higher values lead to slower performance, but higher security. Changing this option will only affect newly created passwords as existing password hashes already have a fixed number of rounds applied, so it is safe to tune this option in a running cluster. The default for bcrypt is 12, must be between 4 and 31, inclusive. The default for scrypt is 16, must be within <b>range(1,32)</b> . The default for pbkdf_sha512 is 60000, must be within <b>range(1,1&lt;&lt;32)</b> WARNING: If using scrypt, increasing this value increases BOTH time AND memory requirements to hash a password.
<b>salt_bytesize = None</b>	integer value	Number of bytes to use in scrypt and pbkdf2_sha512 hashing salt. Default for scrypt is 16 bytes. Default for pbkdf2_sha512 is 16 bytes. Limited to a maximum of 96 bytes due to the size of the column used to store password hashes.
<b>scrypt_block_size = None</b>	integer value	Optional block size to pass to scrypt hash function (the <b>r</b> parameter). Useful for tuning scrypt to optimal performance for your CPU architecture. This option is only used when the <b>password_hash_algorithm</b> option is set to <b>scrypt</b> . Defaults to 8.
<b>scrypt_parallelism = None</b>	integer value	Optional parallelism to pass to scrypt hash function (the <b>p</b> parameter). This option is only used when the <b>password_hash_algorithm</b> option is set to <b>scrypt</b> . Defaults to 1.

### 9.1.18. identity\_mapping

The following table outlines the options available under the **[identity\_mapping]** group in the **/etc/keystone/keystone.conf** file.

Table 9.18. identity\_mapping

Configuration option = Default value	Type	Description
---	------	-------------

Configuration option = Default value	Type	Description
<b>backward_compatible_ids = True</b>	boolean value	The format of user and group IDs changed in Juno for backends that do not generate UUIDs (for example, LDAP), with keystone providing a hash mapping to the underlying attribute in LDAP. By default this mapping is disabled, which ensures that existing IDs will not change. Even when the mapping is enabled by using domain-specific drivers ( <b>[identity] domain_specific_drivers_enabled</b> ), any users and groups from the default domain being handled by LDAP will still not be mapped to ensure their IDs remain backward compatible. Setting this value to false will enable the new mapping for all backends, including the default LDAP driver. It is only guaranteed to be safe to enable this option if you do not already have assignments for users and groups from the default LDAP domain, and you consider it to be acceptable for Keystone to provide the different IDs to clients than it did previously (existing IDs in the API will suddenly change). Typically this means that the only time you can set this value to false is when configuring a fresh installation, although that is the recommended value.
<b>driver = sql</b>	string value	Entry point for the identity mapping backend driver in the <b>keystone.identity.id_mapping</b> namespace. Keystone only provides a <b>sql</b> driver, so there is no reason to change this unless you are providing a custom entry point.
<b>generator = sha256</b>	string value	Entry point for the public ID generator for user and group entities in the <b>keystone.identity.id_generator</b> namespace. The Keystone identity mapper only supports generators that produce 64 bytes or less. Keystone only provides a <b>sha256</b> entry point, so there is no reason to change this value unless you're providing a custom entry point.

### 9.1.19. Idap

The following table outlines the options available under the **[Idap]** group in the **/etc/keystone/keystone.conf** file.

Table 9.19. Idap

Configuration option = Default value	Type	Description
<b>alias_dereferencing = default</b>	string value	The LDAP dereferencing option to use for queries involving aliases. A value of <b>default</b> falls back to using default dereferencing behavior configured by your <b>ldap.conf</b> . A value of <b>never</b> prevents aliases from being dereferenced at all. A value of <b>searching</b> dereferences aliases only after name resolution. A value of <b>finding</b> dereferences aliases only during name resolution. A value of <b>always</b> dereferences aliases in all cases.
<b>auth_pool_connection_lifetime = 60</b>	integer value	The maximum end user authentication connection lifetime to the LDAP server in seconds. When this lifetime is exceeded, the connection will be unbound and removed from the connection pool. This option has no effect unless <b>[ldap] use_auth_pool</b> is also enabled.
<b>auth_pool_size = 100</b>	integer value	The size of the connection pool to use for end user authentication. This option has no effect unless <b>[ldap] use_auth_pool</b> is also enabled.
<b>chase_referrals = None</b>	boolean value	Sets keystone's referral chasing behavior across directory partitions. If left unset, the system's default behavior will be used.
<b>connection_timeout = -1</b>	integer value	The connection timeout to use with the LDAP server. A value of <b>-1</b> means that connections will never timeout.
<b>debug_level = None</b>	integer value	Sets the LDAP debugging level for LDAP calls. A value of 0 means that debugging is not enabled. This value is a bitmask, consult your LDAP documentation for possible values.
<b>group_ad_nesting = False</b>	boolean value	If enabled, group queries will use Active Directory specific filters for nested groups.
<b>group_additional_attribute_mapping = []</b>	list value	A list of LDAP attribute to keystone group attribute pairs used for mapping additional attributes to groups in keystone. The expected format is <b>&lt;ldap_attr&gt;:&lt;group_attr&gt;</b> , where <b>ldap_attr</b> is the attribute in the LDAP object and <b>group_attr</b> is the attribute which should appear in the identity API.
<b>group_attribute_ignore = []</b>	list value	List of group attributes to ignore on create and update. or whether a specific group attribute should be filtered for list or show group.

Configuration option = Default value	Type	Description
<b>group_desc_attribute = description</b>	string value	The LDAP attribute mapped to group descriptions in keystone.
<b>group_filter = None</b>	string value	The LDAP search filter to use for groups.
<b>group_id_attribute = cn</b>	string value	The LDAP attribute mapped to group IDs in keystone. This must NOT be a multivalued attribute. Group IDs are expected to be globally unique across keystone domains and URL-safe.
<b>group_member_attribute = member</b>	string value	The LDAP attribute used to indicate that a user is a member of the group.
<b>group_members_are_ids = False</b>	boolean value	Enable this option if the members of the group object class are keystone user IDs rather than LDAP DNs. This is the case when using <b>posixGroup</b> as the group object class in Open Directory.
<b>group_name_attribute = ou</b>	string value	The LDAP attribute mapped to group names in keystone. Group names are expected to be unique only within a keystone domain and are not expected to be URL-safe.
<b>group_objectclass = groupOfNames</b>	string value	The LDAP object class to use for groups. If setting this option to <b>posixGroup</b> , you may also be interested in enabling the <b>[ldap] group_members_are_ids</b> option.
<b>group_tree_dn = None</b>	string value	The search base to use for groups. Defaults to the <b>[ldap] suffix</b> value.
<b>page_size = 0</b>	integer value	Defines the maximum number of results per page that keystone should request from the LDAP server when listing objects. A value of zero ( <b>0</b> ) disables paging.
<b>password = None</b>	string value	The password of the administrator bind DN to use when querying the LDAP server, if your LDAP server requires it.
<b>pool_connection_lifetime = 600</b>	integer value	The maximum connection lifetime to the LDAP server in seconds. When this lifetime is exceeded, the connection will be unbound and removed from the connection pool. This option has no effect unless <b>[ldap] use_pool</b> is also enabled.

Configuration option = Default value	Type	Description
<b>pool_connection_timeout = -1</b>	integer value	The connection timeout to use when pooling LDAP connections. A value of <b>-1</b> means that connections will never timeout. This option has no effect unless <b>[ldap] use_pool</b> is also enabled.
<b>pool_retry_delay = 0.1</b>	floating point value	The number of seconds to wait before attempting to reconnect to the LDAP server. This option has no effect unless <b>[ldap] use_pool</b> is also enabled.
<b>pool_retry_max = 3</b>	integer value	The maximum number of times to attempt reconnecting to the LDAP server before aborting. A value of zero prevents retries. This option has no effect unless <b>[ldap] use_pool</b> is also enabled.
<b>pool_size = 10</b>	integer value	The size of the LDAP connection pool. This option has no effect unless <b>[ldap] use_pool</b> is also enabled.
<b>query_scope = one</b>	string value	The search scope which defines how deep to search within the search base. A value of <b>one</b> (representing <b>oneLevel</b> or <b>singleLevel</b> ) indicates a search of objects immediately below to the base object, but does not include the base object itself. A value of <b>sub</b> (representing <b>subtree</b> or <b>wholeSubtree</b> ) indicates a search of both the base object itself and the entire subtree below it.
<b>suffix = cn=example,cn=com</b>	string value	The default LDAP server suffix to use, if a DN is not defined via either <b>[ldap] user_tree_dn</b> or <b>[ldap] group_tree_dn</b> .
<b>tls_cacertdir = None</b>	string value	An absolute path to a CA certificate directory to use when communicating with LDAP servers. There is no reason to set this option if you've also set <b>[ldap] tls_cacertfile</b> .
<b>tls_cacertfile = None</b>	string value	An absolute path to a CA certificate file to use when communicating with LDAP servers. This option will take precedence over <b>[ldap] tls_cacertdir</b> , so there is no reason to set both.
<b>tls_req_cert = demand</b>	string value	Specifies which checks to perform against client certificates on incoming TLS sessions. If set to <b>demand</b> , then a certificate will always be requested and required from the LDAP server. If set to <b>allow</b> , then a certificate will always be requested but not required from the LDAP server. If set to <b>never</b> , then a certificate will never be requested.



Configuration option = Default value	Type	Description
<b>url = ldap://localhost</b>	string value	URL(s) for connecting to the LDAP server. Multiple LDAP URLs may be specified as a comma separated string. The first URL to successfully bind is used for the connection.
<b>use_auth_pool = True</b>	boolean value	Enable LDAP connection pooling for end user authentication. There is typically no reason to disable this.
<b>use_pool = True</b>	boolean value	Enable LDAP connection pooling for queries to the LDAP server. There is typically no reason to disable this.
<b>use_tls = False</b>	boolean value	Enable TLS when communicating with LDAP servers. You should also set the <b>[ldap] tls_cacertfile</b> and <b>[ldap] tls_cacertdir</b> options when using this option. Do not set this option if you are using LDAP over SSL (LDAPS) instead of TLS.
<b>user = None</b>	string value	The user name of the administrator bind DN to use when querying the LDAP server, if your LDAP server requires it.
<b>user_additional_attribute_mapping = []</b>	list value	A list of LDAP attribute to keystone user attribute pairs used for mapping additional attributes to users in keystone. The expected format is <b>&lt;ldap_attr&gt;: &lt;user_attr&gt;</b> , where <b>ldap_attr</b> is the attribute in the LDAP object and <b>user_attr</b> is the attribute which should appear in the identity API.
<b>user_attribute_ignore = [u'default_project_id']</b>	list value	List of user attributes to ignore on create and update, or whether a specific user attribute should be filtered for list or show user.
<b>user_default_project_id_attribute = None</b>	string value	The LDAP attribute mapped to a user's default_project_id in keystone. This is most commonly used when keystone has write access to LDAP.
<b>user_description_attribute = description</b>	string value	The LDAP attribute mapped to user descriptions in keystone.
<b>user_enabled_attribute = enabled</b>	string value	The LDAP attribute mapped to the user enabled attribute in keystone. If setting this option to <b>userAccountControl</b> , then you may be interested in setting <b>[ldap] user_enabled_mask</b> and <b>[ldap] user_enabled_default</b> as well.

Configuration option = Default value	Type	Description
<b>user_enabled_default = True</b>	string value	The default value to enable users. This should match an appropriate integer value if the LDAP server uses non-boolean (bitmask) values to indicate if a user is enabled or disabled. If this is not set to <b>True</b> , then the typical value is <b>512</b> . This is typically used when <b>[ldap] user_enabled_attribute = userAccountControl</b> .
<b>user_enabled_emulation = False</b>	boolean value	If enabled, keystone uses an alternative method to determine if a user is enabled or not by checking if they are a member of the group defined by the <b>[ldap] user_enabled_emulation_dn</b> option. Enabling this option causes keystone to ignore the value of <b>[ldap] user_enabled_invert</b> .
<b>user_enabled_emulation_dn = None</b>	string value	DN of the group entry to hold enabled users when using enabled emulation. Setting this option has no effect unless <b>[ldap] user_enabled_emulation</b> is also enabled.
<b>user_enabled_emulation_use_group_config = False</b>	boolean value	Use the <b>[ldap] group_member_attribute</b> and <b>[ldap] group_objectclass</b> settings to determine membership in the emulated enabled group. Enabling this option has no effect unless <b>[ldap] user_enabled_emulation</b> is also enabled.
<b>user_enabled_invert = False</b>	boolean value	Logically negate the boolean value of the enabled attribute obtained from the LDAP server. Some LDAP servers use a boolean lock attribute where "true" means an account is disabled. Setting <b>[ldap] user_enabled_invert = true</b> will allow these lock attributes to be used. This option will have no effect if either the <b>[ldap] user_enabled_mask</b> or <b>[ldap] user_enabled_emulation</b> options are in use.
<b>user_enabled_mask = 0</b>	integer value	Bitmask integer to select which bit indicates the enabled value if the LDAP server represents "enabled" as a bit on an integer rather than as a discrete boolean. A value of <b>0</b> indicates that the mask is not used. If this is not set to <b>0</b> the typical value is <b>2</b> . This is typically used when <b>[ldap] user_enabled_attribute = userAccountControl</b> . Setting this option causes keystone to ignore the value of <b>[ldap] user_enabled_invert</b> .
<b>user_filter = None</b>	string value	The LDAP search filter to use for users.

Configuration option = Default value	Type	Description
<b>user_id_attribute = cn</b>	string value	The LDAP attribute mapped to user IDs in keystone. This must NOT be a multivalued attribute. User IDs are expected to be globally unique across keystone domains and URL-safe.
<b>user_mail_attribute = mail</b>	string value	The LDAP attribute mapped to user emails in keystone.
<b>user_name_attribute = sn</b>	string value	The LDAP attribute mapped to user names in keystone. User names are expected to be unique only within a keystone domain and are not expected to be URL-safe.
<b>user_objectclass = inetOrgPerson</b>	string value	The LDAP object class to use for users.
<b>user_pass_attribute = userPassword</b>	string value	The LDAP attribute mapped to user passwords in keystone.
<b>user_tree_dn = None</b>	string value	The search base to use for users. Defaults to the <b>[ldap] suffix</b> value.

### 9.1.20. matchmaker\_redis

The following table outlines the options available under the **[matchmaker\_redis]** group in the `/etc/keystone/keystone.conf` file.

Table 9.20. matchmaker\_redis

Configuration option = Default value	Type	Description
<b>check_timeout = 20000</b>	integer value	Time in ms to wait before the transaction is killed.
<b>host = 127.0.0.1</b>	string value	Host to locate redis.
<code>`password = `</code>	string value	Password for Redis server (optional).
<b>port = 6379</b>	port value	Use this port to connect to redis host.
<b>sentinel_group_name = oslo-messaging-zeromq</b>	string value	Redis replica set name.
<b>sentinel_hosts = []</b>	list value	List of Redis Sentinel hosts (fault tolerance mode), e.g., [host:port, host1:port ... ]

Configuration option = Default value	Type	Description
<b>socket_timeout = 10000</b>	integer value	Timeout in ms on blocking socket operations.
<b>wait_timeout = 2000</b>	integer value	Time in ms to wait between connection attempts.

### 9.1.21. memcache

The following table outlines the options available under the **[memcache]** group in the `/etc/keystone/keystone.conf` file.

Table 9.21. memcache

Configuration option = Default value	Type	Description
<b>dead_retry = 300</b>	integer value	Number of seconds memcached server is considered dead before it is tried again. This is used by the key value store system.
<b>pool_connection_get_timeout = 10</b>	integer value	Number of seconds that an operation will wait to get a memcache client connection. This is used by the key value store system.
<b>pool_maxsize = 10</b>	integer value	Max total number of open connections to every memcached server. This is used by the key value store system.
<b>pool_unused_timeout = 60</b>	integer value	Number of seconds a connection to memcached is held unused in the pool before it is closed. This is used by the key value store system.
<b>socket_timeout = 3</b>	integer value	Timeout in seconds for every call to a server. This is used by the key value store system.

### 9.1.22. oauth1

The following table outlines the options available under the **[oauth1]** group in the `/etc/keystone/keystone.conf` file.

Table 9.22. oauth1

Configuration option = Default value	Type	Description
---	------	-------------

Configuration option = Default value	Type	Description
<b>access_token_duration = 86400</b>	integer value	Number of seconds for the OAuth Access Token to remain valid after being created. This is the amount of time the consumer has to interact with the service provider (which is typically keystone). Setting this option to zero means that access tokens will last forever.
<b>driver = sql</b>	string value	Entry point for the OAuth backend driver in the <b>keystone.oauth1</b> namespace. Typically, there is no reason to set this option unless you are providing a custom entry point.
<b>request_token_duration = 28800</b>	integer value	Number of seconds for the OAuth Request Token to remain valid after being created. This is the amount of time the user has to authorize the token. Setting this option to zero means that request tokens will last forever.

### 9.1.23. oslo\_messaging\_amqp

The following table outlines the options available under the **[oslo\_messaging\_amqp]** group in the `/etc/keystone/keystone.conf` file.

Table 9.23. oslo\_messaging\_amqp

Configuration option = Default value	Type	Description
<b>addressing_mode = dynamic</b>	string value	Indicates the addressing mode used by the driver. Permitted values: <i>legacy</i> - use legacy non-routable addressing <i>routable</i> - use routable addresses <i>dynamic</i> - use legacy addresses if the message bus does not support routing otherwise use routable addressing
<b>allow_insecure_clients = False</b>	boolean value	Accept clients using either SSL or plain TCP
<b>anycast_address = anycast</b>	string value	Appended to the address prefix when sending to a group of consumers. Used by the message bus to identify messages that should be delivered in a round-robin fashion across consumers.
<b>broadcast_prefix = broadcast</b>	string value	address prefix used when broadcasting to all servers

Configuration option = Default value	Type	Description
<b>connection_retry_backoff = 2</b>	integer value	Increase the <code>connection_retry_interval</code> by this many seconds after each unsuccessful failover attempt.
<b>connection_retry_interval = 1</b>	integer value	Seconds to pause before attempting to re-connect.
<b>connection_retry_interval_max = 30</b>	integer value	Maximum limit for <code>connection_retry_interval</code> + <code>connection_retry_backoff</code>
<b>container_name = None</b>	string value	Name for the AMQP container. must be globally unique. Defaults to a generated UUID
<b>default_notification_exchange = None</b>	string value	Exchange name used in notification addresses. Exchange name resolution precedence: Target.exchange if set else <code>default_notification_exchange</code> if set else <code>control_exchange</code> if set else <i>notify</i>
<b>default_notify_timeout = 30</b>	integer value	The deadline for a sent notification message delivery. Only used when caller does not provide a timeout expiry.
<b>default_reply_retry = 0</b>	integer value	The maximum number of attempts to re-send a reply message which failed due to a recoverable error.
<b>default_reply_timeout = 30</b>	integer value	The deadline for an rpc reply message delivery.
<b>default_rpc_exchange = None</b>	string value	Exchange name used in RPC addresses. Exchange name resolution precedence: Target.exchange if set else <code>default_rpc_exchange</code> if set else <code>control_exchange</code> if set else <i>rpc</i>
<b>default_send_timeout = 30</b>	integer value	The deadline for an rpc cast or call message delivery. Only used when caller does not provide a timeout expiry.
<b>default_sender_link_timeout = 600</b>	integer value	The duration to schedule a purge of idle sender links. Detach link after expiry.
<b>group_request_prefix = unicast</b>	string value	address prefix when sending to any server in group
<b>idle_timeout = 0</b>	integer value	Timeout for inactive connections (in seconds)

Configuration option = Default value	Type	Description
<b>link_retry_delay = 10</b>	integer value	Time to pause between re-connecting an AMQP 1.0 link that failed due to a recoverable error.
<b>multicast_address = multicast</b>	string value	Appended to the address prefix when sending a fanout message. Used by the message bus to identify fanout messages.
<b>notify_address_prefix = openstack.org/om/notify</b>	string value	Address prefix for all generated Notification addresses
<b>notify_server_credit = 100</b>	integer value	Window size for incoming Notification messages
<code>`password = `</code>	string value	Password for message broker authentication
<b>pre_settled = [u'rpc-cast', u'rpc-reply']</b>	multi valued	Send messages of this type pre-settled. Pre-settled messages will not receive acknowledgement from the peer. Note well: pre-settled messages may be silently discarded if the delivery fails. Permitted values: <i>rpc-call</i> - send RPC Calls pre-settled <i>rpc-reply</i> - send RPC Replies pre-settled <i>rpc-cast</i> - Send RPC Casts pre-settled <i>notify</i> - Send Notifications pre-settled
<b>pseudo_vhost = True</b>	boolean value	Enable virtual host support for those message buses that do not natively support virtual hosting (such as <i>qpidd</i> ). When set to true the virtual host name will be added to all message bus addresses, effectively creating a private <i>subnet</i> per virtual host. Set to False if the message bus supports virtual hosting using the <i>hostname</i> field in the AMQP 1.0 Open performative as the name of the virtual host.
<b>reply_link_credit = 200</b>	integer value	Window size for incoming RPC Reply messages.
<b>rpc_address_prefix = openstack.org/om/rpc</b>	string value	Address prefix for all generated RPC addresses
<b>rpc_server_credit = 100</b>	integer value	Window size for incoming RPC Request messages
<code>`sas_config_dir = `</code>	string value	Path to directory that contains the SASL configuration
<code>`sas_config_name = `</code>	string value	Name of configuration file (without <i>.conf</i> suffix)
<code>`sas_default_realm = `</code>	string value	SASL realm to use if no realm present in username

Configuration option = Default value	Type	Description
<code>`sasl_mechanisms = `</code>	string value	Space separated list of acceptable SASL mechanisms
<b>server_request_prefix = exclusive</b>	string value	address prefix used when sending to a specific server
<b>ssl = False</b>	boolean value	Attempt to connect via SSL. If no other ssl-related parameters are given, it will use the system's CA-bundle to verify the server's certificate.
<code>`ssl_ca_file = `</code>	string value	CA certificate PEM file used to verify the server's certificate
<code>`ssl_cert_file = `</code>	string value	Self-identifying certificate PEM file for client authentication
<code>`ssl_key_file = `</code>	string value	Private key PEM file used to sign <code>ssl_cert_file</code> certificate (optional)
<b>ssl_key_password = None</b>	string value	Password for decrypting <code>ssl_key_file</code> (if encrypted)
<b>ssl_verify_vhost = False</b>	boolean value	By default SSL checks that the name in the server's certificate matches the hostname in the <code>transport_url</code> . In some configurations it may be preferable to use the virtual hostname instead, for example if the server uses the Server Name Indication TLS extension (rfc6066) to provide a certificate per virtual host. Set <code>ssl_verify_vhost</code> to True if the server's SSL certificate uses the virtual host name instead of the DNS name.
<b>trace = False</b>	boolean value	Debug: dump AMQP frames to stdout
<b>unicast_address = unicast</b>	string value	Appended to the address prefix when sending to a particular RPC/Notification server. Used by the message bus to identify messages sent to a single destination.
<code>`username = `</code>	string value	User name for message broker authentication

### 9.1.24. oslo\_messaging\_kafka

The following table outlines the options available under the **[oslo\_messaging\_kafka]** group in the `/etc/keystone/keystone.conf` file.

Table 9.24. oslo\_messaging\_kafka



Configuration option = Default value	Type	Description
<b>conn_pool_min_size = 2</b>	integer value	The pool size limit for connections expiration policy
<b>conn_pool_ttl = 1200</b>	integer value	The time-to-live in sec of idle connections in the pool
<b>consumer_group = oslo_messaging_consumer</b>	string value	Group id for Kafka consumer. Consumers in one group will coordinate message consumption
<b>kafka_consumer_timeout = 1.0</b>	floating point value	Default timeout(s) for Kafka consumers
<b>kafka_default_host = localhost</b>	string value	Default Kafka broker Host
<b>kafka_default_port = 9092</b>	port value	Default Kafka broker Port
<b>kafka_max_fetch_bytes = 1048576</b>	integer value	Max fetch bytes of Kafka consumer
<b>pool_size = 10</b>	integer value	Pool Size for Kafka Consumers
<b>producer_batch_size = 16384</b>	integer value	Size of batch for the producer async send
<b>producer_batch_timeout = 0.0</b>	floating point value	Upper bound on the delay for KafkaProducer batching in seconds

### 9.1.25. oslo\_messaging\_notifications

The following table outlines the options available under the **[oslo\_messaging\_notifications]** group in the `/etc/keystone/keystone.conf` file.

Table 9.25. oslo\_messaging\_notifications

Configuration option = Default value	Type	Description
<b>driver = []</b>	multi valued	The Drivers(s) to handle sending notifications. Possible values are messaging, messagingv2, routing, log, test, noop
<b>retry = -1</b>	integer value	The maximum number of attempts to re-send a notification message which failed to be delivered due to a recoverable error. 0 - No retry, -1 - indefinite
<b>topics = [u'notifications']</b>	list value	AMQP topic used for OpenStack notifications.

Configuration option = Default value	Type	Description
<b>transport_url = None</b>	string value	A URL representing the messaging driver to use for notifications. If not set, we fall back to the same configuration used for RPC.

### 9.1.26. oslo\_messaging\_rabbit

The following table outlines the options available under the **[oslo\_messaging\_rabbit]** group in the **/etc/keystone/keystone.conf** file.

Table 9.26. oslo\_messaging\_rabbit

Configuration option = Default value	Type	Description
<b>amqp_auto_delete = False</b>	boolean value	Auto-delete queues in AMQP.
<b>amqp_durable_queues = False</b>	boolean value	Use durable queues in AMQP.
<b>channel_max = None</b>	integer value	Maximum number of channels to allow
<b>connection_factory = single</b>	string value	Connection factory implementation
<b>default_notification_exchange = \${control_exchange}_notification</b>	string value	Exchange name for sending notifications
<b>default_notification_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending notification, -1 means infinite retry.
<b>default_rpc_exchange = \${control_exchange}_rpc</b>	string value	Exchange name for sending RPC messages
<b>default_rpc_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending RPC message, -1 means infinite retry. If actual retry attempts in not 0 the rpc request could be processed more than one time
<b>default_serializer_type = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>fake_rabbit = False</b>	boolean value	Deprecated, use <code>rpc_backend=kombu+memory</code> or <code>rpc_backend=fake</code>

Configuration option = Default value	Type	Description
<b>frame_max = None</b>	integer value	The maximum byte size for an AMQP frame
<b>heartbeat_interval = 3</b>	integer value	How often to send heartbeats for consumer's connections
<b>heartbeat_rate = 2</b>	integer value	How often times during the heartbeat_timeout_threshold we check the heartbeat.
<b>heartbeat_timeout_threshold = 60</b>	integer value	Number of seconds after which the Rabbit broker is considered down if heartbeat's keep-alive fails (0 disable the heartbeat). EXPERIMENTAL
<b>host_connection_reconnect_delay = 0.25</b>	floating point value	Set delay for reconnection to some host which has connection error
<b>kombu_compression = None</b>	string value	EXPERIMENTAL: Possible values are: gzip, bz2. If not set compression will not be used. This option may not be available in future versions.
<b>kombu_failover_strategy = round-robin</b>	string value	Determines how the next RabbitMQ node is chosen in case the one we are currently connected to becomes unavailable. Takes effect only if more than one RabbitMQ node is provided in config.
<b>kombu_missing_consumer_retry_timeout = 60</b>	integer value	How long to wait a missing client before abandoning to send it its replies. This value should not be longer than rpc_response_timeout.
<b>kombu_reconnect_delay = 1.0</b>	floating point value	How long to wait before reconnecting in response to an AMQP consumer cancel notification.
<b>notification_listener_prefetch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to notification listener.
<b>notification_persistence = False</b>	boolean value	Persist notification messages.
<b>notification_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending notification message
<b>pool_max_overflow = 0</b>	integer value	Maximum number of connections to create above <b>pool_max_size</b> .
<b>pool_max_size = 30</b>	integer value	Maximum number of connections to keep queued.

Configuration option = Default value	Type	Description
<b>pool_recycle = 600</b>	integer value	Lifetime of a connection (since creation) in seconds or None for no recycling. Expired connections are closed on acquire.
<b>pool_stale = 60</b>	integer value	Threshold at which inactive (since release) connections are considered stale in seconds or None for no staleness. Stale connections are closed on acquire.
<b>pool_timeout = 30</b>	integer value	Default number of seconds to wait for a connections to available
<b>rabbit_ha_queues = False</b>	boolean value	Try to use HA queues in RabbitMQ (x-ha-policy: all). If you change this option, you must wipe the RabbitMQ database. In RabbitMQ 3.0, queue mirroring is no longer controlled by the x-ha-policy argument when declaring a queue. If you just want to make sure that all queues (except those with auto-generated names) are mirrored across all nodes, run: <code>"rabbitmqctl set_policy HA ^{?!amq\}.*{ha-mode": "all"}"</code>
<b>rabbit_host = localhost</b>	string value	The RabbitMQ broker address where a single node is used.
<b>rabbit_hosts = [u'\$rabbit_host:\$rabbit_port']</b>	list value	RabbitMQ HA cluster host:port pairs.
<b>rabbit_interval_max = 30</b>	integer value	Maximum interval of RabbitMQ connection retries. Default is 30 seconds.
<b>rabbit_login_method = AMQPLAIN</b>	string value	The RabbitMQ login method.
<b>rabbit_max_retries = 0</b>	integer value	Maximum number of RabbitMQ connection retries. Default is 0 (infinite retry count).
<b>rabbit_password = guest</b>	string value	The RabbitMQ password.
<b>rabbit_port = 5672</b>	port value	The RabbitMQ broker port where a single node is used.
<b>rabbit_qos_prefetch_count = 0</b>	integer value	Specifies the number of messages to prefetch. Setting to zero allows unlimited messages.

Configuration option = Default value	Type	Description
<b>rabbit_retry_backoff = 2</b>	integer value	How long to backoff for between retries when connecting to RabbitMQ.
<b>rabbit_retry_interval = 1</b>	integer value	How frequently to retry connecting with RabbitMQ.
<b>rabbit_transient_queues_ttl = 1800</b>	integer value	Positive integer representing duration in seconds for queue TTL (x-expires). Queues which are unused for the duration of the TTL are automatically deleted. The parameter affects only reply and fanout queues.
<b>rabbit_userid = guest</b>	string value	The RabbitMQ userid.
<b>rabbit_virtual_host = /</b>	string value	The RabbitMQ virtual host.
<b>rpc_listener_prefetch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to rpc listener.
<b>rpc_queue_expiration = 60</b>	integer value	Time to live for rpc queues without consumers in seconds.
<b>rpc_reply_exchange = \${control_exchange}_rpc_reply</b>	string value	Exchange name for receiving RPC replies
<b>rpc_reply_listener_prefetch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to rpc reply listener.
<b>rpc_reply_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending reply. -1 means infinite retry during rpc_timeout
<b>rpc_reply_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending reply.
<b>rpc_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending RPC message
<b>socket_timeout = 0.25</b>	floating point value	Set socket timeout in seconds for connection's socket
<b>ssl = None</b>	boolean value	Enable SSL
<b>ssl_ca_file = `</b>	string value	SSL certification authority file (valid only if SSL enabled).

Configuration option = Default value	Type	Description
<code>`ssl_cert_file = `</code>	string value	SSL cert file (valid only if SSL enabled).
<code>`ssl_key_file = `</code>	string value	SSL key file (valid only if SSL enabled).
<b>ssl_options = None</b>	dict value	Arguments passed to <code>ssl.wrap_socket</code>
<code>`ssl_version = `</code>	string value	SSL version to use (valid only if SSL enabled). Valid values are TLSv1 and SSLv23. SSLv2, SSLv3, TLSv1_1, and TLSv1_2 may be available on some distributions.
<b>tcp_user_timeout = 0.25</b>	floating point value	Set <code>TCP_USER_TIMEOUT</code> in seconds for connection's socket

### 9.1.27. oslo\_messaging\_zmq

The following table outlines the options available under the `[oslo_messaging_zmq]` group in the `/etc/keystone/keystone.conf` file.

Table 9.27. oslo\_messaging\_zmq

Configuration option = Default value	Type	Description
<b>rpc_ack_timeout_base = 15</b>	integer value	Number of seconds to wait for an ack from a cast/call. After each retry attempt this timeout is multiplied by some specified multiplier.
<b>rpc_ack_timeout_multiplier = 2</b>	integer value	Number to multiply base ack timeout by after each retry attempt.
<b>rpc_message_ttl = 300</b>	integer value	Expiration timeout in seconds of a sent/received message after which it is not tracked anymore by a client/server.
<b>rpc_poll_timeout = 1</b>	integer value	The default number of seconds that poll should wait. Poll raises timeout exception when timeout expired.
<b>rpc_retry_attempts = 3</b>	integer value	Default number of message sending attempts in case of any problems occurred: positive value N means at most N retries, 0 means no retries, None or -1 (or any other negative values) mean to retry forever. This option is used only if acknowledgments are enabled.
<b>rpc_thread_pool_size = 100</b>	integer value	Maximum number of (green) threads to work concurrently.

Configuration option = Default value	Type	Description
<b>rpc_use_acks = False</b>	boolean value	Wait for message acknowledgements from receivers. This mechanism works only via proxy without PUB/SUB.
<b>rpc_zmq_bind_address = *</b>	string value	ZeroMQ bind address. Should be a wildcard (*), an ethernet interface, or IP. The "host" option should point or resolve to this address.
<b>rpc_zmq_bind_port_retries = 100</b>	integer value	Number of retries to find free port number before fail with ZMQBindError.
<b>rpc_zmq_contexts = 1</b>	integer value	Number of ZeroMQ contexts, defaults to 1.
<b>rpc_zmq_host = ccosp-undercloud.usersys.redhat.com</b>	string value	Name of this node. Must be a valid hostname, FQDN, or IP address. Must match "host" option, if running Nova.
<b>rpc_zmq_ipc_dir = /var/run/openstack</b>	string value	Directory for holding IPC sockets.
<b>rpc_zmq_matchmaker = redis</b>	string value	MatchMaker driver.
<b>rpc_zmq_max_port = 65536</b>	integer value	Maximal port number for random ports range.
<b>rpc_zmq_min_port = 49153</b>	port value	Minimal port number for random ports range.
<b>rpc_zmq_serialization = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>rpc_zmq_topic_backlog = None</b>	integer value	Maximum number of ingress messages to locally buffer per topic. Default is unlimited.
<b>subscribe_on = []</b>	list value	List of publisher hosts SubConsumer can subscribe on. This option has higher priority then the default publishers list taken from the matchmaker.
<b>use_dynamic_connections = False</b>	boolean value	This option makes direct connections dynamic or static. It makes sense only with use_router_proxy=False which means to use direct connections for direct message types (ignored otherwise).

Configuration option = Default value	Type	Description
<b>use_pub_sub = False</b>	boolean value	Use PUB/SUB pattern for fanout methods. PUB/SUB always uses proxy.
<b>use_router_proxy = False</b>	boolean value	Use ROUTER remote proxy.
<b>zmq_failover_connections = 2</b>	integer value	How many additional connections to a host will be made for failover reasons. This option is actual only in dynamic connections mode.
<b>zmq_immediate = True</b>	boolean value	This option configures round-robin mode in zmq socket. True means not keeping a queue when server side disconnects. False means to keep queue and messages even if server is disconnected, when the server appears we send all accumulated messages to it.
<b>zmq_linger = -1</b>	integer value	Number of seconds to wait before all pending messages will be sent after closing a socket. The default value of -1 specifies an infinite linger period. The value of 0 specifies no linger period. Pending messages shall be discarded immediately when the socket is closed. Positive values specify an upper bound for the linger period.
<b>zmq_target_expire = 300</b>	integer value	Expiration timeout in seconds of a name service record about existing target (< 0 means no timeout).
<b>zmq_target_update = 180</b>	integer value	Update period in seconds of a name service record about existing target.
<b>zmq_tcp_keepalive = -1</b>	integer value	Enable/disable TCP keepalive (KA) mechanism. The default value of -1 (or any other negative value) means to skip any overrides and leave it to OS default; 0 and 1 (or any other positive value) mean to disable and enable the option respectively.
<b>zmq_tcp_keepalive_cnt = -1</b>	integer value	The number of retransmissions to be carried out before declaring that remote end is not available. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.



Configuration option = Default value	Type	Description
<b>zmq_tcp_keepalive_idle = -1</b>	integer value	The duration between two keepalive transmissions in idle condition. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_intvl = -1</b>	integer value	The duration between two successive keepalive retransmissions, if acknowledgement to the previous keepalive transmission is not received. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.

### 9.1.28. oslo\_middleware

The following table outlines the options available under the **[oslo\_middleware]** group in the **/etc/keystone/keystone.conf** file.

Table 9.28. oslo\_middleware

Configuration option = Default value	Type	Description
<b>enable_proxy_headers_parsing = False</b>	boolean value	Whether the application is behind a proxy or not. This determines if the middleware should parse the headers or not.
<b>max_request_body_size = 114688</b>	integer value	The maximum body size for each request, in bytes.
<b>secure_proxy_ssl_header = X-Forwarded-Proto</b>	string value	The HTTP Header that will be used to determine what the original request protocol scheme was, even if it was hidden by a SSL termination proxy.

### 9.1.29. oslo\_policy

The following table outlines the options available under the **[oslo\_policy]** group in the **/etc/keystone/keystone.conf** file.

Table 9.29. oslo\_policy

Configuration option = Default value	Type	Description
---	------	-------------

Configuration option = Default value	Type	Description
<b>enforce_scope = False</b>	boolean value	This option controls whether or not to enforce scope when evaluating policies. If <b>True</b> , the scope of the token used in the request is compared to the <b>scope_types</b> of the policy being enforced. If the scopes do not match, an <b>InvalidScope</b> exception will be raised. If <b>False</b> , a message will be logged informing operators that policies are being invoked with mismatching scope.
<b>policy_default_rule = default</b>	string value	Default rule. Enforced when a requested rule is not found.
<b>policy_dirs = [u'policy.d']</b>	multi valued	Directories where policy configuration files are stored. They can be relative to any directory in the search path defined by the config_dir option, or absolute paths. The file defined by policy_file must exist for these directories to be searched. Missing or empty directories are ignored.
<b>policy_file = policy.json</b>	string value	The file that defines policies.
<b>remote_content_type = application/x-www-form-urlencoded</b>	string value	Content Type to send and receive data for REST based policy check
<b>remote_ssl_ca_cert_file = None</b>	string value	Absolute path to ca cert file for REST based policy check
<b>remote_ssl_client_cert_file = None</b>	string value	Absolute path to client cert for REST based policy check
<b>remote_ssl_client_key_file = None</b>	string value	Absolute path client key file REST based policy check
<b>remote_ssl_verify_server_cert = False</b>	boolean value	server identity verification for REST based policy check

### 9.1.30. paste\_deploy

The following table outlines the options available under the **[paste\_deploy]** group in the **/etc/keystone/keystone.conf** file.

Table 9.30. paste\_deploy

Configuration option = Default value	Type	Description
<b>config_file = keystone-paste.ini</b>	string value	Name of (or absolute path to) the Paste Deploy configuration file that composes middleware and the keystone application itself into actual WSGI entry points. See <a href="http://pythonpaste.org/deploy/">http://pythonpaste.org/deploy/</a> for additional documentation on the file's format.

### 9.1.31. policy

The following table outlines the options available under the **[policy]** group in the `/etc/keystone/keystone.conf` file.

Table 9.31. policy

Configuration option = Default value	Type	Description
<b>driver = sql</b>	string value	Entry point for the policy backend driver in the <b>keystone.policy</b> namespace. Supplied drivers are <b>rules</b> (which does not support any CRUD operations for the v3 policy API) and <b>sql</b> . Typically, there is no reason to set this option unless you are providing a custom entry point.
<b>list_limit = None</b>	integer value	Maximum number of entities that will be returned in a policy collection.

### 9.1.32. profiler

The following table outlines the options available under the **[profiler]** group in the `/etc/keystone/keystone.conf` file.

Table 9.32. profiler

Configuration option = Default value	Type	Description
---	------	-------------

Configuration option = Default value	Type	Description
<b>connection_string = messaging://</b>	string value	<p>Connection string for a notifier backend. Default value is messaging:// which sets the notifier to oslo_messaging.</p> <p>Examples of possible values:</p> <ul style="list-style-type: none"> <li>● messaging://: use oslo_messaging driver for sending notifications.</li> <li>● mongodb://127.0.0.1:27017 : use mongoddb driver for sending notifications.</li> <li>● elasticsearch://127.0.0.1:9200 : use elasticsearch driver for sending notifications.</li> </ul>
<b>enabled = False</b>	boolean value	<p>Enables the profiling for all services on this node. Default value is False (fully disable the profiling feature).</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True: Enables the feature</li> <li>● False: Disables the feature. The profiling cannot be started via this project operations. If the profiling is triggered by another project, this project part will be empty.</li> </ul>
<b>es_doc_type = notification</b>	string value	Document type for notification indexing in elasticsearch.
<b>es_scroll_size = 10000</b>	integer value	Elasticsearch splits large requests in batches. This parameter defines maximum size of each batch (for example: es_scroll_size=10000).
<b>es_scroll_time = 2m</b>	string value	This parameter is a time value parameter (for example: es_scroll_time=2m), indicating for how long the nodes that participate in the search will maintain relevant resources in order to continue and support it.

Configuration option = Default value	Type	Description
<b>hmac_keys = SECRET_KEY</b>	string value	<p>Secret key(s) to use for encrypting context data for performance profiling. This string value should have the following format: &lt;key1&gt;[,&lt;key2&gt;,...&lt;keyn&gt;], where each key is some random string. A user who triggers the profiling via the REST API has to set one of these keys in the headers of the REST API call to include profiling results of this node for this particular project.</p> <p>Both "enabled" flag and "hmac_keys" config options should be set to enable profiling. Also, to generate correct profiling information across all services at least one key needs to be consistent between OpenStack projects. This ensures it can be used from client side to generate the trace, containing information from all possible resources.</p>
<b>sentinel_service_name = mymaster</b>	string value	<p>Redis sentinel uses a service name to identify a master redis service. This parameter defines the name (for example: sentinel_service_name=mymaster).</p>
<b>socket_timeout = 0.1</b>	floating point value	<p>Redis sentinel provides a timeout option on the connections. This parameter defines that timeout (for example: socket_timeout=0.1).</p>
<b>trace_sqlalchemy = False</b>	boolean value	<p>Enables SQL requests profiling in services. Default value is False (SQL requests won't be traced).</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True: Enables SQL requests profiling. Each SQL query will be part of the trace and can be analyzed by how much time was spent for that.</li> <li>● False: Disables SQL requests profiling. The spent time is only shown on a higher level of operations. Single SQL queries cannot be analyzed this way.</li> </ul>

### 9.1.33. resource

The following table outlines the options available under the **[resource]** group in the `/etc/keystone/keystone.conf` file.

Table 9.33. resource

Configuration option = Default value	Type	Description
<b>admin_project_domain_name = None</b>	string value	Name of the domain that owns the <b>admin_project_name</b> . If left unset, then there is no admin project. <b>[resource]</b> <b>admin_project_name</b> must also be set to use this option.
<b>admin_project_name = None</b>	string value	This is a special project which represents cloud-level administrator privileges across services. Tokens scoped to this project will contain a true <b>is_admin_project</b> attribute to indicate to policy systems that the role assignments on that specific project should apply equally across every project. If left unset, then there is no admin project, and thus no explicit means of cross-project role assignments. <b>[resource]</b> <b>admin_project_domain_name</b> must also be set to use this option.
<b>cache_time = None</b>	integer value	Time to cache resource data in seconds. This has no effect unless global caching is enabled.
<b>caching = True</b>	boolean value	Toggle for resource caching. This has no effect unless global caching is enabled.
<b>domain_name_url_safe = off</b>	string value	This controls whether the names of domains are restricted from containing URL-reserved characters. If set to <b>new</b> , attempts to create or update a domain with a URL-unsafe name will fail. If set to <b>strict</b> , attempts to scope a token with a URL-unsafe domain name will fail, thereby forcing all domain names to be updated to be URL-safe.
<b>driver = sql</b>	string value	Entry point for the resource driver in the <b>keystone.resource</b> namespace. Only a <b>sql</b> driver is supplied by keystone. Unless you are writing proprietary drivers for keystone, you do not need to set this option.
<b>list_limit = None</b>	integer value	Maximum number of entities that will be returned in a resource collection.
<b>project_name_url_safe = off</b>	string value	This controls whether the names of projects are restricted from containing URL-reserved characters. If set to <b>new</b> , attempts to create or update a project with a URL-unsafe name will fail. If set to <b>strict</b> , attempts to scope a token with a URL-unsafe project name will fail, thereby forcing all project names to be updated to be URL-safe.

### 9.1.34. revoke

The following table outlines the options available under the **[revoke]** group in the `/etc/keystone/keystone.conf` file.

Table 9.34. revoke

Configuration option = Default value	Type	Description
<b>cache_time = 3600</b>	integer value	Time to cache the revocation list and the revocation events (in seconds). This has no effect unless global and <b>[revoke] caching</b> are both enabled.
<b>caching = True</b>	boolean value	Toggle for revocation event caching. This has no effect unless global caching is enabled.
<b>driver = sql</b>	string value	Entry point for the token revocation backend driver in the <b>keystone.revoke</b> namespace. Keystone only provides a <b>sql</b> driver, so there is no reason to set this option unless you are providing a custom entry point.
<b>expiration_buffer = 1800</b>	integer value	The number of seconds after a token has expired before a corresponding revocation event may be purged from the backend.

### 9.1.35. role

The following table outlines the options available under the **[role]** group in the `/etc/keystone/keystone.conf` file.

Table 9.35. role

Configuration option = Default value	Type	Description
<b>cache_time = None</b>	integer value	Time to cache role data, in seconds. This has no effect unless both global caching and <b>[role] caching</b> are enabled.
<b>caching = True</b>	boolean value	Toggle for role caching. This has no effect unless global caching is enabled. In a typical deployment, there is no reason to disable this.
<b>driver = None</b>	string value	Entry point for the role backend driver in the <b>keystone.role</b> namespace. Keystone only provides a <b>sql</b> driver, so there's no reason to change this unless you are providing a custom entry point.

Configuration option = Default value	Type	Description
<b>list_limit = None</b>	integer value	Maximum number of entities that will be returned in a role collection. This may be useful to tune if you have a large number of discrete roles in your deployment.

### 9.1.36. saml

The following table outlines the options available under the **[saml]** group in the **/etc/keystone/keystone.conf** file.

Table 9.36. saml

Configuration option = Default value	Type	Description
<b>assertion_expiration_time = 3600</b>	integer value	Determines the lifetime for any SAML assertions generated by keystone, using <b>NotOnOrAfter</b> attributes.
<b>certfile = /etc/keystone/ssl/certs/signing_cert.pem</b>	string value	Absolute path to the public certificate file to use for SAML signing. The value cannot contain a comma (,).
<b>idp_contact_company = Example, Inc.</b>	string value	This is the company name of the identity provider's contact person.
<b>idp_contact_email = support@example.com</b>	string value	This is the email address of the identity provider's contact person.
<b>idp_contact_name = SAML Identity Provider Support</b>	string value	This is the given name of the identity provider's contact person.
<b>idp_contact_surname = Support</b>	string value	This is the surname of the identity provider's contact person.
<b>idp_contact_telephone = +1 800 555 0100</b>	string value	This is the telephone number of the identity provider's contact person.
<b>idp_contact_type = other</b>	string value	This is the type of contact that best describes the identity provider's contact person.



Configuration option = Default value	Type	Description
<b>idp_entity_id = None</b>	uri value	This is the unique entity identifier of the identity provider (keystone) to use when generating SAML assertions. This value is required to generate identity provider metadata and must be a URI (a URL is recommended). For example: <b>https://keystone.example.com/v3/OS-FEDERATION/saml2/idp.</b>
<b>idp_lang = en</b>	string value	This is the language used by the identity provider's organization.
<b>idp_metadata_path = /etc/keystone/saml2_idp_metadata.xml</b>	string value	Absolute path to the identity provider metadata file. This file should be generated with the <b>keystone-manage saml_idp_metadata</b> command. There is typically no reason to change this value.
<b>idp_organization_display_name = OpenStack SAML Identity Provider</b>	string value	This is the name of the identity provider's organization to be displayed.
<b>idp_organization_name = SAML Identity Provider</b>	string value	This is the name of the identity provider's organization.
<b>idp_organization_url = https://example.com/</b>	uri value	This is the URL of the identity provider's organization. The URL referenced here should be useful to humans.
<b>idp_sso_endpoint = None</b>	uri value	This is the single sign-on (SSO) service location of the identity provider which accepts HTTP POST requests. A value is required to generate identity provider metadata. For example: <b>https://keystone.example.com/v3/OS-FEDERATION/saml2/sso.</b>
<b>keyfile = /etc/keystone/ssl/private/signing_key.pem</b>	string value	Absolute path to the private key file to use for SAML signing. The value cannot contain a comma (,).
<b>relay_state_prefix = ss:mem:</b>	string value	The prefix of the RelayState SAML attribute to use when generating enhanced client and proxy (ECP) assertions. In a typical deployment, there is no reason to change this value.

Configuration option = Default value	Type	Description
<b>xmlsec1_binary = xmlsec1</b>	string value	Name of, or absolute path to, the binary to be used for XML signing. Although only the XML Security Library ( <b>xmlsec1</b> ) is supported, it may have a non-standard name or path on your system. If keystone cannot find the binary itself, you may need to install the appropriate package, use this option to specify an absolute path, or adjust keystone's PATH environment variable.

### 9.1.37. security\_compliance

The following table outlines the options available under the **[security\_compliance]** group in the `/etc/keystone/keystone.conf` file.

Table 9.37. security\_compliance

Configuration option = Default value	Type	Description
<b>change_password_upon_ first_use = False</b>	boolean value	Enabling this option requires users to change their password when the user is created, or upon administrative reset. Before accessing any services, affected users will have to change their password. To ignore this requirement for specific users, such as service users, set the <b>options</b> attribute <b>ignore_change_password_upon_first_use</b> to <b>True</b> for the desired user via the update user API. This feature is disabled by default. This feature is only applicable with the <b>sql</b> backend for the <b>[identity] driver</b> .
<b>disable_user_account_da ys_inactive = None</b>	integer value	The maximum number of days a user can go without authenticating before being considered "inactive" and automatically disabled (locked). This feature is disabled by default; set any value to enable it. This feature depends on the <b>sql</b> backend for the <b>[identity] driver</b> . When a user exceeds this threshold and is considered "inactive", the user's <b>enabled</b> attribute in the HTTP API may not match the value of the user's <b>enabled</b> column in the user table.

Configuration option = Default value	Type	Description
<b>lockout_duration = 1800</b>	integer value	The number of seconds a user account will be locked when the maximum number of failed authentication attempts (as specified by <b>[security_compliance] lockout_failure_attempts</b> ) is exceeded. Setting this option will have no effect unless you also set <b>[security_compliance] lockout_failure_attempts</b> to a non-zero value. This feature depends on the <b>sql</b> backend for the <b>[identity] driver</b> .
<b>lockout_failure_attempts = None</b>	integer value	The maximum number of times that a user can fail to authenticate before the user account is locked for the number of seconds specified by <b>[security_compliance] lockout_duration</b> . This feature is disabled by default. If this feature is enabled and <b>[security_compliance] lockout_duration</b> is not set, then users may be locked out indefinitely until the user is explicitly enabled via the API. This feature depends on the <b>sql</b> backend for the <b>[identity] driver</b> .
<b>minimum_password_age = 0</b>	integer value	The number of days that a password must be used before the user can change it. This prevents users from changing their passwords immediately in order to wipe out their password history and reuse an old password. This feature does not prevent administrators from manually resetting passwords. It is disabled by default and allows for immediate password changes. This feature depends on the <b>sql</b> backend for the <b>[identity] driver</b> . Note: If <b>[security_compliance] password_expires_days</b> is set, then the value for this option should be less than the <b>password_expires_days</b> .
<b>password_expires_days = None</b>	integer value	The number of days for which a password will be considered valid before requiring it to be changed. This feature is disabled by default. If enabled, new password changes will have an expiration date, however existing passwords would not be impacted. This feature depends on the <b>sql</b> backend for the <b>[identity] driver</b> .

Configuration option = Default value	Type	Description
<b>password_regex = None</b>	string value	The regular expression used to validate password strength requirements. By default, the regular expression will match any password. The following is an example of a pattern which requires at least 1 letter, 1 digit, and have a minimum length of 7 characters: <code>^(?=.*\d)(?=.*[a-zA-Z]).{7,}\$</code> This feature depends on the <b>sql</b> backend for the <b>[identity] driver</b> .
<b>password_regex_description = None</b>	string value	Describe your password regular expression here in language for humans. If a password fails to match the regular expression, the contents of this configuration variable will be returned to users to explain why their requested password was insufficient.
<b>unique_last_password_count = 1</b>	integer value	This controls the number of previous user password iterations to keep in history, in order to enforce that newly created passwords are unique. The total number which includes the new password should not be greater or equal to this value. Setting the value to one (the default) disables this feature. Thus, to enable this feature, values must be greater than 1. This feature depends on the <b>sql</b> backend for the <b>[identity] driver</b> .

### 9.1.38. shadow\_users

The following table outlines the options available under the **[shadow\_users]** group in the `/etc/keystone/keystone.conf` file.

Table 9.38. shadow\_users

Configuration option = Default value	Type	Description
<b>driver = sql</b>	string value	Entry point for the shadow users backend driver in the <b>keystone.identity.shadow_users</b> namespace. This driver is used for persisting local user references to externally-managed identities (via federation, LDAP, etc). Keystone only provides a <b>sql</b> driver, so there is no reason to change this option unless you are providing a custom entry point.

### 9.1.39. signing

The following table outlines the options available under the **[signing]** group in the `/etc/keystone/keystone.conf` file.

Table 9.39. signing

Configuration option = Default value	Type	Description
<b>ca_certs =</b> <b>/etc/keystone/ssl/certs/ca.</b> <b>pem</b>	string value	Absolute path to the public certificate authority (CA) file to use when creating self-signed certificates with <b>keystone-manage pki_setup</b> . Set this together with <b>[signing] ca_key</b> . There is no reason to set this option unless you are requesting revocation lists in a non-production environment. Use a <b>[signing] certfile</b> issued from a trusted certificate authority instead.
<b>ca_key =</b> <b>/etc/keystone/ssl/private/c</b> <b>akey.pem</b>	string value	Absolute path to the private certificate authority (CA) key file to use when creating self-signed certificates with <b>keystone-manage pki_setup</b> . Set this together with <b>[signing] ca_certs</b> . There is no reason to set this option unless you are requesting revocation lists in a non-production environment. Use a <b>[signing] certfile</b> issued from a trusted certificate authority instead.
<b>cert_subject =</b> <b>/C=US/ST=Unset/L=Unset/</b> <b>O=Unset/CN=www.examp</b> <b>le.com</b>	string value	The certificate subject to use when generating a self-signed token signing certificate. There is no reason to set this option unless you are requesting revocation lists in a non-production environment. Use a <b>[signing] certfile</b> issued from a trusted certificate authority instead.
<b>certfile =</b> <b>/etc/keystone/ssl/certs/sig</b> <b>ning_cert.pem</b>	string value	Absolute path to the public certificate file to use for signing responses to revocation lists requests. Set this together with <b>[signing] keyfile</b> . For non-production environments, you may be interested in using <b>keystone-manage pki_setup</b> to generate self-signed certificates.
<b>key_size = 2048</b>	integer value	Key size (in bits) to use when generating a self-signed token signing certificate. There is no reason to set this option unless you are requesting revocation lists in a non-production environment. Use a <b>[signing] certfile</b> issued from a trusted certificate authority instead.
<b>keyfile =</b> <b>/etc/keystone/ssl/private/s</b> <b>igning_key.pem</b>	string value	Absolute path to the private key file to use for signing responses to revocation lists requests. Set this together with <b>[signing] certfile</b> .
<b>valid_days = 3650</b>	integer value	The validity period (in days) to use when generating a self-signed token signing certificate. There is no reason to set this option unless you are requesting revocation lists in a non-production environment. Use a <b>[signing] certfile</b> issued from a trusted certificate authority instead.

## 9.1.40. token

The following table outlines the options available under the **[token]** group in the `/etc/keystone/keystone.conf` file.

Table 9.40. token

Configuration option = Default value	Type	Description
<b>allow_expired_window = 172800</b>	integer value	This controls the number of seconds that a token can be retrieved for beyond the built-in expiry time. This allows long running operations to succeed. Defaults to two days.
<b>allow_rescope_scoped_token = True</b>	boolean value	This toggles whether scoped tokens may be re-scoped to a new project or domain, thereby preventing users from exchanging a scoped token (including those with a default project scope) for any other token. This forces users to either authenticate for unscoped tokens (and later exchange that unscoped token for tokens with a more specific scope) or to provide their credentials in every request for a scoped token to avoid re-scoping altogether.
<b>bind = []</b>	list value	This is a list of external authentication mechanisms which should add token binding metadata to tokens, such as <b>kerberos</b> or <b>x509</b> . Binding metadata is enforced according to the <b>[token] enforce_token_bind</b> option.
<b>cache_on_issue = True</b>	boolean value	Enable storing issued token data to token validation cache so that first token validation doesn't actually cause full validation cycle. This option has no effect unless global caching and token caching are enabled.
<b>cache_time = None</b>	integer value	The number of seconds to cache token creation and validation data. This has no effect unless both global and <b>[token] caching</b> are enabled.
<b>caching = True</b>	boolean value	Toggle for caching token creation and validation data. This has no effect unless global caching is enabled.
<b>driver = sql</b>	string value	Entry point for the token persistence backend driver in the <b>keystone.token.persistence</b> namespace. Keystone provides the <b>sql</b> driver. The <b>sql</b> option (default) depends on the options in your <b>[database]</b> section. If you're using the <b>fernet [token] provider</b> , this backend will not be utilized to persist tokens at all.

Configuration option = Default value	Type	Description
<b>enforce_token_bind = permissive</b>	string value	This controls the token binding enforcement policy on tokens presented to keystone with token binding metadata (as specified by the <b>[token] bind</b> option). <b>disabled</b> completely bypasses token binding validation. <b>permissive</b> and <b>strict</b> do not require tokens to have binding metadata (but will validate it if present), whereas <b>required</b> will always demand tokens to having binding metadata. <b>permissive</b> will allow unsupported binding metadata to pass through without validation (usually to be validated at another time by another component), whereas <b>strict</b> and <b>required</b> will demand that the included binding metadata be supported by keystone.
<b>expiration = 3600</b>	integer value	The amount of time that a token should remain valid (in seconds). Drastically reducing this value may break "long-running" operations that involve multiple services to coordinate together, and will force users to authenticate with keystone more frequently. Drastically increasing this value will increase load on the <b>[token] driver</b> , as more tokens will be simultaneously valid. Keystone tokens are also bearer tokens, so a shorter duration will also reduce the potential security impact of a compromised token.
<b>infer_roles = True</b>	boolean value	This controls whether roles should be included with tokens that are not directly assigned to the token's scope, but are instead linked implicitly to other role assignments.
<b>provider = fernet</b>	string value	Entry point for the token provider in the <b>keystone.token.provider</b> namespace. The token provider controls the token construction, validation, and revocation operations. Keystone includes <b>fernet</b> and <b>uuid</b> token providers. <b>uuid</b> tokens must be persisted (using the backend specified in the <b>[token] driver</b> option), but do not require any extra configuration or setup. <b>fernet</b> tokens do not need to be persisted at all, but require that you run <b>keystone-manage fernet_setup</b> (also see the <b>keystone-manage fernet_rotate</b> command).
<b>revoke_by_id = True</b>	boolean value	This toggles support for revoking individual tokens by the token identifier and thus various token enumeration operations (such as listing all tokens issued to a specific user). These operations are used to determine the list of tokens to consider revoked. Do not disable this option if you're using the <b>kvs [revoke] driver</b> .

### 9.1.41. tokenless\_auth

The following table outlines the options available under the **[tokenless\_auth]** group in the `/etc/keystone/keystone.conf` file.

Table 9.41. tokenless\_auth

Configuration option = Default value	Type	Description
<b>issuer_attribute = SSL_CLIENT_I_DN</b>	string value	The name of the WSGI environment variable used to pass the issuer of the client certificate to keystone. This attribute is used as an identity provider ID for the X.509 tokenless authorization along with the protocol to look up its corresponding mapping. In a typical deployment, there is no reason to change this value.
<b>protocol = x509</b>	string value	The federated protocol ID used to represent X.509 tokenless authorization. This is used in combination with the value of <b>[tokenless_auth] issuer_attribute</b> to find a corresponding federated mapping. In a typical deployment, there is no reason to change this value.
<b>trusted_issuer = []</b>	multi valued	The list of distinguished names which identify trusted issuers of client certificates allowed to use X.509 tokenless authorization. If the option is absent then no certificates will be allowed. The format for the values of a distinguished name (DN) must be separated by a comma and contain no spaces. Furthermore, because an individual DN may contain commas, this configuration option may be repeated multiple times to represent multiple values. For example, keystone.conf would include two consecutive lines in order to trust two different DNs, such as <b>trusted_issuer = CN=john,OU=keystone,O=openstack</b> and <b>trusted_issuer = CN=mary,OU=eng,O=abc</b> .

### 9.1.42. trust

The following table outlines the options available under the **[trust]** group in the `/etc/keystone/keystone.conf` file.

Table 9.42. trust

Configuration option = Default value	Type	Description
---	------	-------------



Configuration option = Default value	Type	Description
<b>allow_redelegation = False</b>	boolean value	Allows authorization to be redelegated from one user to another, effectively chaining trusts together. When disabled, the <b>remaining_uses</b> attribute of a trust is constrained to be zero.
<b>driver = sql</b>	string value	Entry point for the trust backend driver in the <b>keystone.trust</b> namespace. Keystone only provides a <b>sql</b> driver, so there is no reason to change this unless you are providing a custom entry point.
<b>enabled = True</b>	boolean value	Delegation and impersonation features using trusts can be optionally disabled.
<b>max_redelegation_count = 3</b>	integer value	Maximum number of times that authorization can be redelegated from one user to another in a chain of trusts. This number may be reduced further for a specific trust.

### 9.1.43. unified\_limit

The following table outlines the options available under the **[unified\_limit]** group in the **/etc/keystone/keystone.conf** file.

Table 9.43. unified\_limit

Configuration option = Default value	Type	Description
<b>cache_time = None</b>	integer value	Time to cache unified limit data, in seconds. This has no effect unless both global caching and <b>[unified_limit] caching</b> are enabled.
<b>caching = True</b>	boolean value	Toggle for unified limit caching. This has no effect unless global caching is enabled. In a typical deployment, there is no reason to disable this.
<b>driver = sql</b>	string value	Entry point for the unified limit backend driver in the <b>keystone.unified_limit</b> namespace. Keystone only provides a <b>sql</b> driver, so there's no reason to change this unless you are providing a custom entry point.
<b>list_limit = None</b>	integer value	Maximum number of entities that will be returned in a role collection. This may be useful to tune if you have a large number of unified limits in your deployment.

## CHAPTER 10. NEUTRON

The following chapter contains information about the configuration options in the **neutron** service.

### 10.1. DHCP\_AGENT.INI

This section contains options for the `/etc/neutron/dhcp_agent.ini` file.

#### 10.1.1. DEFAULT

The following table outlines the options available under the **[DEFAULT]** group in the `/etc/neutron/dhcp_agent.ini` file.

Table 10.1. DEFAULT

Configuration option = Default value	Type	Description
<b>debug = False</b>	boolean value	If set to true, the logging level will be set to DEBUG instead of the default INFO level.
<b>default_log_levels =</b> [u'amqp=WARN', u'amqpplib=WARN', u'boto=WARN', u'qpid=WARN', u'sqlalchemy=WARN', u'suds=INFO', u'oslo.messaging=INFO', u'oslo_messaging=INFO', u'iso8601=WARN', u'requests.packages.urlli b3.connectionpool=WAR N', u'urllib3.connectionpool= WARN', u'websocket=WARN', u'requests.packages.urlli b3.util.retry=WARN', u'urllib3.util.retry=WARN', u'keystonemiddleware=W ARN', u'routes.middleware=WA RN', u'stevedore=WARN', u'taskflow=WARN', u'keystoneauth=WARN', u'oslo.cache=INFO', u'dogpile.core.dogpile=IN FO']	list value	List of package logging levels in logger=LEVEL pairs. This option is ignored if log_config_append is set.
<b>dhcp_broadcast_reply = False</b>	boolean value	Use broadcast in DHCP replies.

Configuration option = Default value	Type	Description
<b>dhcp_confs =</b> <b>\$state_path/dhcp</b>	string value	Location to store DHCP server config files.
<b>dhcp_driver =</b> <b>neutron.agent.linux.dhcp.</b> <b>Dnsmasq</b>	string value	The driver used to manage the DHCP server.
<b>dhcp_rebinding_time = 0</b>	integer value	DHCP rebinding time T2 (in seconds). If set to 0, it will default to 7/8 of the lease time.
<b>dhcp_renewal_time = 0</b>	integer value	DHCP renewal time T1 (in seconds). If set to 0, it will default to half of the lease time.
<b>dnsmasq_base_log_dir =</b> <b>None</b>	string value	Base log dir for dnsmasq logging. The log contains DHCP and DNS log information and is useful for debugging issues with either DHCP or DNS. If this section is null, disable dnsmasq log.
<code>`dnsmasq_config_file = `</code>	string value	Override the default dnsmasq settings with this file.
<b>dnsmasq_dns_servers =</b> <b>[]</b>	list value	Comma-separated list of the DNS servers which will be used as forwarders.
<b>dnsmasq_lease_max =</b> <b>16777216</b>	integer value	Limit number of leases to prevent a denial-of-service.
<b>dnsmasq_local_resolv =</b> <b>False</b>	boolean value	Enables the dnsmasq service to provide name resolution for instances via DNS resolvers on the host running the DHCP agent. Effectively removes the <code>--no-resolv</code> option from the dnsmasq process arguments. Adding custom DNS resolvers to the <code>dnsmasq_dns_servers</code> option disables this feature.
<b>enable_isolated_metadata</b> <b>a = False</b>	boolean value	The DHCP server can assist with providing metadata support on isolated networks. Setting this value to True will cause the DHCP server to append specific host routes to the DHCP request. The metadata service will only be activated when the subnet does not contain any router port. The guest instance must be configured to request host routes via DHCP (Option 121). This option doesn't have any effect when <code>force_metadata</code> is set to True.

Configuration option = Default value	Type	Description
<b>enable_metadata_network = False</b>	boolean value	Allows for serving metadata requests coming from a dedicated metadata access network whose CIDR is 169.254.169.254/16 (or larger prefix), and is connected to a Neutron router from which the VMs send metadata:1 request. In this case DHCP Option 121 will not be injected in VMs, as they will be able to reach 169.254.169.254 through a router. This option requires <code>enable_isolated_metadata = True</code> .
<b>fatal_deprecations = False</b>	boolean value	Enables or disables fatal status of deprecations.
<b>force_metadata = False</b>	boolean value	In some cases the Neutron router is not present to provide the metadata IP but the DHCP server can be used to provide this info. Setting this value will force the DHCP server to append specific host routes to the DHCP request. If this option is set, then the metadata service will be activated for all the networks.
<code>`instance_format = [instance: %(uuid)s]`</code>	string value	The format for an instance that is passed with the log message.
<code>`instance_uuid_format = [instance: %(uuid)s]`</code>	string value	The format for an instance UUID that is passed with the log message.
<b>interface_driver = None</b>	string value	The driver used to manage the virtual interface.
<b>log-config-append = None</b>	string value	The name of a logging configuration file. This file is appended to any existing logging configuration files. For details about logging configuration files, see the Python logging module documentation. Note that when logging configuration files are used then all logging configuration is set in the configuration file and other logging configuration options are ignored (for example, <code>logging_context_format_string</code> ).
<b>log-date-format = %Y-%m-%d %H:%M:%S</b>	string value	Defines the format string for <code>%(asctime)s</code> in log records. Default: <code>%(default)s</code> . This option is ignored if <code>log_config_append</code> is set.
<b>log-dir = None</b>	string value	(Optional) The base directory used for relative <code>log_file</code> paths. This option is ignored if <code>log_config_append</code> is set.

Configuration option = Default value	Type	Description
<b>log-file = None</b>	string value	(Optional) Name of log file to send logging output to. If no default is set, logging will go to stderr as defined by use_stderr. This option is ignored if log_config_append is set.
<b>logging_context_format_string = %(asctime)s.% (msecs)03d %(process)d %(levelname)s %(name)s [%(request_id)s % (user_identity)s] % (instance)s%(message)s</b>	string value	Format string to use for log messages with context.
<b>logging_debug_format_s uffix = %(funcName)s % (pathname)s:%(lineno)d</b>	string value	Additional data to append to log message when logging level for the message is DEBUG.
<b>logging_default_format_s tring = %(asctime)s.% (msecs)03d %(process)d %(levelname)s %(name)s [-] %(instance)s% (message)s</b>	string value	Format string to use for log messages when context is undefined.
<b>logging_exception_prefix = %(asctime)s.% (msecs)03d %(process)d ERROR %(name)s % (instance)s</b>	string value	Prefix each line of exception output with this format.
<b>logging_user_identity_for mat = %(user)s % (tenant)s %(domain)s % (user_domain)s % (project_domain)s</b>	string value	Defines the format string for %(user_identity)s that is used in logging_context_format_string.
<b>num_sync_threads = 4</b>	integer value	Number of threads to use during sync process. Should not exceed connection pool size configured on server.
<b>ovs_integration_bridge = br-int</b>	string value	Name of Open vSwitch bridge to use

Configuration option = Default value	Type	Description
<b>ovs_use_veth = False</b>	boolean value	Uses veth for an OVS interface or not. Support kernels with limited namespace support (e.g. RHEL 6.5) and rate limiting on router's gateway port so long as ovs_use_veth is set to True.
<b>publish_errors = False</b>	boolean value	Enables or disables publication of error events.
<b>rate_limit_burst = 0</b>	integer value	Maximum number of logged messages per rate_limit_interval.
<b>rate_limit_except_level = CRITICAL</b>	string value	Log level name used by rate limiting: CRITICAL, ERROR, INFO, WARNING, DEBUG or empty string. Logs with level greater or equal to rate_limit_except_level are not filtered. An empty string means that all levels are filtered.
<b>rate_limit_interval = 0</b>	integer value	Interval, number of seconds, of log rate limiting.
<b>resync_interval = 5</b>	integer value	The DHCP agent will resync its state with Neutron to recover from any transient notification or RPC errors. The interval is number of seconds between attempts.
<b>syslog-log-facility = LOG_USER</b>	string value	Syslog facility to receive log lines. This option is ignored if log_config_append is set.
<b>use-journal = False</b>	boolean value	Enable journald for logging. If running in a systemd environment you may wish to enable journal support. Doing so will use the journal native protocol which includes structured metadata in addition to log messages. This option is ignored if log_config_append is set.
<b>use-json = False</b>	boolean value	Use JSON formatting for logging. This option is ignored if log_config_append is set.
<b>use-syslog = False</b>	boolean value	Use syslog for logging. Existing syslog format is DEPRECATED and will be changed later to honor RFC5424. This option is ignored if log_config_append is set.
<b>use_stderr = False</b>	boolean value	Log output to standard error. This option is ignored if log_config_append is set.

Configuration option = Default value	Type	Description
<b>watch-log-file = False</b>	boolean value	Uses logging handler designed to watch file system. When log file is moved or removed this handler will open a new log file with specified path instantaneously. It makes sense only if log_file option is specified and Linux platform is used. This option is ignored if log_config_append is set.

### 10.1.2. agent

The following table outlines the options available under the **[agent]** group in the `/etc/neutron/dhcp_agent.ini` file.

Table 10.2. agent

Configuration option = Default value	Type	Description
<b>availability_zone = nova</b>	string value	Availability zone of this node
<b>log_agent_heartbeats = False</b>	boolean value	Log agent heartbeats
<b>report_interval = 30</b>	floating point value	Seconds between nodes reporting state to server; should be less than agent_down_time, best if it is half or less than agent_down_time.

### 10.1.3. ovs

The following table outlines the options available under the **[ovs]** group in the `/etc/neutron/dhcp_agent.ini` file.

Table 10.3. ovs

Configuration option = Default value	Type	Description
<b>bridge_mac_table_size = 50000</b>	integer value	The maximum number of MAC addresses to learn on a bridge managed by the Neutron OVS agent. Values outside a reasonable range (10 to 1,000,000) might be overridden by Open vSwitch according to the documentation.
<b>ovsdb_connection = tcp:127.0.0.1:6640</b>	string value	The connection string for the OVSDB backend. Will be used by ovsdb-client when monitoring and used for the all ovsdb commands when native ovsdb_interface is enabled

Configuration option = Default value	Type	Description
<b>ovsdb_interface = native</b>	string value	The interface for interacting with the OVSDB
<b>ovsdb_timeout = 10</b>	integer value	Timeout in seconds for ovsdb commands. If the timeout expires, ovsdb commands will fail with ALARMCLOCK error.
<b>ssl_ca_cert_file = None</b>	string value	The Certificate Authority (CA) certificate to use when interacting with OVSDB. Required when using an "ssl:" prefixed ovsdb_connection
<b>ssl_cert_file = None</b>	string value	The SSL certificate file to use when interacting with OVSDB. Required when using an "ssl:" prefixed ovsdb_connection
<b>ssl_key_file = None</b>	string value	The SSL private key file to use when interacting with OVSDB. Required when using an "ssl:" prefixed ovsdb_connection

## 10.2. L3\_AGENT.INI

This section contains options for the `/etc/neutron/l3_agent.ini` file.

### 10.2.1. DEFAULT

The following table outlines the options available under the **[DEFAULT]** group in the `/etc/neutron/l3_agent.ini` file.

Table 10.4. DEFAULT

Configuration option = Default value	Type	Description
---	------	-------------



Configuration option = Default value	Type	Description
<b>agent_mode = legacy</b>	string value	The working mode for the agent. Allowed modes are: <i>legacy</i> - this preserves the existing behavior where the L3 agent is deployed on a centralized networking node to provide L3 services like DNAT, and SNAT. Use this mode if you do not want to adopt DVR. <i>dvr</i> - this mode enables DVR functionality and must be used for an L3 agent that runs on a compute host. <i>dvr_snat</i> - this enables centralized SNAT support in conjunction with DVR. This mode must be used for an L3 agent running on a centralized node (or in single-host deployments, e.g. devstack). <i>dvr_no_external</i> - this mode enables only East/West DVR routing functionality for a L3 agent that runs on a compute host, the North/South functionality such as DNAT and SNAT will be provided by the centralized network node that is running in <i>dvr_snat</i> mode. This mode should be used when there is no external network connectivity on the compute host.
<b>api_workers = None</b>	integer value	Number of separate API worker processes for service. If not specified, the default is equal to the number of CPUs available for best performance.
<b>debug = False</b>	boolean value	If set to true, the logging level will be set to DEBUG instead of the default INFO level.

Configuration option = Default value	Type	Description
<b>default_log_levels =</b> [u'amqp=WARN', u'amqplib=WARN', u'boto=WARN', u'qpid=WARN', u'sqlalchemy=WARN', u'suds=INFO', u'oslo.messaging=INFO', u'oslo_messaging=INFO', u'iso8601=WARN', u'requests.packages.urlli b3.connectionpool=WAR N', u'urllib3.connectionpool= WARN', u'websocket=WARN', u'requests.packages.urlli b3.util.retry=WARN', u'urllib3.util.retry=WARN', u'keystonemiddleware=W ARN', u'routes.middleware=WA RN', u'stevedore=WARN', u'taskflow=WARN', u'keystoneauth=WARN', u'oslo.cache=INFO', u'dogpile.core.dogpile=IN FO']	list value	List of package logging levels in logger=LEVEL pairs. This option is ignored if log_config_append is set.
<b>enable_metadata_proxy =</b> <b>True</b>	boolean value	Allow running metadata proxy.
<b>external_ingress_mark =</b> <b>0x2</b>	string value	Iptables mangle mark used to mark ingress from external network. This mark will be masked with 0xffff so that only the lower 16 bits will be used.
<code>`external_network_bridge = `</code>	string value	Name of bridge used for external network traffic. When this parameter is set, the L3 agent will plug an interface directly into an external bridge which will not allow any wiring by the L2 agent. Using this will result in incorrect port statuses. This option is deprecated and will be removed in Ocata.
<b>fatal_deprecations =</b> <b>False</b>	boolean value	Enables or disables fatal status of deprecations.

Configuration option = Default value	Type	Description
<code>`gateway_external_network_id = `</code>	string value	When <code>external_network_bridge</code> is set, each L3 agent can be associated with no more than one external network. This value should be set to the UUID of that external network. To allow L3 agent support multiple external networks, both the <code>external_network_bridge</code> and <code>gateway_external_network_id</code> must be left empty.
<b><code>ha_confs_path = \$state_path/ha_confs</code></b>	string value	Location to store <code>keepalived/conntrackd</code> config files
<b><code>ha_keepalived_state_change_server_threads = 4</code></b>	integer value	Number of concurrent threads for <code>keepalived</code> server connection requests. More threads create a higher CPU load on the agent node.
<b><code>ha_vrrp_advert_int = 2</code></b>	integer value	The advertisement interval in seconds
<b><code>ha_vrrp_auth_password = None</code></b>	string value	VRRP authentication password
<b><code>ha_vrrp_auth_type = PASS</code></b>	string value	VRRP authentication type
<b><code>ha_vrrp_health_check_interval = 0</code></b>	integer value	The VRRP health check interval in seconds. Values > 0 enable VRRP health checks. Setting it to 0 disables VRRP health checks. Recommended value is 5. This will cause pings to be sent to the gateway IP address(es) - requires <code>ICMP_ECHO_REQUEST</code> to be enabled on the gateway. If gateway fails, all routers will be reported as master, and master election will be repeated in round-robin fashion, until one of the router restore the gateway connection.
<b><code>handle_internal_only_routers = True</code></b>	boolean value	Indicates that this L3 agent should also handle routers that do not have an external network gateway configured. This option should be <code>True</code> only for a single agent in a Neutron deployment, and may be <code>False</code> for all agents if all routers must have an external network gateway.
<code>`instance_format = [instance: %(uuid)s] `</code>	string value	The format for an instance that is passed with the log message.
<code>`instance_uuid_format = [instance: %(uuid)s] `</code>	string value	The format for an instance UUID that is passed with the log message.
<b><code>interface_driver = None</code></b>	string value	The driver used to manage the virtual interface.

Configuration option = Default value	Type	Description
<code>`ipv6_gateway = `</code>	string value	With IPv6, the network used for the external gateway does not need to have an associated subnet, since the automatically assigned link-local address (LLA) can be used. However, an IPv6 gateway address is needed for use as the next-hop for the default route. If no IPv6 gateway address is configured here, (and only then) the neutron router will be configured to get its default route from router advertisements (RAs) from the upstream router; in which case the upstream router must also be configured to send these RAs. The <code>ipv6_gateway</code> , when configured, should be the LLA of the interface on the upstream router. If a next-hop using a global unique address (GUA) is desired, it needs to be done via a subnet allocated to the network and not through this parameter.
<b>log-config-append = None</b>	string value	The name of a logging configuration file. This file is appended to any existing logging configuration files. For details about logging configuration files, see the Python logging module documentation. Note that when logging configuration files are used then all logging configuration is set in the configuration file and other logging configuration options are ignored (for example, <code>logging_context_format_string</code> ).
<b>log-date-format = %Y-%m-%d %H:%M:%S</b>	string value	Defines the format string for <code>%(asctime)s</code> in log records. Default: <code>%(default)s</code> . This option is ignored if <code>log_config_append</code> is set.
<b>log-dir = None</b>	string value	(Optional) The base directory used for relative <code>log_file</code> paths. This option is ignored if <code>log_config_append</code> is set.
<b>log-file = None</b>	string value	(Optional) Name of log file to send logging output to. If no default is set, logging will go to <code>stderr</code> as defined by <code>use_stderr</code> . This option is ignored if <code>log_config_append</code> is set.
<b>logging_context_format_string = %(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [%(request_id)s %(user_identity)s] %(instance)s%(message)s</b>	string value	Format string to use for log messages with context.

Configuration option = Default value	Type	Description
<b>logging_debug_format_s uffix = %(funcName)s % (pathname)s:%(lineno)d</b>	string value	Additional data to append to log message when logging level for the message is DEBUG.
<b>logging_default_format_s tring = %(asctime)s.% (msecs)03d %(process)d %(levelname)s %(name)s [-] %(instance)s% (message)s</b>	string value	Format string to use for log messages when context is undefined.
<b>logging_exception_prefix = %(asctime)s.% (msecs)03d %(process)d ERROR %(name)s % (instance)s</b>	string value	Prefix each line of exception output with this format.
<b>logging_user_identity_for mat = %(user)s % (tenant)s %(domain)s % (user_domain)s % (project_domain)s</b>	string value	Defines the format string for %(user_identity)s that is used in logging_context_format_string.
<b>max_rtr_adv_interval = 100</b>	integer value	MaxRtrAdvInterval setting for radvd.conf
<b>metadata_access_mark = 0x1</b>	string value	Iptables mangle mark used to mark metadata valid requests. This mark will be masked with 0xffff so that only the lower 16 bits will be used.
<b>metadata_port = 9697</b>	port value	TCP Port used by Neutron metadata namespace proxy.
<b>min_rtr_adv_interval = 30</b>	integer value	MinRtrAdvInterval setting for radvd.conf
<b>ovs_integration_bridge = br-int</b>	string value	Name of Open vSwitch bridge to use
<b>ovs_use_veth = False</b>	boolean value	Uses veth for an OVS interface or not. Support kernels with limited namespace support (e.g. RHEL 6.5) and rate limiting on router's gateway port so long as ovs_use_veth is set to True.
<b>pd_confs = \$state_path/pd</b>	string value	Location to store IPv6 PD files.

Configuration option = Default value	Type	Description
<b>periodic_fuzzy_delay = 5</b>	integer value	Range of seconds to randomly delay when starting the periodic task scheduler to reduce stampeding. (Disable by setting to 0)
<b>periodic_interval = 40</b>	integer value	Seconds between running periodic tasks.
<b>prefix_delegation_driver = dibbler</b>	string value	Driver used for ipv6 prefix delegation. This needs to be an entry point defined in the neutron.agent.linux.pd_drivers namespace. See setup.cfg for entry points included with the neutron source.
<b>publish_errors = False</b>	boolean value	Enables or disables publication of error events.
<b>ra_confs = \$state_path/ra</b>	string value	Location to store IPv6 RA config files
<b>rate_limit_burst = 0</b>	integer value	Maximum number of logged messages per rate_limit_interval.
<b>rate_limit_except_level = CRITICAL</b>	string value	Log level name used by rate limiting: CRITICAL, ERROR, INFO, WARNING, DEBUG or empty string. Logs with level greater or equal to rate_limit_except_level are not filtered. An empty string means that all levels are filtered.
<b>rate_limit_interval = 0</b>	integer value	Interval, number of seconds, of log rate limiting.
<b>rpc_state_report_workers = 1</b>	integer value	Number of RPC worker processes dedicated to state reports queue.
<b>rpc_workers = 1</b>	integer value	Number of RPC worker processes for service.
<b>syslog-log-facility = LOG_USER</b>	string value	Syslog facility to receive log lines. This option is ignored if log_config_append is set.
<b>use-journal = False</b>	boolean value	Enable journald for logging. If running in a systemd environment you may wish to enable journal support. Doing so will use the journal native protocol which includes structured metadata in addition to log messages. This option is ignored if log_config_append is set.
<b>use-json = False</b>	boolean value	Use JSON formatting for logging. This option is ignored if log_config_append is set.

Configuration option = Default value	Type	Description
<b>use-syslog = False</b>	boolean value	Use syslog for logging. Existing syslog format is DEPRECATED and will be changed later to honor RFC5424. This option is ignored if log_config_append is set.
<b>use_stderr = False</b>	boolean value	Log output to standard error. This option is ignored if log_config_append is set.
<b>vendor_pen = 8888</b>	string value	A decimal value as Vendor's Registered Private Enterprise Number as required by RFC3315 DUID-EN.
<b>watch-log-file = False</b>	boolean value	Uses logging handler designed to watch file system. When log file is moved or removed this handler will open a new log file with specified path instantaneously. It makes sense only if log_file option is specified and Linux platform is used. This option is ignored if log_config_append is set.

### 10.2.2. agent

The following table outlines the options available under the **[agent]** group in the `/etc/neutron/l3_agent.ini` file.

Table 10.5. agent

Configuration option = Default value	Type	Description
<b>availability_zone = nova</b>	string value	Availability zone of this node
<b>extensions = []</b>	list value	Extensions list to use
<b>log_agent_heartbeats = False</b>	boolean value	Log agent heartbeats
<b>report_interval = 30</b>	floating point value	Seconds between nodes reporting state to server; should be less than agent_down_time, best if it is half or less than agent_down_time.

### 10.2.3. ovs

The following table outlines the options available under the **[ovs]** group in the `/etc/neutron/l3_agent.ini` file.

Table 10.6. ovs

Configuration option = Default value	Type	Description
<b>bridge_mac_table_size = 50000</b>	integer value	The maximum number of MAC addresses to learn on a bridge managed by the Neutron OVS agent. Values outside a reasonable range (10 to 1,000,000) might be overridden by Open vSwitch according to the documentation.
<b>ovsdb_connection = tcp:127.0.0.1:6640</b>	string value	The connection string for the OVSDB backend. Will be used by ovsdb-client when monitoring and used for the all ovsdb commands when native ovsdb_interface is enabled
<b>ovsdb_interface = native</b>	string value	The interface for interacting with the OVSDB
<b>ovsdb_timeout = 10</b>	integer value	Timeout in seconds for ovsdb commands. If the timeout expires, ovsdb commands will fail with ALARMCLOCK error.
<b>ssl_ca_cert_file = None</b>	string value	The Certificate Authority (CA) certificate to use when interacting with OVSDB. Required when using an "ssl:" prefixed ovsdb_connection
<b>ssl_cert_file = None</b>	string value	The SSL certificate file to use when interacting with OVSDB. Required when using an "ssl:" prefixed ovsdb_connection
<b>ssl_key_file = None</b>	string value	The SSL private key file to use when interacting with OVSDB. Required when using an "ssl:" prefixed ovsdb_connection

## 10.3. LINUXBRIDGE\_AGENT.INI

This section contains options for the `/etc/neutron/plugins/ml2/linuxbridge_agent.ini` file.

### 10.3.1. DEFAULT

The following table outlines the options available under the **[DEFAULT]** group in the `/etc/neutron/plugins/ml2/linuxbridge_agent.ini` file.

Table 10.7. DEFAULT

Configuration option = Default value	Type	Description
<b>debug = False</b>	boolean value	If set to true, the logging level will be set to DEBUG instead of the default INFO level.



Configuration option = Default value	Type	Description
<b>default_log_levels =</b> [u'amqp=WARN', u'amqplib=WARN', u'boto=WARN', u'qpid=WARN', u'sqlalchemy=WARN', u'suds=INFO', u'oslo.messaging=INFO', u'oslo_messaging=INFO', u'iso8601=WARN', u'requests.packages.urlli b3.connectionpool=WAR N', u'urllib3.connectionpool= WARN', u'websocket=WARN', u'requests.packages.urlli b3.util.retry=WARN', u'urllib3.util.retry=WARN', u'keystonemiddleware=W ARN', u'routes.middleware=WA RN', u'stevedore=WARN', u'taskflow=WARN', u'keystoneauth=WARN', u'oslo.cache=INFO', u'dogpile.core.dogpile=IN FO']	list value	List of package logging levels in logger=LEVEL pairs. This option is ignored if log_config_append is set.
<b>fatal_deprecations =</b> <b>False</b>	boolean value	Enables or disables fatal status of deprecations.
<code>`instance_format = [instance:            %(uuid)s] `</code>	string value	The format for an instance that is passed with the log message.
<code>`instance_uuid_format =            [instance: %(uuid)s] `</code>	string value	The format for an instance UUID that is passed with the log message.
<b>log-config-append = None</b>	string value	The name of a logging configuration file. This file is appended to any existing logging configuration files. For details about logging configuration files, see the Python logging module documentation. Note that when logging configuration files are used then all logging configuration is set in the configuration file and other logging configuration options are ignored (for example, logging_context_format_string).

Configuration option = Default value	Type	Description
<b>log-date-format = %Y-%m-%d %H:%M:%S</b>	string value	Defines the format string for <code>%(asctime)s</code> in log records. Default: <code>%(default)s</code> . This option is ignored if <code>log_config_append</code> is set.
<b>log-dir = None</b>	string value	(Optional) The base directory used for relative <code>log_file</code> paths. This option is ignored if <code>log_config_append</code> is set.
<b>log-file = None</b>	string value	(Optional) Name of log file to send logging output to. If no default is set, logging will go to <code>stderr</code> as defined by <code>use_stderr</code> . This option is ignored if <code>log_config_append</code> is set.
<b>logging_context_format_string = %(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [%(request_id)s %(user_identity)s] %(instance)s%(message)s</b>	string value	Format string to use for log messages with context.
<b>logging_debug_format_s uffix = %(funcName)s % (pathname)s:%(lineno)d</b>	string value	Additional data to append to log message when logging level for the message is <code>DEBUG</code> .
<b>logging_default_format_s tring = %(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [-] %(instance)s%(message)s</b>	string value	Format string to use for log messages when context is undefined.
<b>logging_exception_prefix = %(asctime)s.%(msecs)03d %(process)d ERROR %(name)s % (instance)s</b>	string value	Prefix each line of exception output with this format.
<b>logging_user_identity_for mat = %(user)s % (tenant)s %(domain)s % (user_domain)s % (project_domain)s</b>	string value	Defines the format string for <code>%(user_identity)s</code> that is used in <code>logging_context_format_string</code> .
<b>publish_errors = False</b>	boolean value	Enables or disables publication of error events.

Configuration option = Default value	Type	Description
<b>rate_limit_burst = 0</b>	integer value	Maximum number of logged messages per rate_limit_interval.
<b>rate_limit_except_level = CRITICAL</b>	string value	Log level name used by rate limiting: CRITICAL, ERROR, INFO, WARNING, DEBUG or empty string. Logs with level greater or equal to rate_limit_except_level are not filtered. An empty string means that all levels are filtered.
<b>rate_limit_interval = 0</b>	integer value	Interval, number of seconds, of log rate limiting.
<b>syslog-log-facility = LOG_USER</b>	string value	Syslog facility to receive log lines. This option is ignored if log_config_append is set.
<b>use-journal = False</b>	boolean value	Enable journald for logging. If running in a systemd environment you may wish to enable journal support. Doing so will use the journal native protocol which includes structured metadata in addition to log messages. This option is ignored if log_config_append is set.
<b>use-json = False</b>	boolean value	Use JSON formatting for logging. This option is ignored if log_config_append is set.
<b>use-syslog = False</b>	boolean value	Use syslog for logging. Existing syslog format is DEPRECATED and will be changed later to honor RFC5424. This option is ignored if log_config_append is set.
<b>use_stderr = False</b>	boolean value	Log output to standard error. This option is ignored if log_config_append is set.
<b>watch-log-file = False</b>	boolean value	Uses logging handler designed to watch file system. When log file is moved or removed this handler will open a new log file with specified path instantaneously. It makes sense only if log_file option is specified and Linux platform is used. This option is ignored if log_config_append is set.

### 10.3.2. agent

The following table outlines the options available under the **[agent]** group in the `/etc/neutron/plugins/ml2/linuxbridge_agent.ini` file.

Table 10.8. agent

Configuration option = Default value	Type	Description
<b>dscp = None</b>	integer value	The DSCP value to use for outer headers during tunnel encapsulation.
<b>dscp_inherit = False</b>	boolean value	If set to True, the DSCP value of tunnel interfaces is overwritten and set to inherit. The DSCP value of the inner header is then copied to the outer header.
<b>extensions = []</b>	list value	Extensions list to use
<b>polling_interval = 2</b>	integer value	The number of seconds the agent will wait between polling for local device changes.
<b>quitting_rpc_timeout = 10</b>	integer value	Set new timeout in seconds for new rpc calls after agent receives SIGTERM. If value is set to 0, rpc timeout won't be changed

### 10.3.3. linux\_bridge

The following table outlines the options available under the **[linux\_bridge]** group in the `/etc/neutron/plugins/ml2/linuxbridge_agent.ini` file.

Table 10.9. linux\_bridge

Configuration option = Default value	Type	Description
<b>bridge_mappings = []</b>	list value	List of <physical_network>:<physical_bridge>
<b>physical_interface_mappings = []</b>	list value	Comma-separated list of <physical_network>:<physical_interface> tuples mapping physical network names to the agent's node-specific physical network interfaces to be used for flat and VLAN networks. All physical networks listed in <code>network_vlan_ranges</code> on the server should have mappings to appropriate interfaces on each agent.

### 10.3.4. network\_log

The following table outlines the options available under the **[network\_log]** group in the `/etc/neutron/plugins/ml2/linuxbridge_agent.ini` file.

Table 10.10. network\_log

Configuration option = Default value	Type	Description
<b>burst_limit = 25</b>	integer value	Maximum number of packets per rate_limit.

Configuration option = Default value	Type	Description
<b>local_output_log_base = None</b>	string value	Output logfile path on agent side, default syslog file.
<b>rate_limit = 100</b>	integer value	Maximum packets logging per second.

### 10.3.5. securitygroup

The following table outlines the options available under the **[securitygroup]** group in the `/etc/neutron/plugins/ml2/linuxbridge_agent.ini` file.

Table 10.11. securitygroup

Configuration option = Default value	Type	Description
<b>enable_ipset = True</b>	boolean value	Use ipset to speed-up the iptables based security groups. Enabling ipset support requires that ipset is installed on L2 agent node.
<b>enable_security_group = True</b>	boolean value	Controls whether the neutron security group API is enabled in the server. It should be false when using no security groups or using the nova security group API.
<b>firewall_driver = None</b>	string value	Driver for security groups firewall in the L2 agent

### 10.3.6. vxlan

The following table outlines the options available under the **[vxlan]** group in the `/etc/neutron/plugins/ml2/linuxbridge_agent.ini` file.

Table 10.12. vxlan

Configuration option = Default value	Type	Description
<b>arp_responder = False</b>	boolean value	Enable local ARP responder which provides local responses instead of performing ARP broadcast into the overlay. Enabling local ARP responder is not fully compatible with the allowed-address-pairs extension.
<b>enable_vxlan = True</b>	boolean value	Enable VXLAN on the agent. Can be enabled when agent is managed by ml2 plugin using linuxbridge mechanism driver

Configuration option = Default value	Type	Description
<b>l2_population = False</b>	boolean value	Extension to use alongside ml2 plugin's l2population mechanism driver. It enables the plugin to populate VXLAN forwarding table.
<b>local_ip = None</b>	IP address value	IP address of local overlay (tunnel) network endpoint. Use either an IPv4 or IPv6 address that resides on one of the host network interfaces. The IP version of this value must match the value of the <i>overlay_ip_version</i> option in the ML2 plug-in configuration file on the neutron server node(s).
<b>multicast_ranges = []</b>	list value	Optional comma-separated list of <multicast address>:<vni_min>:<vni_max> triples describing how to assign a multicast address to VXLAN according to its VNI ID.
<b>tos = None</b>	integer value	TOS for vxlan interface protocol packets. This option is deprecated in favor of the dscp option in the AGENT section and will be removed in a future release. To convert the TOS value to DSCP, divide by 4.
<b>ttl = None</b>	integer value	TTL for vxlan interface protocol packets.
<b>udp_dstport = None</b>	port value	The UDP port used for VXLAN communication. By default, the Linux kernel doesn't use the IANA assigned standard value, so if you want to use it, this option must be set to 4789. It is not set by default because of backward compatibility.
<b>udp_srcport_max = 0</b>	port value	The maximum of the UDP source port range used for VXLAN communication.
<b>udp_srcport_min = 0</b>	port value	The minimum of the UDP source port range used for VXLAN communication.
<b>vxlan_group = 224.0.0.1</b>	string value	Multicast group(s) for vxlan interface. A range of group addresses may be specified by using CIDR notation. Specifying a range allows different VNIs to use different group addresses, reducing or eliminating spurious broadcast traffic to the tunnel endpoints. To reserve a unique group for each possible (24-bit) VNI, use a /8 such as 239.0.0.0/8. This setting must be the same on all the agents.

## 10.4. METADATA\_AGENT.INI

This section contains options for the `/etc/neutron/metadata_agent.ini` file.

### 10.4.1. DEFAULT

The following table outlines the options available under the **[DEFAULT]** group in the `/etc/neutron/metadata_agent.ini` file.

Table 10.13. DEFAULT

Configuration option = Default value	Type	Description
<b>auth_ca_cert = None</b>	string value	Certificate Authority public key (CA cert) file for ssl
<b>debug = False</b>	boolean value	If set to true, the logging level will be set to DEBUG instead of the default INFO level.
<b>default_log_levels =</b> [u'amqp=WARN', u'amqplib=WARN', u'boto=WARN', u'qpid=WARN', u'sqlalchemy=WARN', u'suds=INFO', u'oslo.messaging=INFO', u'oslo_messaging=INFO', u'iso8601=WARN', u'requests.packages.urlli b3.connectionpool=WAR N', u'urllib3.connectionpool= WARN', u'websocket=WARN', u'requests.packages.urlli b3.util.retry=WARN', u'urllib3.util.retry=WARN', u'keystonemiddleware=W ARN', u'routes.middleware=WA RN', u'stevedore=WARN', u'taskflow=WARN', u'keystoneauth=WARN', u'oslo.cache=INFO', u'dogpile.core.dogpile=IN FO']	list value	List of package logging levels in logger=LEVEL pairs. This option is ignored if log_config_append is set.
<b>fatal_deprecations =</b> <b>False</b>	boolean value	Enables or disables fatal status of deprecations.
<code>`instance_format = [instance: %(uuid)s] `</code>	string value	The format for an instance that is passed with the log message.

Configuration option = Default value	Type	Description
<code>`instance_uuid_format = [instance: %(uuid)s]`</code>	string value	The format for an instance UUID that is passed with the log message.
<b>log-config-append = None</b>	string value	The name of a logging configuration file. This file is appended to any existing logging configuration files. For details about logging configuration files, see the Python logging module documentation. Note that when logging configuration files are used then all logging configuration is set in the configuration file and other logging configuration options are ignored (for example, <code>logging_context_format_string</code> ).
<b>log-date-format = %Y-%m-%d %H:%M:%S</b>	string value	Defines the format string for <code>%(asctime)s</code> in log records. Default: <code>%(default)s</code> . This option is ignored if <code>log_config_append</code> is set.
<b>log-dir = None</b>	string value	(Optional) The base directory used for relative <code>log_file</code> paths. This option is ignored if <code>log_config_append</code> is set.
<b>log-file = None</b>	string value	(Optional) Name of log file to send logging output to. If no default is set, logging will go to <code>stderr</code> as defined by <code>use_stderr</code> . This option is ignored if <code>log_config_append</code> is set.
<b>logging_context_format_string = %(asctime)s.% (msecs)03d %(process)d %(levelname)s %(name)s [% (request_id)s % (user_identity)s] % (instance)s%(message)s</b>	string value	Format string to use for log messages with context.
<b>logging_debug_format_s uffix = %(funcName)s % (pathname)s:%(lineno)d</b>	string value	Additional data to append to log message when logging level for the message is <code>DEBUG</code> .
<b>logging_default_format_s tring = %(asctime)s.% (msecs)03d %(process)d %(levelname)s %(name)s [-] %(instance)s% (message)s</b>	string value	Format string to use for log messages when context is undefined.
<b>logging_exception_prefix = %(asctime)s.% (msecs)03d %(process)d ERROR %(name)s % (instance)s</b>	string value	Prefix each line of exception output with this format.



Configuration option = Default value	Type	Description
<b>logging_user_identity_for_mat = %(user)s %(tenant)s %(domain)s %(user_domain)s %(project_domain)s</b>	string value	Defines the format string for %(user_identity)s that is used in logging_context_format_string.
<b>metadata_backlog = 4096</b>	integer value	Number of backlog requests to configure the metadata server socket with
<code>`metadata_proxy_group = `</code>	string value	Group (gid or name) running metadata proxy after its initialization (if empty: agent effective group).
<code>`metadata_proxy_shared_secret = `</code>	string value	When proxying metadata requests, Neutron signs the Instance-ID header with a shared secret to prevent spoofing. You may select any string for a secret, but it must match here and in the configuration used by the Nova Metadata Server. NOTE: Nova uses the same config key, but in [neutron] section.
<b>metadata_proxy_socket = \$state_path/metadata_proxy</b>	string value	Location for Metadata Proxy UNIX domain socket.
<b>metadata_proxy_socket_mode = deduce</b>	string value	Metadata Proxy UNIX domain socket mode, 4 values allowed: <i>deduce</i> : deduce mode from metadata_proxy_user/group values, <i>user</i> : set metadata proxy socket mode to 0o644, to use when metadata_proxy_user is agent effective user or root, <i>group</i> : set metadata proxy socket mode to 0o664, to use when metadata_proxy_group is agent effective group or root, <i>all</i> : set metadata proxy socket mode to 0o666, to use otherwise.
<code>`metadata_proxy_user = `</code>	string value	User (uid or name) running metadata proxy after its initialization (if empty: agent effective user).
<b>metadata_workers = 4</b>	integer value	Number of separate worker processes for metadata server (defaults to half of the number of CPUs)
<code>`nova_client_cert = `</code>	string value	Client certificate for nova metadata api server.
<code>`nova_client_priv_key = `</code>	string value	Private key of client certificate.
<b>nova_metadata_host = 127.0.0.1</b>	unknown value	IP address or DNS name of Nova metadata server.
<b>nova_metadata_insecure = False</b>	boolean value	Allow to perform insecure SSL (https) requests to nova metadata

Configuration option = Default value	Type	Description
<b>nova_metadata_port = 8775</b>	port value	TCP Port used by Nova metadata server.
<b>nova_metadata_protocol = http</b>	string value	Protocol to access nova metadata, http or https
<b>publish_errors = False</b>	boolean value	Enables or disables publication of error events.
<b>rate_limit_burst = 0</b>	integer value	Maximum number of logged messages per rate_limit_interval.
<b>rate_limit_except_level = CRITICAL</b>	string value	Log level name used by rate limiting: CRITICAL, ERROR, INFO, WARNING, DEBUG or empty string. Logs with level greater or equal to rate_limit_except_level are not filtered. An empty string means that all levels are filtered.
<b>rate_limit_interval = 0</b>	integer value	Interval, number of seconds, of log rate limiting.
<b>syslog-log-facility = LOG_USER</b>	string value	Syslog facility to receive log lines. This option is ignored if log_config_append is set.
<b>use-journal = False</b>	boolean value	Enable journald for logging. If running in a systemd environment you may wish to enable journal support. Doing so will use the journal native protocol which includes structured metadata in addition to log messages. This option is ignored if log_config_append is set.
<b>use-json = False</b>	boolean value	Use JSON formatting for logging. This option is ignored if log_config_append is set.
<b>use-syslog = False</b>	boolean value	Use syslog for logging. Existing syslog format is DEPRECATED and will be changed later to honor RFC5424. This option is ignored if log_config_append is set.
<b>use_stderr = False</b>	boolean value	Log output to standard error. This option is ignored if log_config_append is set.
<b>watch-log-file = False</b>	boolean value	Uses logging handler designed to watch file system. When log file is moved or removed this handler will open a new log file with specified path instantaneously. It makes sense only if log_file option is specified and Linux platform is used. This option is ignored if log_config_append is set.

## 10.4.2. agent

The following table outlines the options available under the **[agent]** group in the `/etc/neutron/metadata_agent.ini` file.

Table 10.14. agent

Configuration option = Default value	Type	Description
<b>log_agent_heartbeats = False</b>	boolean value	Log agent heartbeats
<b>report_interval = 30</b>	floating point value	Seconds between nodes reporting state to server; should be less than agent_down_time, best if it is half or less than agent_down_time.

## 10.4.3. cache

The following table outlines the options available under the **[cache]** group in the `/etc/neutron/metadata_agent.ini` file.

Table 10.15. cache

Configuration option = Default value	Type	Description
<b>backend = dogpile.cache.null</b>	string value	Cache backend module. For eventlet-based or environments with hundreds of threaded servers, Memcache with pooling ( <code>oslo_cache.memcache_pool</code> ) is recommended. For environments with less than 100 threaded servers, Memcached ( <code>dogpile.cache.memcached</code> ) or Redis ( <code>dogpile.cache.redis</code> ) is recommended. Test environments with a single instance of the server can use the <code>dogpile.cache.memory</code> backend.
<b>backend_argument = []</b>	multi valued	Arguments supplied to the backend module. Specify this option once per argument to be passed to the <code>dogpile.cache</code> backend. Example format: " <code>&lt;argname&gt;: &lt;value&gt;</code> ".
<b>config_prefix = cache.oslo</b>	string value	Prefix for building the configuration dictionary for the cache region. This should not need to be changed unless there is another <code>dogpile.cache</code> region with the same configuration name.
<b>debug_cache_backend = False</b>	boolean value	Extra debugging from the cache backend (cache keys, get/set/delete/etc calls). This is only really useful if you need to see the specific cache-backend get/set/delete calls with the keys/values. Typically this should be left set to false.

Configuration option = Default value	Type	Description
<b>enabled = False</b>	boolean value	Global toggle for caching.
<b>expiration_time = 600</b>	integer value	Default TTL, in seconds, for any cached item in the dogpile.cache region. This applies to any cached method that doesn't have an explicit cache expiration time defined for it.
<b>memcache_dead_retry = 300</b>	integer value	Number of seconds memcached server is considered dead before it is tried again. (dogpile.cache.memcache and oslo_cache.memcache_pool backends only).
<b>memcache_pool_connection_get_timeout = 10</b>	integer value	Number of seconds that an operation will wait to get a memcache client connection.
<b>memcache_pool_maxsize = 10</b>	integer value	Max total number of open connections to every memcached server. (oslo_cache.memcache_pool backend only).
<b>memcache_pool_unused_timeout = 60</b>	integer value	Number of seconds a connection to memcached is held unused in the pool before it is closed. (oslo_cache.memcache_pool backend only).
<b>memcache_servers = [u'localhost:11211']</b>	list value	Memcache servers in the format of "host:port". (dogpile.cache.memcache and oslo_cache.memcache_pool backends only).
<b>memcache_socket_timeout = 3</b>	integer value	Timeout in seconds for every call to a server. (dogpile.cache.memcache and oslo_cache.memcache_pool backends only).
<b>proxies = []</b>	list value	Proxy classes to import that will affect the way the dogpile.cache backend functions. See the dogpile.cache documentation on changing-backend-behavior.

## 10.5. METERING\_AGENT.INI

This section contains options for the `/etc/neutron/metering_agent.ini` file.

### 10.5.1. DEFAULT

The following table outlines the options available under the **[DEFAULT]** group in the `/etc/neutron/metering_agent.ini` file.

Table 10.16. DEFAULT

Configuration option = Default value	Type	Description
<b>debug = False</b>	boolean value	If set to true, the logging level will be set to DEBUG instead of the default INFO level.
<b>default_log_levels =</b> [u'amqp=WARN', u'amqpplib=WARN', u'boto=WARN', u'qpid=WARN', u'sqlalchemy=WARN', u'suds=INFO', u'oslo.messaging=INFO', u'oslo_messaging=INFO', u'iso8601=WARN', u'requests.packages.urlli b3.connectionpool=WAR N', u'urllib3.connectionpool= WARN', u'websocket=WARN', u'requests.packages.urlli b3.util.retry=WARN', u'urllib3.util.retry=WARN', u'keystonemiddleware=W ARN', u'routes.middleware=WA RN', u'stevedore=WARN', u'taskflow=WARN', u'keystoneauth=WARN', u'oslo.cache=INFO', u'dogpile.core.dogpile=IN FO']	list value	List of package logging levels in logger=LEVEL pairs. This option is ignored if log_config_append is set.
<b>driver =</b> neutron.services.meterin g.drivers.noop.noop_driv er.NoopMeteringDriver	string value	Metering driver
<b>fatal_deprecations =</b> False	boolean value	Enables or disables fatal status of deprecations.
instance_format = [instance: %(uuid)s]`	string value	The format for an instance that is passed with the log message.
instance_uuid_format = [instance: %(uuid)s]`	string value	The format for an instance UUID that is passed with the log message.
<b>interface_driver = None</b>	string value	The driver used to manage the virtual interface.

Configuration option = Default value	Type	Description
<b>log-config-append = None</b>	string value	The name of a logging configuration file. This file is appended to any existing logging configuration files. For details about logging configuration files, see the Python logging module documentation. Note that when logging configuration files are used then all logging configuration is set in the configuration file and other logging configuration options are ignored (for example, <code>logging_context_format_string</code> ).
<b>log-date-format = %Y-%m-%d %H:%M:%S</b>	string value	Defines the format string for <code>%(asctime)s</code> in log records. Default: <code>%(default)s</code> . This option is ignored if <code>log_config_append</code> is set.
<b>log-dir = None</b>	string value	(Optional) The base directory used for relative <code>log_file</code> paths. This option is ignored if <code>log_config_append</code> is set.
<b>log-file = None</b>	string value	(Optional) Name of log file to send logging output to. If no default is set, logging will go to <code>stderr</code> as defined by <code>use_stderr</code> . This option is ignored if <code>log_config_append</code> is set.
<b>logging_context_format_string = %(asctime)s.% (msecs)03d %(process)d %(levelname)s %(name)s [% (request_id)s % (user_identity)s] % (instance)s%(message)s</b>	string value	Format string to use for log messages with context.
<b>logging_debug_format_s uffix = %(funcName)s % (pathname)s:%(lineno)d</b>	string value	Additional data to append to log message when logging level for the message is <code>DEBUG</code> .
<b>logging_default_format_s tring = %(asctime)s.% (msecs)03d %(process)d %(levelname)s %(name)s [-] %(instance)s% (message)s</b>	string value	Format string to use for log messages when context is undefined.
<b>logging_exception_prefix = %(asctime)s.% (msecs)03d %(process)d ERROR %(name)s % (instance)s</b>	string value	Prefix each line of exception output with this format.

Configuration option = Default value	Type	Description
<b>logging_user_identity_format =</b> <code>%(user)s %(tenant)s %(domain)s %(user_domain)s %(project_domain)s</code>	string value	Defines the format string for <code>%(user_identity)s</code> that is used in <code>logging_context_format_string</code> .
<b>measure_interval =</b> <code>30</code>	integer value	Interval between two metering measures
<b>ovs_integration_bridge =</b> <code>br-int</code>	string value	Name of Open vSwitch bridge to use
<b>ovs_use_veth =</b> <code>False</code>	boolean value	Uses veth for an OVS interface or not. Support kernels with limited namespace support (e.g. RHEL 6.5) and rate limiting on router's gateway port so long as <code>ovs_use_veth</code> is set to <code>True</code> .
<b>publish_errors =</b> <code>False</code>	boolean value	Enables or disables publication of error events.
<b>rate_limit_burst =</b> <code>0</code>	integer value	Maximum number of logged messages per <code>rate_limit_interval</code> .
<b>rate_limit_except_level =</b> <code>CRITICAL</code>	string value	Log level name used by rate limiting: <code>CRITICAL</code> , <code>ERROR</code> , <code>INFO</code> , <code>WARNING</code> , <code>DEBUG</code> or empty string. Logs with level greater or equal to <code>rate_limit_except_level</code> are not filtered. An empty string means that all levels are filtered.
<b>rate_limit_interval =</b> <code>0</code>	integer value	Interval, number of seconds, of log rate limiting.
<b>report_interval =</b> <code>300</code>	integer value	Interval between two metering reports
<b>syslog-log-facility =</b> <code>LOG_USER</code>	string value	Syslog facility to receive log lines. This option is ignored if <code>log_config_append</code> is set.
<b>use-journal =</b> <code>False</code>	boolean value	Enable journald for logging. If running in a systemd environment you may wish to enable journal support. Doing so will use the journal native protocol which includes structured metadata in addition to log messages. This option is ignored if <code>log_config_append</code> is set.
<b>use-json =</b> <code>False</code>	boolean value	Use JSON formatting for logging. This option is ignored if <code>log_config_append</code> is set.

Configuration option = Default value	Type	Description
<b>use-syslog = False</b>	boolean value	Use syslog for logging. Existing syslog format is DEPRECATED and will be changed later to honor RFC5424. This option is ignored if log_config_append is set.
<b>use_stderr = False</b>	boolean value	Log output to standard error. This option is ignored if log_config_append is set.
<b>watch-log-file = False</b>	boolean value	Uses logging handler designed to watch file system. When log file is moved or removed this handler will open a new log file with specified path instantaneously. It makes sense only if log_file option is specified and Linux platform is used. This option is ignored if log_config_append is set.

### 10.5.2. agent

The following table outlines the options available under the **[agent]** group in the `/etc/neutron/metering_agent.ini` file.

Table 10.17. agent

Configuration option = Default value	Type	Description
<b>log_agent_heartbeats = False</b>	boolean value	Log agent heartbeats
<b>report_interval = 30</b>	floating point value	Seconds between nodes reporting state to server; should be less than agent_down_time, best if it is half or less than agent_down_time.

### 10.5.3. ovs

The following table outlines the options available under the **[ovs]** group in the `/etc/neutron/metering_agent.ini` file.

Table 10.18. ovs

Configuration option = Default value	Type	Description
<b>bridge_mac_table_size = 50000</b>	integer value	The maximum number of MAC addresses to learn on a bridge managed by the Neutron OVS agent. Values outside a reasonable range (10 to 1,000,000) might be overridden by Open vSwitch according to the documentation.



Configuration option = Default value	Type	Description
<b>ovsdb_connection = tcp:127.0.0.1:6640</b>	string value	The connection string for the OVSDB backend. Will be used by ovsdb-client when monitoring and used for the all ovsdb commands when native ovsdb_interface is enabled
<b>ovsdb_interface = native</b>	string value	The interface for interacting with the OVSDB
<b>ovsdb_timeout = 10</b>	integer value	Timeout in seconds for ovsdb commands. If the timeout expires, ovsdb commands will fail with ALARMCLOCK error.
<b>ssl_ca_cert_file = None</b>	string value	The Certificate Authority (CA) certificate to use when interacting with OVSDB. Required when using an "ssl:" prefixed ovsdb_connection
<b>ssl_cert_file = None</b>	string value	The SSL certificate file to use when interacting with OVSDB. Required when using an "ssl:" prefixed ovsdb_connection
<b>ssl_key_file = None</b>	string value	The SSL private key file to use when interacting with OVSDB. Required when using an "ssl:" prefixed ovsdb_connection

## 10.6. ML2\_CONF.INI

This section contains options for the `/etc/neutron/plugins/ml2/ml2_conf.ini` file.

### 10.6.1. DEFAULT

The following table outlines the options available under the **[DEFAULT]** group in the `/etc/neutron/plugins/ml2/ml2_conf.ini` file.

Table 10.19. DEFAULT

Configuration option = Default value	Type	Description
<b>debug = False</b>	boolean value	If set to true, the logging level will be set to DEBUG instead of the default INFO level.

Configuration option = Default value	Type	Description
<b>default_log_levels =</b> [u'amqp=WARN', u'amqplib=WARN', u'boto=WARN', u'qpid=WARN', u'sqlalchemy=WARN', u'suds=INFO', u'oslo.messaging=INFO', u'oslo_messaging=INFO', u'iso8601=WARN', u'requests.packages.urlli b3.connectionpool=WAR N', u'urllib3.connectionpool= WARN', u'websocket=WARN', u'requests.packages.urlli b3.util.retry=WARN', u'urllib3.util.retry=WARN', u'keystonemiddleware=W ARN', u'routes.middleware=WA RN', u'stevedore=WARN', u'taskflow=WARN', u'keystoneauth=WARN', u'oslo.cache=INFO', u'dogpile.core.dogpile=IN FO']	list value	List of package logging levels in logger=LEVEL pairs. This option is ignored if log_config_append is set.
<b>fatal_deprecations =</b> <b>False</b>	boolean value	Enables or disables fatal status of deprecations.
<code>`instance_format = [instance: %(uuid)s] `</code>	string value	The format for an instance that is passed with the log message.
<code>`instance_uuid_format = [instance:%(uuid)s] `</code>	string value	The format for an instance UUID that is passed with the log message.
<b>log-config-append = None</b>	string value	The name of a logging configuration file. This file is appended to any existing logging configuration files. For details about logging configuration files, see the Python logging module documentation. Note that when logging configuration files are used then all logging configuration is set in the configuration file and other logging configuration options are ignored (for example, logging_context_format_string).

Configuration option = Default value	Type	Description
<b>log-date-format = %Y-%m-%d %H:%M:%S</b>	string value	Defines the format string for <code>%(asctime)s</code> in log records. Default: <code>%(default)s</code> . This option is ignored if <code>log_config_append</code> is set.
<b>log-dir = None</b>	string value	(Optional) The base directory used for relative <code>log_file</code> paths. This option is ignored if <code>log_config_append</code> is set.
<b>log-file = None</b>	string value	(Optional) Name of log file to send logging output to. If no default is set, logging will go to <code>stderr</code> as defined by <code>use_stderr</code> . This option is ignored if <code>log_config_append</code> is set.
<b>logging_context_format_string = %(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [%(request_id)s %(user_identity)s] %(instance)s%(message)s</b>	string value	Format string to use for log messages with context.
<b>logging_debug_format_sufffix = %(funcName)s %(pathname)s:%(lineno)d</b>	string value	Additional data to append to log message when logging level for the message is <code>DEBUG</code> .
<b>logging_default_format_string = %(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [-] %(instance)s%(message)s</b>	string value	Format string to use for log messages when context is undefined.
<b>logging_exception_prefix = %(asctime)s.%(msecs)03d %(process)d ERROR %(name)s %(instance)s</b>	string value	Prefix each line of exception output with this format.
<b>logging_user_identity_format = %(user)s %(tenant)s %(domain)s %(user_domain)s %(project_domain)s</b>	string value	Defines the format string for <code>%(user_identity)s</code> that is used in <code>logging_context_format_string</code> .
<b>publish_errors = False</b>	boolean value	Enables or disables publication of error events.

Configuration option = Default value	Type	Description
<b>rate_limit_burst = 0</b>	integer value	Maximum number of logged messages per rate_limit_interval.
<b>rate_limit_except_level = CRITICAL</b>	string value	Log level name used by rate limiting: CRITICAL, ERROR, INFO, WARNING, DEBUG or empty string. Logs with level greater or equal to rate_limit_except_level are not filtered. An empty string means that all levels are filtered.
<b>rate_limit_interval = 0</b>	integer value	Interval, number of seconds, of log rate limiting.
<b>syslog-log-facility = LOG_USER</b>	string value	Syslog facility to receive log lines. This option is ignored if log_config_append is set.
<b>use-journal = False</b>	boolean value	Enable journald for logging. If running in a systemd environment you may wish to enable journal support. Doing so will use the journal native protocol which includes structured metadata in addition to log messages. This option is ignored if log_config_append is set.
<b>use-json = False</b>	boolean value	Use JSON formatting for logging. This option is ignored if log_config_append is set.
<b>use-syslog = False</b>	boolean value	Use syslog for logging. Existing syslog format is DEPRECATED and will be changed later to honor RFC5424. This option is ignored if log_config_append is set.
<b>use_stderr = False</b>	boolean value	Log output to standard error. This option is ignored if log_config_append is set.
<b>watch-log-file = False</b>	boolean value	Uses logging handler designed to watch file system. When log file is moved or removed this handler will open a new log file with specified path instantaneously. It makes sense only if log_file option is specified and Linux platform is used. This option is ignored if log_config_append is set.

### 10.6.2. l2pop

The following table outlines the options available under the **[l2pop]** group in the `/etc/neutron/plugins/ml2/ml2_conf.ini` file.

Table 10.20. l2pop

Configuration option = Default value	Type	Description
<b>agent_boot_time = 180</b>	integer value	Delay within which agent is expected to update existing ports when it restarts

### 10.6.3. ml2

The following table outlines the options available under the **[ml2]** group in the **/etc/neutron/plugins/ml2/ml2\_conf.ini** file.

Table 10.21. ml2

Configuration option = Default value	Type	Description
<b>extension_drivers = []</b>	list value	An ordered list of extension driver entrypoints to be loaded from the neutron.ml2.extension_drivers namespace. For example: extension_drivers = port_security,qos
<b>external_network_type = None</b>	string value	Default network type for external networks when no provider attributes are specified. By default it is None, which means that if provider attributes are not specified while creating external networks then they will have the same type as tenant networks. Allowed values for external_network_type config option depend on the network type values configured in type_drivers config option.
<b>mechanism_drivers = []</b>	list value	An ordered list of networking mechanism driver entrypoints to be loaded from the neutron.ml2.mechanism_drivers namespace.
<b>overlay_ip_version = 4</b>	integer value	IP version of all overlay (tunnel) network endpoints. Use a value of 4 for IPv4 or 6 for IPv6.
<b>path_mtu = 0</b>	integer value	Maximum size of an IP packet (MTU) that can traverse the underlying physical network infrastructure without fragmentation when using an overlay/tunnel protocol. This option allows specifying a physical network MTU value that differs from the default global_physnet_mtu value.
<b>physical_network_mtu = []</b>	list value	A list of mappings of physical networks to MTU values. The format of the mapping is <physnet>:<mtu val>. This mapping allows specifying a physical network MTU value that differs from the default global_physnet_mtu value.

Configuration option = Default value	Type	Description
<b>tenant_network_types = [u'local']</b>	list value	Ordered list of network_types to allocate as tenant networks. The default value <i>local</i> is useful for single-box testing but provides no connectivity between hosts.
<b>type_drivers = [u'local', u'flat', u'vlan', u'gre', u'vxlan', u'geneve']</b>	list value	List of network type driver entrypoints to be loaded from the neutron.ml2.type_drivers namespace.

#### 10.6.4. ml2\_type\_flat

The following table outlines the options available under the **[ml2\_type\_flat]** group in the `/etc/neutron/plugins/ml2/ml2_conf.ini` file.

Table 10.22. ml2\_type\_flat

Configuration option = Default value	Type	Description
<b>flat_networks = *</b>	list value	List of physical_network names with which flat networks can be created. Use default <i>*</i> to allow flat networks with arbitrary physical_network names. Use an empty list to disable flat networks.

#### 10.6.5. ml2\_type\_geneve

The following table outlines the options available under the **[ml2\_type\_geneve]** group in the `/etc/neutron/plugins/ml2/ml2_conf.ini` file.

Table 10.23. ml2\_type\_geneve

Configuration option = Default value	Type	Description
<b>max_header_size = 30</b>	integer value	Geneve encapsulation header size is dynamic, this value is used to calculate the maximum MTU for the driver. This is the sum of the sizes of the outer ETH + IP + UDP + GENEVE header sizes. The default size for this field is 50, which is the size of the Geneve header without any additional option headers.
<b>vni_ranges = []</b>	list value	Comma-separated list of <vni_min>:<vni_max> tuples enumerating ranges of Geneve VNI IDs that are available for tenant network allocation

### 10.6.6. ml2\_type\_gre

The following table outlines the options available under the **[ml2\_type\_gre]** group in the `/etc/neutron/plugins/ml2/ml2_conf.ini` file.

Table 10.24. ml2\_type\_gre

Configuration option = Default value	Type	Description
<b>tunnel_id_ranges = []</b>	list value	Comma-separated list of <tun_min>:<tun_max> tuples enumerating ranges of GRE tunnel IDs that are available for tenant network allocation

### 10.6.7. ml2\_type\_vlan

The following table outlines the options available under the **[ml2\_type\_vlan]** group in the `/etc/neutron/plugins/ml2/ml2_conf.ini` file.

Table 10.25. ml2\_type\_vlan

Configuration option = Default value	Type	Description
<b>network_vlan_ranges = []</b>	list value	List of <physical_network>:<vlan_min>:<vlan_max> or <physical_network> specifying physical_network names usable for VLAN provider and tenant networks, as well as ranges of VLAN tags on each available for allocation to tenant networks.

### 10.6.8. ml2\_type\_vxlan

The following table outlines the options available under the **[ml2\_type\_vxlan]** group in the `/etc/neutron/plugins/ml2/ml2_conf.ini` file.

Table 10.26. ml2\_type\_vxlan

Configuration option = Default value	Type	Description
<b>vni_ranges = []</b>	list value	Comma-separated list of <vni_min>:<vni_max> tuples enumerating ranges of VXLAN VNI IDs that are available for tenant network allocation
<b>vxlan_group = None</b>	string value	Multicast group for VXLAN. When configured, will enable sending all broadcast traffic to this multicast group. When left unconfigured, will disable multicast VXLAN mode.

### 10.6.9. securitygroup

The following table outlines the options available under the **[securitygroup]** group in the `/etc/neutron/plugins/ml2/ml2_conf.ini` file.

Table 10.27. securitygroup

Configuration option = Default value	Type	Description
<b>enable_ipset = True</b>	boolean value	Use ipset to speed-up the iptables based security groups. Enabling ipset support requires that ipset is installed on L2 agent node.
<b>enable_security_group = True</b>	boolean value	Controls whether the neutron security group API is enabled in the server. It should be false when using no security groups or using the nova security group API.
<b>firewall_driver = None</b>	string value	Driver for security groups firewall in the L2 agent

## 10.7. NEUTRON.CONF

This section contains options for the `/etc/neutron/neutron.conf` file.

### 10.7.1. DEFAULT

The following table outlines the options available under the **[DEFAULT]** group in the `/etc/neutron/neutron.conf` file.

Table 10.28. DEFAULT

Configuration option = Default value	Type	Description
<b>agent_down_time = 75</b>	integer value	Seconds to regard the agent is down; should be at least twice report_interval, to be sure the agent is down for good.
<b>allow_automatic_dhcp_failover = True</b>	boolean value	Automatically remove networks from offline DHCP agents.
<b>allow_automatic_l3agent_failover = False</b>	boolean value	Automatically reschedule routers from offline L3 agents to online L3 agents.
<b>allow_bulk = True</b>	boolean value	Allow the usage of the bulk API
<b>allow_overlapping_ips = False</b>	boolean value	Allow overlapping IP support in Neutron. Attention: the following parameter MUST be set to False if Neutron is being used in conjunction with Nova security groups.



Configuration option = Default value	Type	Description
<code>`api_extensions_path = `</code>	string value	The path for API extensions. Note that this can be a colon-separated list of paths. For example: <code>api_extensions_path = extensions:/path/to/more/exts:/even/more/exts.</code> The <i>path</i> of <code>neutron.extensions</code> is appended to this, so if your extensions are in there you don't need to specify them here.
<b>api_paste_config = api-paste.ini</b>	string value	File name for the <code>paste.deploy</code> config for api service
<b>api_workers = None</b>	integer value	Number of separate API worker processes for service. If not specified, the default is equal to the number of CPUs available for best performance.
<b>auth_strategy = keystone</b>	string value	The type of authentication to use
<b>backlog = 4096</b>	integer value	Number of backlog requests to configure the socket with
<b>base_mac = fa:16:3e:00:00:00</b>	string value	The base MAC address Neutron will use for VIFs. The first 3 octets will remain unchanged. If the 4th octet is not 00, it will also be used. The others will be randomly generated.
<b>bind_host = 0.0.0.0</b>	unknown value	The host IP to bind to.
<b>bind_port = 9696</b>	port value	The port to bind to
<b>client_socket_timeout = 900</b>	integer value	Timeout for client connections' socket operations. If an incoming connection is idle for this number of seconds it will be closed. A value of 0 means wait forever.
<b>conn_pool_min_size = 2</b>	integer value	The pool size limit for connections expiration policy
<b>conn_pool_ttl = 1200</b>	integer value	The time-to-live in sec of idle connections in the pool
<b>control_exchange = neutron</b>	string value	The default exchange under which topics are scoped. May be overridden by an exchange name specified in the <code>transport_url</code> option.
<b>core_plugin = None</b>	string value	The core plugin Neutron will use

Configuration option = Default value	Type	Description
<b>debug = False</b>	boolean value	If set to true, the logging level will be set to DEBUG instead of the default INFO level.
<b>default_availability_zones = []</b>	list value	Default value of availability zone hints. The availability zone aware schedulers use this when the resources availability_zone_hints is empty. Multiple availability zones can be specified by a comma separated string. This value can be empty. In this case, even if availability_zone_hints for a resource is empty, availability zone is considered for high availability while scheduling the resource.
<b>default_log_levels = [u'amqp=WARN', u'amqp-lib=WARN', u'boto=WARN', u'qpid=WARN', u'sqlalchemy=WARN', u'suds=INFO', u'oslo.messaging=INFO', u'oslo_messaging=INFO', u'iso8601=WARN', u'requests.packages.urllib3.connectionpool=WARN', u'urllib3.connectionpool=WARN', u'websocket=WARN', u'requests.packages.urllib3.util.retry=WARN', u'urllib3.util.retry=WARN', u'keystonemiddleware=WARN', u'routes.middleware=WARN', u'stevedore=WARN', u'taskflow=WARN', u'keystoneauth=WARN', u'oslo.cache=INFO', u'dogpile.core.dogpile=INFO']</b>	list value	List of package logging levels in logger=LEVEL pairs. This option is ignored if log_config_append is set.
<b>dhcp_agent_notification = True</b>	boolean value	Allow sending resource operation notification to DHCP agent

Configuration option = Default value	Type	Description
<b>dhcp_agents_per_network = 1</b>	integer value	Number of DHCP agents scheduled to host a tenant network. If this number is greater than 1, the scheduler automatically assigns multiple DHCP agents for a given tenant network, providing high availability for DHCP service.
<b>dhcp_lease_duration = 86400</b>	integer value	DHCP lease duration (in seconds). Use -1 to tell dnsmasq to use infinite lease times.
<b>dhcp_load_type = networks</b>	string value	Representing the resource type whose load is being reported by the agent. This can be "networks", "subnets" or "ports". When specified (Default is networks), the server will extract particular load sent as part of its agent configuration object from the agent report state, which is the number of resources being consumed, at every report_interval.dhcp_load_type can be used in combination with network_scheduler_driver = neutron.scheduler.dhcp_agent_scheduler.WeightScheduler When the network_scheduler_driver is WeightScheduler, dhcp_load_type can be configured to represent the choice for the resource being balanced. Example: dhcp_load_type=networks
<b>dns_domain = openstacklocal</b>	string value	Domain to use for building the hostnames
<b>dvr_base_mac = fa:16:3f:00:00:00</b>	string value	The base mac address used for unique DVR instances by Neutron. The first 3 octets will remain unchanged. If the 4th octet is not 00, it will also be used. The others will be randomly generated. The <i>dvr_base_mac</i> <b>must</b> be different from <i>base_mac</i> to avoid mixing them up with MAC's allocated for tenant ports. A 4 octet example would be <i>dvr_base_mac</i> = fa:16:3f:4f:00:00. The default is 3 octet
<b>enable_dvr = True</b>	boolean value	Determine if setup is configured for DVR. If False, DVR API extension will be disabled.
<b>enable_new_agents = True</b>	boolean value	Agent starts with admin_state_up=False when enable_new_agents=False. In the case, user's resources will not be scheduled automatically to the agent until admin changes admin_state_up to True.

Configuration option = Default value	Type	Description
<b>enable_services_on_agents_with_admin_state_down = False</b>	boolean value	Enable services on an agent with <code>admin_state_up</code> False. If this option is False, when <code>admin_state_up</code> of an agent is turned False, services on it will be disabled. Agents with <code>admin_state_up</code> False are not selected for automatic scheduling regardless of this option. But manual scheduling to such agents is available if this option is True.
<b>enable_snat_by_default = True</b>	boolean value	Define the default value of <code>enable_snat</code> if not provided in <code>external_gateway_info</code> .
<b>executor_thread_pool_size = 64</b>	integer value	Size of executor thread pool when executor is threading or eventlet.
<b>external_dns_driver = None</b>	string value	Driver for external DNS integration.
<b>fatal_deprecations = False</b>	boolean value	Enables or disables fatal status of deprecations.
<b>global_physnet_mtu = 1500</b>	integer value	MTU of the underlying physical network. Neutron uses this value to calculate MTU for all virtual network components. For flat and VLAN networks, neutron uses this value without modification. For overlay networks such as VXLAN, neutron automatically subtracts the overlay protocol overhead from this value. Defaults to 1500, the standard value for Ethernet.
<b>host = ccsosp-undercloud.usersys.redhat.com</b>	unknown value	Hostname to be used by the Neutron server, agents and services running on this machine. All the agents and services running on this machine must use the same host value.
<code>`instance_format = [instance: %(uuid)s] `</code>	string value	The format for an instance that is passed with the log message.
<code>`instance_uuid_format = [instance: %(uuid)s] `</code>	string value	The format for an instance UUID that is passed with the log message.
<b>interface_driver = None</b>	string value	The driver used to manage the virtual interface.
<b>ipam_driver = internal</b>	string value	Neutron IPAM (IP address management) driver to use. By default, the reference implementation of the Neutron IPAM driver is used.

Configuration option = Default value	Type	Description
<b>ipv6_pd_enabled = False</b>	boolean value	Enables IPv6 Prefix Delegation for automatic subnet CIDR allocation. Set to True to enable IPv6 Prefix Delegation for subnet allocation in a PD-capable environment. Users making subnet creation requests for IPv6 subnets without providing a CIDR or subnetpool ID will be given a CIDR via the Prefix Delegation mechanism. Note that enabling PD will override the behavior of the default IPv6 subnetpool.
<b>l3_ha = False</b>	boolean value	Enable HA mode for virtual routers.
<b>l3_ha_net_cidr = 169.254.192.0/18</b>	string value	Subnet used for the l3 HA admin network.
<code>l3_ha_network_physical_name = `</code>	string value	The physical network name with which the HA network can be created.
<code>l3_ha_network_type = `</code>	string value	The network type to use when creating the HA network for an HA router. By default or if empty, the first <i>tenant_network_types</i> is used. This is helpful when the VRRP traffic should use a specific network which is not the default one.
<b>log-config-append = None</b>	string value	The name of a logging configuration file. This file is appended to any existing logging configuration files. For details about logging configuration files, see the Python logging module documentation. Note that when logging configuration files are used then all logging configuration is set in the configuration file and other logging configuration options are ignored (for example, <code>logging_context_format_string</code> ).
<b>log-date-format = %Y-%m-%d %H:%M:%S</b>	string value	Defines the format string for <code>%(asctime)s</code> in log records. Default: <code>%(default)s</code> . This option is ignored if <code>log_config_append</code> is set.
<b>log-dir = None</b>	string value	(Optional) The base directory used for relative <code>log_file</code> paths. This option is ignored if <code>log_config_append</code> is set.
<b>log-file = None</b>	string value	(Optional) Name of log file to send logging output to. If no default is set, logging will go to <code>stderr</code> as defined by <code>use_stderr</code> . This option is ignored if <code>log_config_append</code> is set.

Configuration option = Default value	Type	Description
<b>logging_context_format_string</b> = <code>%(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [%(request_id)s %(user_identity)s] %(instance)s%(message)s</code>	string value	Format string to use for log messages with context.
<b>logging_debug_format_s uffix</b> = <code>%(funcName)s %(pathname)s:%(lineno)d</code>	string value	Additional data to append to log message when logging level for the message is DEBUG.
<b>logging_default_format_s tring</b> = <code>%(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [-] %(instance)s%(message)s</code>	string value	Format string to use for log messages when context is undefined.
<b>logging_exception_prefix</b> = <code>%(asctime)s.%(msecs)03d %(process)d ERROR %(name)s %(instance)s</code>	string value	Prefix each line of exception output with this format.
<b>logging_user_identity_for mat</b> = <code>%(user)s %(tenant)s %(domain)s %(user_domain)s %(project_domain)s</code>	string value	Defines the format string for <code>%(user_identity)s</code> that is used in <code>logging_context_format_string</code> .
<b>max_allowed_address_pa ir</b> = 10	integer value	Maximum number of allowed address pairs
<b>max_dns_nameservers</b> = 5	integer value	Maximum number of DNS nameservers per subnet
<b>max_header_line</b> = 16384	integer value	Maximum line size of message headers to be accepted. <code>max_header_line</code> may need to be increased when using large tokens (typically those generated when keystone is configured to use PKI tokens with big service catalogs).
<b>max_l3_agents_per_rout er</b> = 3	integer value	Maximum number of L3 agents which a HA router will be scheduled on. If it is set to 0 then the router will be scheduled on every agent.

Configuration option = Default value	Type	Description
<b>max_routes = 30</b>	integer value	Maximum number of routes per router
<b>max_subnet_host_routes = 20</b>	integer value	Maximum number of host routes per subnet
<code>`metadata_proxy_group = `</code>	string value	Group (gid or name) running metadata proxy after its initialization (if empty: agent effective group).
<b>metadata_proxy_socket = \$state_path/metadata_proxy</b>	string value	Location for Metadata Proxy UNIX domain socket.
<code>`metadata_proxy_user = `</code>	string value	User (uid or name) running metadata proxy after its initialization (if empty: agent effective user).
<b>network_auto_schedule = True</b>	boolean value	Allow auto scheduling networks to DHCP agent.
<b>network_link_prefix = None</b>	string value	This string is prepended to the normal URL that is returned in links to the OpenStack Network API. If it is empty (the default), the URLs are returned unchanged.
<b>network_scheduler_driver = neutron.scheduler.dhcp_agent_scheduler.WeightScheduler</b>	string value	Driver to use for scheduling network to DHCP agent
<b>notify_nova_on_port_data_changes = True</b>	boolean value	Send notification to nova when port data (fixed_ips/floatingip) changes so nova can update its cache.
<b>notify_nova_on_port_status_changes = True</b>	boolean value	Send notification to nova when port status changes
<b>pagination_max_limit = -1</b>	string value	The maximum number of items returned in a single response, value was <i>infinite</i> or negative integer means no limit
<b>periodic_fuzzy_delay = 5</b>	integer value	Range of seconds to randomly delay when starting the periodic task scheduler to reduce stampeding. (Disable by setting to 0)
<b>periodic_interval = 40</b>	integer value	Seconds between running periodic tasks.
<b>publish_errors = False</b>	boolean value	Enables or disables publication of error events.

Configuration option = Default value	Type	Description
<b>rate_limit_burst = 0</b>	integer value	Maximum number of logged messages per rate_limit_interval.
<b>rate_limit_except_level = CRITICAL</b>	string value	Log level name used by rate limiting: CRITICAL, ERROR, INFO, WARNING, DEBUG or empty string. Logs with level greater or equal to rate_limit_except_level are not filtered. An empty string means that all levels are filtered.
<b>rate_limit_interval = 0</b>	integer value	Interval, number of seconds, of log rate limiting.
<b>retry_until_window = 30</b>	integer value	Number of seconds to keep retrying to listen
<b>router_auto_schedule = True</b>	boolean value	Allow auto scheduling of routers to L3 agent.
<b>router_distributed = False</b>	boolean value	System-wide flag to determine the type of router that tenants can create. Only admin can override.
<b>router_scheduler_driver = neutron.scheduler.l3_agent_scheduler.LeastRoutersScheduler</b>	string value	Driver to use for scheduling router to a default L3 agent
<b>rpc_ack_timeout_base = 15</b>	integer value	Number of seconds to wait for an ack from a cast/call. After each retry attempt this timeout is multiplied by some specified multiplier.
<b>rpc_ack_timeout_multiplier = 2</b>	integer value	Number to multiply base ack timeout by after each retry attempt.
<b>rpc_backend = rabbit</b>	string value	The messaging driver to use, defaults to rabbit. Other drivers include amqp and zmq.
<b>rpc_conn_pool_size = 30</b>	integer value	Size of RPC connection pool.
<b>rpc_message_ttl = 300</b>	integer value	Expiration timeout in seconds of a sent/received message after which it is not tracked anymore by a client/server.
<b>rpc_poll_timeout = 1</b>	integer value	The default number of seconds that poll should wait. Poll raises timeout exception when timeout expired.
<b>rpc_response_timeout = 60</b>	integer value	Seconds to wait for a response from a call.



Configuration option = Default value	Type	Description
<b>rpc_retry_attempts = 3</b>	integer value	Default number of message sending attempts in case of any problems occurred: positive value N means at most N retries, 0 means no retries, None or -1 (or any other negative values) mean to retry forever. This option is used only if acknowledgments are enabled.
<b>rpc_state_report_workers = 1</b>	integer value	Number of RPC worker processes dedicated to state reports queue.
<b>rpc_thread_pool_size = 100</b>	integer value	Maximum number of (green) threads to work concurrently.
<b>rpc_use_acks = False</b>	boolean value	Wait for message acknowledgements from receivers. This mechanism works only via proxy without PUB/SUB.
<b>rpc_workers = 1</b>	integer value	Number of RPC worker processes for service.
<b>rpc_zmq_bind_address = *</b>	string value	ZeroMQ bind address. Should be a wildcard (*), an ethernet interface, or IP. The "host" option should point or resolve to this address.
<b>rpc_zmq_bind_port_retries = 100</b>	integer value	Number of retries to find free port number before fail with ZMQBindError.
<b>rpc_zmq_contexts = 1</b>	integer value	Number of ZeroMQ contexts, defaults to 1.
<b>rpc_zmq_host = ccosp-undercloud.usersys.redhat.com</b>	string value	Name of this node. Must be a valid hostname, FQDN, or IP address. Must match "host" option, if running Nova.
<b>rpc_zmq_ipc_dir = /var/run/openstack</b>	string value	Directory for holding IPC sockets.
<b>rpc_zmq_matchmaker = redis</b>	string value	MatchMaker driver.
<b>rpc_zmq_max_port = 65536</b>	integer value	Maximal port number for random ports range.
<b>rpc_zmq_min_port = 49153</b>	port value	Minimal port number for random ports range.
<b>rpc_zmq_serialization = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages

Configuration option = Default value	Type	Description
<b>rpc_zmq_topic_backlog = None</b>	integer value	Maximum number of ingress messages to locally buffer per topic. Default is unlimited.
<b>send_events_interval = 2</b>	integer value	Number of seconds between sending events to nova if there are any events to send.
<b>service_plugins = []</b>	list value	The service plugins Neutron will use
<b>state_path = /var/lib/neutron</b>	string value	Where to store Neutron state files. This directory must be writable by the agent.
<b>subscribe_on = []</b>	list value	List of publisher hosts SubConsumer can subscribe on. This option has higher priority then the default publishers list taken from the matchmaker.
<b>syslog-log-facility = LOG_USER</b>	string value	Syslog facility to receive log lines. This option is ignored if log_config_append is set.
<b>tcp_keepidle = 600</b>	integer value	Sets the value of TCP_KEEPIDLE in seconds for each server socket. Not supported on OS X.
<b>transport_url = None</b>	string value	<p>The network address and optional user credentials for connecting to the messaging backend, in URL format. The expected format is:</p> <pre>driver://[user:pass@]host:port[, [userN:passN@]hostN:portN]/virtual_host?query</pre> <p>Example: rabbit://rabbitmq:password@127.0.0.1:5672//</p> <p>For full details on the fields in the URL see the documentation of oslo_messaging.TransportURL at <a href="https://docs.openstack.org/oslo.messaging/latest/reference/transport.html">https://docs.openstack.org/oslo.messaging/latest/reference/transport.html</a></p>
<b>use-journal = False</b>	boolean value	Enable journald for logging. If running in a systemd environment you may wish to enable journal support. Doing so will use the journal native protocol which includes structured metadata in addition to log messages. This option is ignored if log_config_append is set.
<b>use-json = False</b>	boolean value	Use JSON formatting for logging. This option is ignored if log_config_append is set.

Configuration option = Default value	Type	Description
<b>use-syslog = False</b>	boolean value	Use syslog for logging. Existing syslog format is DEPRECATED and will be changed later to honor RFC5424. This option is ignored if log_config_append is set.
<b>use_dynamic_connections = False</b>	boolean value	This option makes direct connections dynamic or static. It makes sense only with use_router_proxy=False which means to use direct connections for direct message types (ignored otherwise).
<b>use_pub_sub = False</b>	boolean value	Use PUB/SUB pattern for fanout methods. PUB/SUB always uses proxy.
<b>use_router_proxy = False</b>	boolean value	Use ROUTER remote proxy.
<b>use_ssl = False</b>	boolean value	Enable SSL on the API server
<b>use_stderr = False</b>	boolean value	Log output to standard error. This option is ignored if log_config_append is set.
<b>vlan_transparent = False</b>	boolean value	If True, then allow plugins that support it to create VLAN transparent networks.
<b>watch-log-file = False</b>	boolean value	Uses logging handler designed to watch file system. When log file is moved or removed this handler will open a new log file with specified path instantaneously. It makes sense only if log_file option is specified and Linux platform is used. This option is ignored if log_config_append is set.
<b>wsgi_default_pool_size = 100</b>	integer value	Size of the pool of greenthreads used by wsgi
<b>wsgi_keep_alive = True</b>	boolean value	If False, closes the client socket connection explicitly.
<b>wsgi_log_format = %(client_ip)s "%(request_line)s" status: %(status_code)s len: %(body_length)s time: %(wall_seconds).7f</b>	string value	A python format string that is used as the template to generate log lines. The following values can be formatted into it: client_ip, date_time, request_line, status_code, body_length, wall_seconds.
<b>zmq_failover_connections = 2</b>	integer value	How many additional connections to a host will be made for failover reasons. This option is actual only in dynamic connections mode.

Configuration option = Default value	Type	Description
<b>zmq_immediate = True</b>	boolean value	This option configures round-robin mode in zmq socket. True means not keeping a queue when server side disconnects. False means to keep queue and messages even if server is disconnected, when the server appears we send all accumulated messages to it.
<b>zmq_linger = -1</b>	integer value	Number of seconds to wait before all pending messages will be sent after closing a socket. The default value of -1 specifies an infinite linger period. The value of 0 specifies no linger period. Pending messages shall be discarded immediately when the socket is closed. Positive values specify an upper bound for the linger period.
<b>zmq_target_expire = 300</b>	integer value	Expiration timeout in seconds of a name service record about existing target (< 0 means no timeout).
<b>zmq_target_update = 180</b>	integer value	Update period in seconds of a name service record about existing target.
<b>zmq_tcp_keepalive = -1</b>	integer value	Enable/disable TCP keepalive (KA) mechanism. The default value of -1 (or any other negative value) means to skip any overrides and leave it to OS default; 0 and 1 (or any other positive value) mean to disable and enable the option respectively.
<b>zmq_tcp_keepalive_cnt = -1</b>	integer value	The number of retransmissions to be carried out before declaring that remote end is not available. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_idle = -1</b>	integer value	The duration between two keepalive transmissions in idle condition. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_intvl = -1</b>	integer value	The duration between two successive keepalive retransmissions, if acknowledgement to the previous keepalive transmission is not received. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.

## 10.7.2. agent

The following table outlines the options available under the **[agent]** group in the `/etc/neutron/neutron.conf` file.

Table 10.29. agent

Configuration option = Default value	Type	Description
<b>availability_zone = nova</b>	string value	Availability zone of this node
<b>check_child_processes_action = respawn</b>	string value	Action to be executed when a child process dies
<b>check_child_processes_interval = 60</b>	integer value	Interval between checks of child process liveness (seconds), use 0 to disable
<b>comment iptables_rules = True</b>	boolean value	Add comments to iptables rules. Set to false to disallow the addition of comments to generated iptables rules that describe each rule's purpose. System must support the iptables comments module for addition of comments.
<b>debug iptables_rules = False</b>	boolean value	Duplicate every iptables difference calculation to ensure the format being generated matches the format of iptables-save. This option should not be turned on for production systems because it imposes a performance penalty.
<b>log_agent_heartbeats = False</b>	boolean value	Log agent heartbeats
<b>report_interval = 30</b>	floating point value	Seconds between nodes reporting state to server; should be less than agent_down_time, best if it is half or less than agent_down_time.
<b>root_helper = sudo</b>	string value	Root helper application. Use <code>sudo neutron-rootwrap /etc/neutron/rootwrap.conf</code> to use the real root filter facility. Change to <code>sudo</code> to skip the filtering and just run the command directly.
<b>root_helper_daemon = None</b>	string value	Root helper daemon application to use when possible. For the agent which needs to execute commands in Dom0 in the hypervisor of XenServer, this item should be set to <code>xenapi_root_helper</code> , so that it will keep a XenAPI session to pass commands to Dom0.

Configuration option = Default value	Type	Description
<b>use_helper_for_ns_read = True</b>	boolean value	Use the root helper when listing the namespaces on a system. This may not be required depending on the security configuration. If the root helper is not required, set this to False for a performance improvement.

### 10.7.3. cors

The following table outlines the options available under the **[cors]** group in the `/etc/neutron/neutron.conf` file.

Table 10.30. cors

Configuration option = Default value	Type	Description
<b>allow_credentials = True</b>	boolean value	Indicate that the actual request can include user credentials
<b>allow_headers = [u'X-Auth-Token', u'X-Identity-Status', u'X-Roles', u'X-Service-Catalog', u'X-User-Id', u'X-Tenant-Id', u'X-OpenStack-Request-ID']</b>	list value	Indicate which header field names may be used during the actual request.
<b>allow_methods = [u'GET', u'PUT', u'POST', u'DELETE', u'PATCH']</b>	list value	Indicate which methods can be used during the actual request.
<b>allowed_origin = None</b>	list value	Indicate whether this resource may be shared with the domain received in the requests "origin" header. Format: "<protocol>://<host>[:<port>]", no trailing slash. Example: <a href="https://horizon.example.com">https://horizon.example.com</a>
<b>expose_headers = [u'X-Auth-Token', u'X-Subject-Token', u'X-Service-Token', u'X-OpenStack-Request-ID', u'OpenStack-Volume-microversion']</b>	list value	Indicate which headers are safe to expose to the API. Defaults to HTTP Simple Headers.
<b>max_age = 3600</b>	integer value	Maximum cache age of CORS preflight requests.

## 10.7.4. database

The following table outlines the options available under the **[database]** group in the `/etc/neutron/neutron.conf` file.

Table 10.31. database

Configuration option = Default value	Type	Description
<b>backend = sqlalchemy</b>	string value	The back end to use for the database.
<b>connection = None</b>	string value	The SQLAlchemy connection string to use to connect to the database.
<b>connection_debug = 0</b>	integer value	Verbosity of SQL debugging information: 0=None, 100=Everything.
<b>connection_recycle_time = 3600</b>	integer value	Connections which have been present in the connection pool longer than this number of seconds will be replaced with a new one the next time they are checked out from the pool.
<b>connection_trace = False</b>	boolean value	Add Python stack traces to SQL as comment strings.
<b>db_inc_retry_interval = True</b>	boolean value	If True, increases the interval between retries of a database operation up to <code>db_max_retry_interval</code> .
<b>db_max_retries = 20</b>	integer value	Maximum retries in case of connection error or deadlock error before error is raised. Set to -1 to specify an infinite retry count.
<b>db_max_retry_interval = 10</b>	integer value	If <code>db_inc_retry_interval</code> is set, the maximum seconds between retries of a database operation.
<b>db_retry_interval = 1</b>	integer value	Seconds between retries of a database transaction.
<code>engine = `</code>	string value	Database engine for which script will be generated when using offline migration.
<b>max_overflow = 50</b>	integer value	If set, use this value for <code>max_overflow</code> with SQLAlchemy.
<b>max_pool_size = 5</b>	integer value	Maximum number of SQL connections to keep open in a pool. Setting a value of 0 indicates no limit.
<b>max_retries = 10</b>	integer value	Maximum number of database connection retries during startup. Set to -1 to specify an infinite retry count.

Configuration option = Default value	Type	Description
<b>min_pool_size = 1</b>	integer value	Minimum number of SQL connections to keep open in a pool.
<b>mysql_enable_ndb = False</b>	boolean value	If True, transparently enables support for handling MySQL Cluster (NDB).
<b>mysql_sql_mode = TRADITIONAL</b>	string value	The SQL mode to be used for MySQL sessions. This option, including the default, overrides any server-set SQL mode. To use whatever SQL mode is set by the server configuration, set this to no value. Example: <code>mysql_sql_mode=</code>
<b>pool_timeout = None</b>	integer value	If set, use this value for <code>pool_timeout</code> with SQLAlchemy.
<b>retry_interval = 10</b>	integer value	Interval between retries of opening a SQL connection.
<b>slave_connection = None</b>	string value	The SQLAlchemy connection string to use to connect to the slave database.
<b>sqlite_synchronous = True</b>	boolean value	If True, SQLite uses synchronous mode.
<b>use_db_reconnect = False</b>	boolean value	Enable the experimental use of database reconnect on connection lost.

### 10.7.5. keystone\_authtoken

The following table outlines the options available under the **[keystone\_authtoken]** group in the `/etc/neutron/neutron.conf` file.

Table 10.32. keystone\_authtoken

Configuration option = Default value	Type	Description
<b>auth_section = None</b>	string value	Config Section from which to load plugin specific options
<b>auth_type = None</b>	string value	Authentication type to load



Configuration option = Default value	Type	Description
<b>auth_uri = None</b>	string value	Complete "public" Identity API endpoint. This endpoint should not be an "admin" endpoint, as it should be accessible by all end users. Unauthenticated clients are redirected to this endpoint to authenticate. Although this endpoint should ideally be unversioned, client support in the wild varies. If you're using a versioned v2 endpoint here, then this should <b>not</b> be the same endpoint the service user utilizes for validating tokens, because normal end users may not be able to reach that endpoint. This option is deprecated in favor of <code>www_authenticate_uri</code> and will be removed in the S release.
<b>auth_version = None</b>	string value	API version of the admin Identity API endpoint.
<b>cache = None</b>	string value	Request environment key where the Swift cache object is stored. When <code>auth_token</code> middleware is deployed with a Swift cache, use this option to have the middleware share a caching backend with swift. Otherwise, use the <b>memcached_servers</b> option instead.
<b>cafile = None</b>	string value	A PEM encoded Certificate Authority to use when verifying HTTPs connections. Defaults to system CAs.
<b>certfile = None</b>	string value	Required if identity server requires client certificate
<b>check_revocations_for_cached = False</b>	boolean value	If true, the revocation list will be checked for cached tokens. This requires that PKI tokens are configured on the identity server.
<b>delay_auth_decision = False</b>	boolean value	Do not handle authorization requests within the middleware, but delegate the authorization decision to downstream WSGI components.
<b>enforce_token_bind = permissive</b>	string value	Used to control the use and type of token binding. Can be set to: "disabled" to not check token binding. "permissive" (default) to validate binding information if the bind type is of a form known to the server and ignore it if not. "strict" like "permissive" but if the bind type is unknown the token will be rejected. "required" any form of token binding is needed to be allowed. Finally the name of a binding method that must be present in tokens.

Configuration option = Default value	Type	Description
<b>hash_algorithms = [u'md5']</b>	list value	Hash algorithms to use for hashing PKI tokens. This may be a single algorithm or multiple. The algorithms are those supported by Python standard <code>hashlib.new()</code> . The hashes will be tried in the order given, so put the preferred one first for performance. The result of the first hash will be stored in the cache. This will typically be set to multiple values only while migrating from a less secure algorithm to a more secure one. Once all the old tokens are expired this option should be set to a single value for better performance.
<b>http_connect_timeout = None</b>	integer value	Request timeout value for communicating with Identity API server.
<b>http_request_max_retries = 3</b>	integer value	How many times are we trying to reconnect when communicating with Identity API Server.
<b>include_service_catalog = True</b>	boolean value	(Optional) Indicate whether to set the X-Service-Catalog header. If False, middleware will not ask for service catalog on token validation and will not set the X-Service-Catalog header.
<b>insecure = False</b>	boolean value	Verify HTTPS connections.
<b>keyfile = None</b>	string value	Required if identity server requires client certificate
<b>memcache_pool_connection_timeout = 10</b>	integer value	(Optional) Number of seconds that an operation will wait to get a memcached client connection from the pool.
<b>memcache_pool_dead_retry = 300</b>	integer value	(Optional) Number of seconds memcached server is considered dead before it is tried again.
<b>memcache_pool_maxsize = 10</b>	integer value	(Optional) Maximum total number of open connections to every memcached server.
<b>memcache_pool_socket_timeout = 3</b>	integer value	(Optional) Socket timeout in seconds for communicating with a memcached server.
<b>memcache_pool_unused_timeout = 60</b>	integer value	(Optional) Number of seconds a connection to memcached is held unused in the pool before it is closed.
<b>memcache_secret_key = None</b>	string value	(Optional, mandatory if <code>memcache_security_strategy</code> is defined) This string is used for key derivation.

Configuration option = Default value	Type	Description
<b>memcache_security_strategy = None</b>	string value	(Optional) If defined, indicate whether token data should be authenticated or authenticated and encrypted. If MAC, token data is authenticated (with HMAC) in the cache. If ENCRYPT, token data is encrypted and authenticated in the cache. If the value is not one of these options or empty, auth_token will raise an exception on initialization.
<b>memcache_use_advanced_pool = False</b>	boolean value	(Optional) Use the advanced (eventlet safe) memcached client pool. The advanced pool will only work under python 2.x.
<b>memcached_servers = None</b>	list value	Optionally specify a list of memcached server(s) to use for caching. If left undefined, tokens will instead be cached in-process.
<b>region_name = None</b>	string value	The region in which the identity server can be found.
<b>revocation_cache_time = 10</b>	integer value	Determines the frequency at which the list of revoked tokens is retrieved from the Identity service (in seconds). A high number of revocation events combined with a low cache duration may significantly reduce performance. Only valid for PKI tokens. This option has been deprecated in the Ocata release and will be removed in the P release.
<b>service_token_roles = [u'service']</b>	list value	A choice of roles that must be present in a service token. Service tokens are allowed to request that an expired token can be used and so this check should tightly control that only actual services should be sending this token. Roles here are applied as an ANY check so any role in this list must be present. For backwards compatibility reasons this currently only affects the allow_expired check.
<b>service_token_roles_required = False</b>	boolean value	For backwards compatibility reasons we must let valid service tokens pass that don't pass the service_token_roles check as valid. Setting this true will become the default in a future release and should be enabled if possible.
<b>signing_dir = None</b>	string value	Directory used to cache files related to PKI tokens. This option has been deprecated in the Ocata release and will be removed in the P release.
<b>token_cache_time = 300</b>	integer value	In order to prevent excessive effort spent validating tokens, the middleware caches previously-seen tokens for a configurable duration (in seconds). Set to -1 to disable caching completely.

Configuration option = Default value	Type	Description
<b>www_authenticate_uri = None</b>	string value	Complete "public" Identity API endpoint. This endpoint should not be an "admin" endpoint, as it should be accessible by all end users. Unauthenticated clients are redirected to this endpoint to authenticate. Although this endpoint should ideally be unversioned, client support in the wild varies. If you're using a versioned v2 endpoint here, then this should <b>not</b> be the same endpoint the service user utilizes for validating tokens, because normal end users may not be able to reach that endpoint.

### 10.7.6. matchmaker\_redis

The following table outlines the options available under the **[matchmaker\_redis]** group in the **/etc/neutron/neutron.conf** file.

Table 10.33. matchmaker\_redis

Configuration option = Default value	Type	Description
<b>check_timeout = 20000</b>	integer value	Time in ms to wait before the transaction is killed.
<b>host = 127.0.0.1</b>	string value	Host to locate redis.
<code>`password = `</code>	string value	Password for Redis server (optional).
<b>port = 6379</b>	port value	Use this port to connect to redis host.
<b>sentinel_group_name = oslo-messaging-zeromq</b>	string value	Redis replica set name.
<b>sentinel_hosts = []</b>	list value	List of Redis Sentinel hosts (fault tolerance mode), e.g., [host:port, host1:port ... ]
<b>socket_timeout = 10000</b>	integer value	Timeout in ms on blocking socket operations.
<b>wait_timeout = 2000</b>	integer value	Time in ms to wait between connection attempts.

### 10.7.7. nova

The following table outlines the options available under the **[nova]** group in the **/etc/neutron/neutron.conf** file.

Table 10.34. nova

Configuration option = Default value	Type	Description
<b>auth-url = None</b>	string value	Authentication URL
<b>auth_type = None</b>	string value	Authentication type to load
<b>cafile = None</b>	string value	PEM encoded Certificate Authority to use when verifying HTTPs connections.
<b>certfile = None</b>	string value	PEM encoded client certificate cert file
<b>default-domain-id = None</b>	string value	Optional domain ID to use with v3 and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.
<b>default-domain-name = None</b>	string value	Optional domain name to use with v3 API and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.
<b>domain-id = None</b>	string value	Domain ID to scope to
<b>domain-name = None</b>	string value	Domain name to scope to
<b>endpoint_type = public</b>	string value	Type of the nova endpoint to use. This endpoint will be looked up in the keystone catalog and should be one of public, internal or admin.
<b>insecure = False</b>	boolean value	Verify HTTPS connections.
<b>keyfile = None</b>	string value	PEM encoded client certificate key file
<b>password = None</b>	string value	User's password
<b>project-domain-id = None</b>	string value	Domain ID containing project
<b>project-domain-name = None</b>	string value	Domain name containing project
<b>project-id = None</b>	string value	Project ID to scope to
<b>project-name = None</b>	string value	Project name to scope to
<b>region_name = None</b>	string value	Name of nova region to use. Useful if keystone manages more than one region.

Configuration option = Default value	Type	Description
<b>system-scope = None</b>	string value	Scope for system operations
<b>tenant-id = None</b>	string value	Tenant ID
<b>tenant-name = None</b>	string value	Tenant Name
<b>timeout = None</b>	integer value	Timeout value for http requests
<b>trust-id = None</b>	string value	Trust ID
<b>user-domain-id = None</b>	string value	User's domain id
<b>user-domain-name = None</b>	string value	User's domain name
<b>user-id = None</b>	string value	User id
<b>username = None</b>	string value	Username

### 10.7.8. oslo\_concurrency

The following table outlines the options available under the **[oslo\_concurrency]** group in the `/etc/neutron/neutron.conf` file.

Table 10.35. oslo\_concurrency

Configuration option = Default value	Type	Description
<b>disable_process_locking = False</b>	boolean value	Enables or disables inter-process locks.
<b>lock_path = None</b>	string value	Directory to use for lock files. For security, the specified directory should only be writable by the user running the processes that need locking. Defaults to environment variable OSLO_LOCK_PATH. If external locks are used, a lock path must be set.

### 10.7.9. oslo\_messaging\_amqp

The following table outlines the options available under the **[oslo\_messaging\_amqp]** group in the `/etc/neutron/neutron.conf` file.

Table 10.36. oslo\_messaging\_amqp

Configuration option = Default value	Type	Description
<b>addressing_mode = dynamic</b>	string value	Indicates the addressing mode used by the driver. Permitted values: <i>legacy</i> - use legacy non-routable addressing <i>routable</i> - use routable addresses <i>dynamic</i> - use legacy addresses if the message bus does not support routing otherwise use routable addressing
<b>allow_insecure_clients = False</b>	boolean value	Accept clients using either SSL or plain TCP
<b>anycast_address = anycast</b>	string value	Appended to the address prefix when sending to a group of consumers. Used by the message bus to identify messages that should be delivered in a round-robin fashion across consumers.
<b>broadcast_prefix = broadcast</b>	string value	address prefix used when broadcasting to all servers
<b>connection_retry_backoff = 2</b>	integer value	Increase the <code>connection_retry_interval</code> by this many seconds after each unsuccessful failover attempt.
<b>connection_retry_interval = 1</b>	integer value	Seconds to pause before attempting to re-connect.
<b>connection_retry_interval _max = 30</b>	integer value	Maximum limit for <code>connection_retry_interval</code> + <code>connection_retry_backoff</code>
<b>container_name = None</b>	string value	Name for the AMQP container. must be globally unique. Defaults to a generated UUID
<b>default_notification_exch ange = None</b>	string value	Exchange name used in notification addresses. Exchange name resolution precedence: Target.exchange if set else default_notification_exchange if set else control_exchange if set else <i>notify</i>
<b>default_notify_timeout = 30</b>	integer value	The deadline for a sent notification message delivery. Only used when caller does not provide a timeout expiry.
<b>default_reply_retry = 0</b>	integer value	The maximum number of attempts to re-send a reply message which failed due to a recoverable error.
<b>default_reply_timeout = 30</b>	integer value	The deadline for an rpc reply message delivery.

Configuration option = Default value	Type	Description
<b>default_rpc_exchange = None</b>	string value	Exchange name used in RPC addresses. Exchange name resolution precedence: Target.exchange if set else default_rpc_exchange if set else control_exchange if set else <i>rpc</i>
<b>default_send_timeout = 30</b>	integer value	The deadline for an rpc cast or call message delivery. Only used when caller does not provide a timeout expiry.
<b>default_sender_link_timeout = 600</b>	integer value	The duration to schedule a purge of idle sender links. Detach link after expiry.
<b>group_request_prefix = unicast</b>	string value	address prefix when sending to any server in group
<b>idle_timeout = 0</b>	integer value	Timeout for inactive connections (in seconds)
<b>link_retry_delay = 10</b>	integer value	Time to pause between re-connecting an AMQP 1.0 link that failed due to a recoverable error.
<b>multicast_address = multicast</b>	string value	Appended to the address prefix when sending a fanout message. Used by the message bus to identify fanout messages.
<b>notify_address_prefix = openstack.org/om/notify</b>	string value	Address prefix for all generated Notification addresses
<b>notify_server_credit = 100</b>	integer value	Window size for incoming Notification messages
<code>`password = `</code>	string value	Password for message broker authentication
<b>pre_settled = [u'rpc-cast', u'rpc-reply']</b>	multi valued	Send messages of this type pre-settled. Pre-settled messages will not receive acknowledgement from the peer. Note well: pre-settled messages may be silently discarded if the delivery fails. Permitted values: <i>rpc-call</i> - send RPC Calls pre-settled <i>rpc-reply</i> - send RPC Replies pre-settled <i>rpc-cast</i> - Send RPC Casts pre-settled <i>notify</i> - Send Notifications pre-settled
<b>pseudo_vhost = True</b>	boolean value	Enable virtual host support for those message buses that do not natively support virtual hosting (such as <i>qpidd</i> ). When set to true the virtual host name will be added to all message bus addresses, effectively creating a private <i>subnet</i> per virtual host. Set to False if the message bus supports virtual hosting using the <i>hostname</i> field in the AMQP 1.0 Open performative as the name of the virtual host.



Configuration option = Default value	Type	Description
<b>reply_link_credit = 200</b>	integer value	Window size for incoming RPC Reply messages.
<b>rpc_address_prefix = openstack.org/om/rpc</b>	string value	Address prefix for all generated RPC addresses
<b>rpc_server_credit = 100</b>	integer value	Window size for incoming RPC Request messages
<code>`sasl_config_dir = `</code>	string value	Path to directory that contains the SASL configuration
<code>`sasl_config_name = `</code>	string value	Name of configuration file (without .conf suffix)
<code>`sasl_default_realm = `</code>	string value	SASL realm to use if no realm present in username
<code>`sasl_mechanisms = `</code>	string value	Space separated list of acceptable SASL mechanisms
<b>server_request_prefix = exclusive</b>	string value	address prefix used when sending to a specific server
<b>ssl = False</b>	boolean value	Attempt to connect via SSL. If no other ssl-related parameters are given, it will use the system's CA-bundle to verify the server's certificate.
<code>`ssl_ca_file = `</code>	string value	CA certificate PEM file used to verify the server's certificate
<code>`ssl_cert_file = `</code>	string value	Self-identifying certificate PEM file for client authentication
<code>`ssl_key_file = `</code>	string value	Private key PEM file used to sign ssl_cert_file certificate (optional)
<b>ssl_key_password = None</b>	string value	Password for decrypting ssl_key_file (if encrypted)
<b>ssl_verify_vhost = False</b>	boolean value	By default SSL checks that the name in the server's certificate matches the hostname in the transport_url. In some configurations it may be preferable to use the virtual hostname instead, for example if the server uses the Server Name Indication TLS extension (rfc6066) to provide a certificate per virtual host. Set ssl_verify_vhost to True if the server's SSL certificate uses the virtual host name instead of the DNS name.
<b>trace = False</b>	boolean value	Debug: dump AMQP frames to stdout

Configuration option = Default value	Type	Description
<b>unicast_address = unicast</b>	string value	Appended to the address prefix when sending to a particular RPC/Notification server. Used by the message bus to identify messages sent to a single destination.
<code>`username = `</code>	string value	User name for message broker authentication

### 10.7.10. oslo\_messaging\_kafka

The following table outlines the options available under the **[oslo\_messaging\_kafka]** group in the `/etc/neutron/neutron.conf` file.

Table 10.37. oslo\_messaging\_kafka

Configuration option = Default value	Type	Description
<b>conn_pool_min_size = 2</b>	integer value	The pool size limit for connections expiration policy
<b>conn_pool_ttl = 1200</b>	integer value	The time-to-live in sec of idle connections in the pool
<b>consumer_group = oslo_messaging_consumer</b>	string value	Group id for Kafka consumer. Consumers in one group will coordinate message consumption
<b>kafka_consumer_timeout = 1.0</b>	floating point value	Default timeout(s) for Kafka consumers
<b>kafka_default_host = localhost</b>	string value	Default Kafka broker Host
<b>kafka_default_port = 9092</b>	port value	Default Kafka broker Port
<b>kafka_max_fetch_bytes = 1048576</b>	integer value	Max fetch bytes of Kafka consumer
<b>pool_size = 10</b>	integer value	Pool Size for Kafka Consumers
<b>producer_batch_size = 16384</b>	integer value	Size of batch for the producer async send
<b>producer_batch_timeout = 0.0</b>	floating point value	Upper bound on the delay for KafkaProducer batching in seconds

### 10.7.11. oslo\_messaging\_notifications

The following table outlines the options available under the **[oslo\_messaging\_notifications]** group in the **/etc/neutron/neutron.conf** file.

Table 10.38. oslo\_messaging\_notifications

Configuration option = Default value	Type	Description
<b>driver = []</b>	multi valued	The Drivers(s) to handle sending notifications. Possible values are messaging, messagingv2, routing, log, test, noop
<b>retry = -1</b>	integer value	The maximum number of attempts to re-send a notification message which failed to be delivered due to a recoverable error. 0 - No retry, -1 - indefinite
<b>topics = [u'notifications']</b>	list value	AMQP topic used for OpenStack notifications.
<b>transport_url = None</b>	string value	A URL representing the messaging driver to use for notifications. If not set, we fall back to the same configuration used for RPC.

### 10.7.12. oslo\_messaging\_rabbit

The following table outlines the options available under the **[oslo\_messaging\_rabbit]** group in the **/etc/neutron/neutron.conf** file.

Table 10.39. oslo\_messaging\_rabbit

Configuration option = Default value	Type	Description
<b>amqp_auto_delete = False</b>	boolean value	Auto-delete queues in AMQP.
<b>amqp_durable_queues = False</b>	boolean value	Use durable queues in AMQP.
<b>channel_max = None</b>	integer value	Maximum number of channels to allow
<b>connection_factory = single</b>	string value	Connection factory implementation
<b>default_notification_exchange = \${control_exchange}_notification</b>	string value	Exchange name for sending notifications

Configuration option = Default value	Type	Description
<b>default_notification_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending notification, -1 means infinite retry.
<b>default_rpc_exchange = \${control_exchange}_rpc</b>	string value	Exchange name for sending RPC messages
<b>default_rpc_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending RPC message, -1 means infinite retry. If actual retry attempts in not 0 the rpc request could be processed more than one time
<b>default_serializer_type = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>fake_rabbit = False</b>	boolean value	Deprecated, use <code>rpc_backend=kombu+memory</code> or <code>rpc_backend=fake</code>
<b>frame_max = None</b>	integer value	The maximum byte size for an AMQP frame
<b>heartbeat_interval = 3</b>	integer value	How often to send heartbeats for consumer's connections
<b>heartbeat_rate = 2</b>	integer value	How often times during the <code>heartbeat_timeout_threshold</code> we check the heartbeat.
<b>heartbeat_timeout_threshold = 60</b>	integer value	Number of seconds after which the Rabbit broker is considered down if heartbeat's keep-alive fails (0 disable the heartbeat). EXPERIMENTAL
<b>host_connection_reconnect_delay = 0.25</b>	floating point value	Set delay for reconnection to some host which has connection error
<b>kombu_compression = None</b>	string value	EXPERIMENTAL: Possible values are: <code>gzip</code> , <code>bz2</code> . If not set compression will not be used. This option may not be available in future versions.
<b>kombu_failover_strategy = round-robin</b>	string value	Determines how the next RabbitMQ node is chosen in case the one we are currently connected to becomes unavailable. Takes effect only if more than one RabbitMQ node is provided in config.
<b>kombu_missing_consumer_retry_timeout = 60</b>	integer value	How long to wait a missing client before abandoning to send it its replies. This value should not be longer than <code>rpc_response_timeout</code> .

Configuration option = Default value	Type	Description
<b>kombu_reconnect_delay = 1.0</b>	floating point value	How long to wait before reconnecting in response to an AMQP consumer cancel notification.
<b>notification_listener_prefetch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to notification listener.
<b>notification_persistence = False</b>	boolean value	Persist notification messages.
<b>notification_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending notification message
<b>pool_max_overflow = 0</b>	integer value	Maximum number of connections to create above <b>pool_max_size</b> .
<b>pool_max_size = 30</b>	integer value	Maximum number of connections to keep queued.
<b>pool_recycle = 600</b>	integer value	Lifetime of a connection (since creation) in seconds or None for no recycling. Expired connections are closed on acquire.
<b>pool_stale = 60</b>	integer value	Threshold at which inactive (since release) connections are considered stale in seconds or None for no staleness. Stale connections are closed on acquire.
<b>pool_timeout = 30</b>	integer value	Default number of seconds to wait for a connections to available
<b>rabbit_ha_queues = False</b>	boolean value	Try to use HA queues in RabbitMQ (x-ha-policy: all). If you change this option, you must wipe the RabbitMQ database. In RabbitMQ 3.0, queue mirroring is no longer controlled by the x-ha-policy argument when declaring a queue. If you just want to make sure that all queues (except those with auto-generated names) are mirrored across all nodes, run: <code>"rabbitmqctl set_policy HA ^{?!amq\\.}.* {\"ha-mode\": \"all\"}"</code>
<b>rabbit_host = localhost</b>	string value	The RabbitMQ broker address where a single node is used.
<b>rabbit_hosts = [u'\$rabbit_host:\$rabbit_port']</b>	list value	RabbitMQ HA cluster host:port pairs.

Configuration option = Default value	Type	Description
<b>rabbit_interval_max = 30</b>	integer value	Maximum interval of RabbitMQ connection retries. Default is 30 seconds.
<b>rabbit_login_method = AMQPLAIN</b>	string value	The RabbitMQ login method.
<b>rabbit_max_retries = 0</b>	integer value	Maximum number of RabbitMQ connection retries. Default is 0 (infinite retry count).
<b>rabbit_password = guest</b>	string value	The RabbitMQ password.
<b>rabbit_port = 5672</b>	port value	The RabbitMQ broker port where a single node is used.
<b>rabbit_qos_prefetch_count = 0</b>	integer value	Specifies the number of messages to prefetch. Setting to zero allows unlimited messages.
<b>rabbit_retry_backoff = 2</b>	integer value	How long to backoff for between retries when connecting to RabbitMQ.
<b>rabbit_retry_interval = 1</b>	integer value	How frequently to retry connecting with RabbitMQ.
<b>rabbit_transient_queues_ttl = 1800</b>	integer value	Positive integer representing duration in seconds for queue TTL (x-expires). Queues which are unused for the duration of the TTL are automatically deleted. The parameter affects only reply and fanout queues.
<b>rabbit_userid = guest</b>	string value	The RabbitMQ userid.
<b>rabbit_virtual_host = /</b>	string value	The RabbitMQ virtual host.
<b>rpc_listener_prefetch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to rpc listener.
<b>rpc_queue_expiration = 60</b>	integer value	Time to live for rpc queues without consumers in seconds.
<b>rpc_reply_exchange = \${control_exchange}_rpc_reply</b>	string value	Exchange name for receiving RPC replies
<b>rpc_reply_listener_prefetch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to rpc reply listener.

Configuration option = Default value	Type	Description
<b>rpc_reply_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending reply. -1 means infinite retry during <code>rpc_timeout</code>
<b>rpc_reply_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending reply.
<b>rpc_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending RPC message
<b>socket_timeout = 0.25</b>	floating point value	Set socket timeout in seconds for connection's socket
<b>ssl = None</b>	boolean value	Enable SSL
<code>`ssl_ca_file = `</code>	string value	SSL certification authority file (valid only if SSL enabled).
<code>`ssl_cert_file = `</code>	string value	SSL cert file (valid only if SSL enabled).
<code>`ssl_key_file = `</code>	string value	SSL key file (valid only if SSL enabled).
<b>ssl_options = None</b>	dict value	Arguments passed to <code>ssl.wrap_socket</code>
<code>`ssl_version = `</code>	string value	SSL version to use (valid only if SSL enabled). Valid values are <code>TLSv1</code> and <code>SSLv23</code> . <code>SSLv2</code> , <code>SSLv3</code> , <code>TLSv1_1</code> , and <code>TLSv1_2</code> may be available on some distributions.
<b>tcp_user_timeout = 0.25</b>	floating point value	Set <code>TCP_USER_TIMEOUT</code> in seconds for connection's socket

### 10.7.13. oslo\_messaging\_zmq

The following table outlines the options available under the **[oslo\_messaging\_zmq]** group in the `/etc/neutron/neutron.conf` file.

Table 10.40. oslo\_messaging\_zmq

Configuration option = Default value	Type	Description
<b>rpc_ack_timeout_base = 15</b>	integer value	Number of seconds to wait for an ack from a cast/call. After each retry attempt this timeout is multiplied by some specified multiplier.

Configuration option = Default value	Type	Description
<b>rpc_ack_timeout_multiplier = 2</b>	integer value	Number to multiply base ack timeout by after each retry attempt.
<b>rpc_message_ttl = 300</b>	integer value	Expiration timeout in seconds of a sent/received message after which it is not tracked anymore by a client/server.
<b>rpc_poll_timeout = 1</b>	integer value	The default number of seconds that poll should wait. Poll raises timeout exception when timeout expired.
<b>rpc_retry_attempts = 3</b>	integer value	Default number of message sending attempts in case of any problems occurred: positive value N means at most N retries, 0 means no retries, None or -1 (or any other negative values) mean to retry forever. This option is used only if acknowledgments are enabled.
<b>rpc_thread_pool_size = 100</b>	integer value	Maximum number of (green) threads to work concurrently.
<b>rpc_use_acks = False</b>	boolean value	Wait for message acknowledgements from receivers. This mechanism works only via proxy without PUB/SUB.
<b>rpc_zmq_bind_address = *</b>	string value	ZeroMQ bind address. Should be a wildcard (*), an ethernet interface, or IP. The "host" option should point or resolve to this address.
<b>rpc_zmq_bind_port_retries = 100</b>	integer value	Number of retries to find free port number before fail with ZMQBindError.
<b>rpc_zmq_contexts = 1</b>	integer value	Number of ZeroMQ contexts, defaults to 1.
<b>rpc_zmq_host = ccosp-undercloud.usersys.redhat.com</b>	string value	Name of this node. Must be a valid hostname, FQDN, or IP address. Must match "host" option, if running Nova.
<b>rpc_zmq_ipc_dir = /var/run/openstack</b>	string value	Directory for holding IPC sockets.
<b>rpc_zmq_matchmaker = redis</b>	string value	MatchMaker driver.
<b>rpc_zmq_max_port = 65536</b>	integer value	Maximal port number for random ports range.



Configuration option = Default value	Type	Description
<b>rpc_zmq_min_port = 49153</b>	port value	Minimal port number for random ports range.
<b>rpc_zmq_serialization = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>rpc_zmq_topic_backlog = None</b>	integer value	Maximum number of ingress messages to locally buffer per topic. Default is unlimited.
<b>subscribe_on = []</b>	list value	List of publisher hosts SubConsumer can subscribe on. This option has higher priority then the default publishers list taken from the matchmaker.
<b>use_dynamic_connections = False</b>	boolean value	This option makes direct connections dynamic or static. It makes sense only with use_router_proxy=False which means to use direct connections for direct message types (ignored otherwise).
<b>use_pub_sub = False</b>	boolean value	Use PUB/SUB pattern for fanout methods. PUB/SUB always uses proxy.
<b>use_router_proxy = False</b>	boolean value	Use ROUTER remote proxy.
<b>zmq_failover_connections = 2</b>	integer value	How many additional connections to a host will be made for failover reasons. This option is actual only in dynamic connections mode.
<b>zmq_immediate = True</b>	boolean value	This option configures round-robin mode in zmq socket. True means not keeping a queue when server side disconnects. False means to keep queue and messages even if server is disconnected, when the server appears we send all accumulated messages to it.
<b>zmq_linger = -1</b>	integer value	Number of seconds to wait before all pending messages will be sent after closing a socket. The default value of -1 specifies an infinite linger period. The value of 0 specifies no linger period. Pending messages shall be discarded immediately when the socket is closed. Positive values specify an upper bound for the linger period.
<b>zmq_target_expire = 300</b>	integer value	Expiration timeout in seconds of a name service record about existing target (< 0 means no timeout).

Configuration option = Default value	Type	Description
<b>zmq_target_update = 180</b>	integer value	Update period in seconds of a name service record about existing target.
<b>zmq_tcp_keepalive = -1</b>	integer value	Enable/disable TCP keepalive (KA) mechanism. The default value of -1 (or any other negative value) means to skip any overrides and leave it to OS default; 0 and 1 (or any other positive value) mean to disable and enable the option respectively.
<b>zmq_tcp_keepalive_cnt = -1</b>	integer value	The number of retransmissions to be carried out before declaring that remote end is not available. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_idle = -1</b>	integer value	The duration between two keepalive transmissions in idle condition. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_intvl = -1</b>	integer value	The duration between two successive keepalive retransmissions, if acknowledgement to the previous keepalive transmission is not received. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.

#### 10.7.14. oslo\_middlewares

The following table outlines the options available under the **[oslo\_middlewares]** group in the **/etc/neutron/neutron.conf** file.

Table 10.41. oslo\_middlewares

Configuration option = Default value	Type	Description
<b>enable_proxy_headers_parsing = False</b>	boolean value	Whether the application is behind a proxy or not. This determines if the middleware should parse the headers or not.

#### 10.7.15. oslo\_policy

The following table outlines the options available under the **[oslo\_policy]** group in the `/etc/neutron/neutron.conf` file.

Table 10.42. oslo\_policy

Configuration option = Default value	Type	Description
<b>enforce_scope = False</b>	boolean value	This option controls whether or not to enforce scope when evaluating policies. If <b>True</b> , the scope of the token used in the request is compared to the <b>scope_types</b> of the policy being enforced. If the scopes do not match, an <b>InvalidScope</b> exception will be raised. If <b>False</b> , a message will be logged informing operators that policies are being invoked with mismatching scope.
<b>policy_default_rule = default</b>	string value	Default rule. Enforced when a requested rule is not found.
<b>policy_dirs = [u'policy.d']</b>	multi valued	Directories where policy configuration files are stored. They can be relative to any directory in the search path defined by the <code>config_dir</code> option, or absolute paths. The file defined by <code>policy_file</code> must exist for these directories to be searched. Missing or empty directories are ignored.
<b>policy_file = policy.json</b>	string value	The file that defines policies.
<b>remote_content_type = application/x-www-form-urlencoded</b>	string value	Content Type to send and receive data for REST based policy check
<b>remote_ssl_ca_cert_file = None</b>	string value	Absolute path to ca cert file for REST based policy check
<b>remote_ssl_client_cert_file = None</b>	string value	Absolute path to client cert for REST based policy check
<b>remote_ssl_client_key_file = None</b>	string value	Absolute path client key file REST based policy check
<b>remote_ssl_verify_server_cert = False</b>	boolean value	server identity verification for REST based policy check

### 10.7.16. quotas

The following table outlines the options available under the **[quotas]** group in the `/etc/neutron/neutron.conf` file.

Table 10.43. quotas

Configuration option = Default value	Type	Description
<b>default_quota = -1</b>	integer value	Default number of resource allowed per tenant. A negative value means unlimited.
<b>quota_driver = neutron.db.quota.driver.DbQuotaDriver</b>	string value	Default driver to use for quota checks.
<b>quota_floatingip = 50</b>	integer value	Number of floating IPs allowed per tenant. A negative value means unlimited.
<b>quota_network = 100</b>	integer value	Number of networks allowed per tenant. A negative value means unlimited.
<b>quota_port = 500</b>	integer value	Number of ports allowed per tenant. A negative value means unlimited.
<b>quota_router = 10</b>	integer value	Number of routers allowed per tenant. A negative value means unlimited.
<b>quota_security_group = 10</b>	integer value	Number of security groups allowed per tenant. A negative value means unlimited.
<b>quota_security_group_rule = 100</b>	integer value	Number of security rules allowed per tenant. A negative value means unlimited.
<b>quota_subnet = 100</b>	integer value	Number of subnets allowed per tenant, A negative value means unlimited.
<b>track_quota_usage = True</b>	boolean value	Keep in track in the database of current resource quota usage. Plugins which do not leverage the neutron database should set this flag to False.

### 10.7.17. ssl

The following table outlines the options available under the **[ssl]** group in the `/etc/neutron/neutron.conf` file.

Table 10.44. ssl

Configuration option = Default value	Type	Description
<b>ca_file = None</b>	string value	CA certificate file to use to verify connecting clients.
<b>cert_file = None</b>	string value	Certificate file to use when starting the server securely.

Configuration option = Default value	Type	Description
<b>ciphers = None</b>	string value	Sets the list of available ciphers. value should be a string in the OpenSSL cipher list format.
<b>key_file = None</b>	string value	Private key file to use when starting the server securely.
<b>version = None</b>	string value	SSL version to use (valid only if SSL enabled). Valid values are TLSv1 and SSLv23. SSLv2, SSLv3, TLSv1_1, and TLSv1_2 may be available on some distributions.

## 10.8. OPENVSWITCH\_AGENT.INI

This section contains options for the `/etc/neutron/plugins/ml2/openvswitch_agent.ini` file.

### 10.8.1. DEFAULT

The following table outlines the options available under the **[DEFAULT]** group in the `/etc/neutron/plugins/ml2/openvswitch_agent.ini` file.

Table 10.45. DEFAULT

Configuration option = Default value	Type	Description
<b>debug = False</b>	boolean value	If set to true, the logging level will be set to DEBUG instead of the default INFO level.

Configuration option = Default value	Type	Description
<b>default_log_levels =</b> [u'amqp=WARN', u'amqplib=WARN', u'boto=WARN', u'qpid=WARN', u'sqlalchemy=WARN', u'suds=INFO', u'oslo.messaging=INFO', u'oslo_messaging=INFO', u'iso8601=WARN', u'requests.packages.urlli b3.connectionpool=WAR N', u'urllib3.connectionpool= WARN', u'websocket=WARN', u'requests.packages.urlli b3.util.retry=WARN', u'urllib3.util.retry=WARN', u'keystonemiddleware=W ARN', u'routes.middleware=WA RN', u'stevedore=WARN', u'taskflow=WARN', u'keystoneauth=WARN', u'oslo.cache=INFO', u'dogpile.core.dogpile=IN FO']	list value	List of package logging levels in logger=LEVEL pairs. This option is ignored if log_config_append is set.
<b>fatal_deprecations =</b> <b>False</b>	boolean value	Enables or disables fatal status of deprecations.
<code>`instance_format = [instance: %(uuid)s] `</code>	string value	The format for an instance that is passed with the log message.
<code>`instance_uuid_format = [instance:%(uuid)s] `</code>	string value	The format for an instance UUID that is passed with the log message.
<b>log-config-append = None</b>	string value	The name of a logging configuration file. This file is appended to any existing logging configuration files. For details about logging configuration files, see the Python logging module documentation. Note that when logging configuration files are used then all logging configuration is set in the configuration file and other logging configuration options are ignored (for example, logging_context_format_string).

Configuration option = Default value	Type	Description
<b>log-date-format = %Y-%m-%d %H:%M:%S</b>	string value	Defines the format string for <code>%(asctime)s</code> in log records. Default: <code>%(default)s</code> . This option is ignored if <code>log_config_append</code> is set.
<b>log-dir = None</b>	string value	(Optional) The base directory used for relative <code>log_file</code> paths. This option is ignored if <code>log_config_append</code> is set.
<b>log-file = None</b>	string value	(Optional) Name of log file to send logging output to. If no default is set, logging will go to <code>stderr</code> as defined by <code>use_stderr</code> . This option is ignored if <code>log_config_append</code> is set.
<b>logging_context_format_string = %(asctime)s.% (msecs)03d %(process)d %(levelname)s %(name)s [% (request_id)s % (user_identity)s] % (instance)s%(message)s</b>	string value	Format string to use for log messages with context.
<b>logging_debug_format_s uffix = %(funcName)s % (pathname)s:%(lineno)d</b>	string value	Additional data to append to log message when logging level for the message is <code>DEBUG</code> .
<b>logging_default_format_s tring = %(asctime)s.% (msecs)03d %(process)d %(levelname)s %(name)s [-] %(instance)s% (message)s</b>	string value	Format string to use for log messages when context is undefined.
<b>logging_exception_prefix = %(asctime)s.% (msecs)03d %(process)d ERROR %(name)s % (instance)s</b>	string value	Prefix each line of exception output with this format.
<b>logging_user_identity_for mat = %(user)s % (tenant)s %(domain)s % (user_domain)s % (project_domain)s</b>	string value	Defines the format string for <code>%(user_identity)s</code> that is used in <code>logging_context_format_string</code> .
<b>publish_errors = False</b>	boolean value	Enables or disables publication of error events.

Configuration option = Default value	Type	Description
<b>rate_limit_burst = 0</b>	integer value	Maximum number of logged messages per <code>rate_limit_interval</code> .
<b>rate_limit_except_level = CRITICAL</b>	string value	Log level name used by rate limiting: CRITICAL, ERROR, INFO, WARNING, DEBUG or empty string. Logs with level greater or equal to <code>rate_limit_except_level</code> are not filtered. An empty string means that all levels are filtered.
<b>rate_limit_interval = 0</b>	integer value	Interval, number of seconds, of log rate limiting.
<b>syslog-log-facility = LOG_USER</b>	string value	Syslog facility to receive log lines. This option is ignored if <code>log_config_append</code> is set.
<b>use-journal = False</b>	boolean value	Enable journald for logging. If running in a systemd environment you may wish to enable journal support. Doing so will use the journal native protocol which includes structured metadata in addition to log messages. This option is ignored if <code>log_config_append</code> is set.
<b>use-json = False</b>	boolean value	Use JSON formatting for logging. This option is ignored if <code>log_config_append</code> is set.
<b>use-syslog = False</b>	boolean value	Use syslog for logging. Existing syslog format is DEPRECATED and will be changed later to honor RFC5424. This option is ignored if <code>log_config_append</code> is set.
<b>use_stderr = False</b>	boolean value	Log output to standard error. This option is ignored if <code>log_config_append</code> is set.
<b>watch-log-file = False</b>	boolean value	Uses logging handler designed to watch file system. When log file is moved or removed this handler will open a new log file with specified path instantaneously. It makes sense only if <code>log_file</code> option is specified and Linux platform is used. This option is ignored if <code>log_config_append</code> is set.

### 10.8.2. agent

The following table outlines the options available under the **[agent]** group in the `/etc/neutron/plugins/ml2/openvswitch_agent.ini` file.

Table 10.46. agent



Configuration option = Default value	Type	Description
<b>agent_type = Open vSwitch agent</b>	string value	Selects the Agent Type reported
<b>arp_responder = False</b>	boolean value	Enable local ARP responder if it is supported. Requires OVS 2.1 and ML2 I2population driver. Allows the switch (when supporting an overlay) to respond to an ARP request locally without performing a costly ARP broadcast into the overlay.
<b>dont_fragment = True</b>	boolean value	Set or un-set the don't fragment (DF) bit on outgoing IP packet carrying GRE/VXLAN tunnel.
<b>drop_flows_on_start = False</b>	boolean value	Reset flow table on start. Setting this to True will cause brief traffic interruption.
<b>enable_distributed_routing = False</b>	boolean value	Make the I2 agent run in DVR mode.
<b>extensions = []</b>	list value	Extensions list to use
<b>I2_population = False</b>	boolean value	Use ML2 I2population mechanism driver to learn remote MAC and IPs and improve tunnel scalability.
<b>minimize_polling = True</b>	boolean value	Minimize polling by monitoring ovssdb for interface changes.
<b>ovsdb_monitor_respawn_interval = 30</b>	integer value	The number of seconds to wait before respawning the ovsdb monitor after losing communication with it.
<b>tunnel_csum = False</b>	boolean value	Set or un-set the tunnel header checksum on outgoing IP packet carrying GRE/VXLAN tunnel.
<b>tunnel_types = []</b>	list value	Network types supported by the agent (gre and/or vxlan).
<b>veth_mtu = 9000</b>	integer value	MTU size of veth interfaces
<b>vxlan_udp_port = 4789</b>	port value	The UDP port to use for VXLAN tunnels.

### 10.8.3. network\_log

The following table outlines the options available under the **[network\_log]** group in the **/etc/neutron/plugins/ml2/openvswitch\_agent.ini** file.

Table 10.47. network\_log

Configuration option = Default value	Type	Description
<b>burst_limit = 25</b>	integer value	Maximum number of packets per rate_limit.
<b>local_output_log_base = None</b>	string value	Output logfile path on agent side, default syslog file.
<b>rate_limit = 100</b>	integer value	Maximum packets logging per second.

#### 10.8.4. ovs

The following table outlines the options available under the **[ovs]** group in the `/etc/neutron/plugins/ml2/openvswitch_agent.ini` file.

Table 10.48. ovs

Configuration option = Default value	Type	Description
<b>bridge_mappings = []</b>	list value	Comma-separated list of <physical_network>: <bridge> tuples mapping physical network names to the agent's node-specific Open vSwitch bridge names to be used for flat and VLAN networks. The length of bridge names should be no more than 11. Each bridge must exist, and should have a physical network interface configured as a port. All physical networks configured on the server should have mappings to appropriate bridges on each agent. Note: If you remove a bridge from this mapping, make sure to disconnect it from the integration bridge as it won't be managed by the agent anymore.
<b>datapath_type = system</b>	string value	OVS datapath to use. <i>system</i> is the default value and corresponds to the kernel datapath. To enable the userspace datapath set this value to <i>netdev</i> .
<b>int_peer_patch_port = patch-tun</b>	string value	Peer patch port in integration bridge for tunnel bridge.
<b>integration_bridge = br-int</b>	string value	Integration bridge to use. Do not change this parameter unless you have a good reason to. This is the name of the OVS integration bridge. There is one per hypervisor. The integration bridge acts as a virtual <i>patch bay</i> . All VM VIFs are attached to this bridge and then <i>patched</i> according to their network connectivity.

Configuration option = Default value	Type	Description
<b>local_ip = None</b>	IP address value	IP address of local overlay (tunnel) network endpoint. Use either an IPv4 or IPv6 address that resides on one of the host network interfaces. The IP version of this value must match the value of the <i>overlay_ip_version</i> option in the ML2 plug-in configuration file on the neutron server node(s).
<b>of_connect_timeout = 300</b>	integer value	Timeout in seconds to wait for the local switch connecting the controller. Used only for <i>native</i> driver.
<b>of_interface = native</b>	string value	OpenFlow interface to use.
<b>of_listen_address = 127.0.0.1</b>	IP address value	Address to listen on for OpenFlow connections. Used only for <i>native</i> driver.
<b>of_listen_port = 6633</b>	port value	Port to listen on for OpenFlow connections. Used only for <i>native</i> driver.
<b>of_request_timeout = 300</b>	integer value	Timeout in seconds to wait for a single OpenFlow request. Used only for <i>native</i> driver.
<b>ovsdb_connection = tcp:127.0.0.1:6640</b>	string value	The connection string for the OVSDB backend. Will be used by ovsdb-client when monitoring and used for the all ovsdb commands when native ovsdb_interface is enabled
<b>ovsdb_interface = native</b>	string value	The interface for interacting with the OVSDB
<b>ssl_ca_cert_file = None</b>	string value	The Certificate Authority (CA) certificate to use when interacting with OVSDB. Required when using an "ssl:" prefixed ovsdb_connection
<b>ssl_cert_file = None</b>	string value	The SSL certificate file to use when interacting with OVSDB. Required when using an "ssl:" prefixed ovsdb_connection
<b>ssl_key_file = None</b>	string value	The SSL private key file to use when interacting with OVSDB. Required when using an "ssl:" prefixed ovsdb_connection
<b>tun_peer_patch_port = patch-int</b>	string value	Peer patch port in tunnel bridge for integration bridge.
<b>tunnel_bridge = br-tun</b>	string value	Tunnel bridge to use.

Configuration option = Default value	Type	Description
<b>use_veth_interconnection = False</b>	boolean value	Use veths instead of patch ports to interconnect the integration bridge to physical networks. Support kernel without Open vSwitch patch port support so long as it is set to True.
<b>vhostuser_socket_dir = /var/run/openvswitch</b>	string value	OVS vhost-user socket directory.

### 10.8.5. securitygroup

The following table outlines the options available under the **[securitygroup]** group in the **/etc/neutron/plugins/ml2/openvswitch\_agent.ini** file.

Table 10.49. securitygroup

Configuration option = Default value	Type	Description
<b>enable_ipset = True</b>	boolean value	Use ipset to speed-up the iptables based security groups. Enabling ipset support requires that ipset is installed on L2 agent node.
<b>enable_security_group = True</b>	boolean value	Controls whether the neutron security group API is enabled in the server. It should be false when using no security groups or using the nova security group API.
<b>firewall_driver = None</b>	string value	Driver for security groups firewall in the L2 agent

### 10.8.6. xenapi

The following table outlines the options available under the **[xenapi]** group in the **/etc/neutron/plugins/ml2/openvswitch\_agent.ini** file.

Table 10.50. xenapi

Configuration option = Default value	Type	Description
<b>connection_password = None</b>	string value	Password for connection to XenServer/Xen Cloud Platform.
<b>connection_url = None</b>	string value	URL for connection to XenServer/Xen Cloud Platform.
<b>connection_username = None</b>	string value	Username for connection to XenServer/Xen Cloud Platform.

## 10.9. SRIOV\_AGENT.INI

This section contains options for the `/etc/neutron/plugins/ml2/sriov_agent.ini` file.

### 10.9.1. DEFAULT

The following table outlines the options available under the **[DEFAULT]** group in the `/etc/neutron/plugins/ml2/sriov_agent.ini` file.

Table 10.51. DEFAULT

Configuration option = Default value	Type	Description
<b>debug = False</b>	boolean value	If set to true, the logging level will be set to DEBUG instead of the default INFO level.
<b>default_log_levels =</b> [u'amqp=WARN', u'amqplib=WARN', u'boto=WARN', u'qpid=WARN', u'sqlalchemy=WARN', u'suds=INFO', u'oslo.messaging=INFO', u'oslo_messaging=INFO', u'iso8601=WARN', u'requests.packages.urlli b3.connectionpool=WAR N', u'urllib3.connectionpool= WARN', u'websocket=WARN', u'requests.packages.urlli b3.util.retry=WARN', u'urllib3.util.retry=WARN', u'keystonemiddleware=W ARN', u'routes.middleware=WA RN', u'stevedore=WARN', u'taskflow=WARN', u'keystoneauth=WARN', u'oslo.cache=INFO', u'dogpile.core.dogpile=IN FO']	list value	List of package logging levels in logger=LEVEL pairs. This option is ignored if log_config_append is set.
<b>fatal_deprecations =</b> <b>False</b>	boolean value	Enables or disables fatal status of deprecations.
<code>`instance_format = [instance: %(uuid)s] `</code>	string value	The format for an instance that is passed with the log message.

Configuration option = Default value	Type	Description
<code>`instance_uuid_format = [instance: %(uuid)s]`</code>	string value	The format for an instance UUID that is passed with the log message.
<b>log-config-append = None</b>	string value	The name of a logging configuration file. This file is appended to any existing logging configuration files. For details about logging configuration files, see the Python logging module documentation. Note that when logging configuration files are used then all logging configuration is set in the configuration file and other logging configuration options are ignored (for example, <code>logging_context_format_string</code> ).
<b>log-date-format = %Y-%m-%d %H:%M:%S</b>	string value	Defines the format string for <code>%(asctime)s</code> in log records. Default: <code>%(default)s</code> . This option is ignored if <code>log_config_append</code> is set.
<b>log-dir = None</b>	string value	(Optional) The base directory used for relative <code>log_file</code> paths. This option is ignored if <code>log_config_append</code> is set.
<b>log-file = None</b>	string value	(Optional) Name of log file to send logging output to. If no default is set, logging will go to <code>stderr</code> as defined by <code>use_stderr</code> . This option is ignored if <code>log_config_append</code> is set.
<b>logging_context_format_string = %(asctime)s.% (msecs)03d %(process)d %(levelname)s %(name)s [% (request_id)s % (user_identity)s] % (instance)s%(message)s</b>	string value	Format string to use for log messages with context.
<b>logging_debug_format_s uffix = %(funcName)s % (pathname)s:%(lineno)d</b>	string value	Additional data to append to log message when logging level for the message is <code>DEBUG</code> .
<b>logging_default_format_s tring = %(asctime)s.% (msecs)03d %(process)d %(levelname)s %(name)s [-] %(instance)s% (message)s</b>	string value	Format string to use for log messages when context is undefined.
<b>logging_exception_prefix = %(asctime)s.% (msecs)03d %(process)d ERROR %(name)s % (instance)s</b>	string value	Prefix each line of exception output with this format.

Configuration option = Default value	Type	Description
<b>logging_user_identity_format =</b> <b>%(user)s %(tenant)s %(domain)s</b> <b>%(user_domain)s</b> <b>%(project_domain)s</b>	string value	Defines the format string for <code>%(user_identity)s</code> that is used in <code>logging_context_format_string</code> .
<b>publish_errors = False</b>	boolean value	Enables or disables publication of error events.
<b>rate_limit_burst = 0</b>	integer value	Maximum number of logged messages per <code>rate_limit_interval</code> .
<b>rate_limit_except_level =</b> <b>CRITICAL</b>	string value	Log level name used by rate limiting: CRITICAL, ERROR, INFO, WARNING, DEBUG or empty string. Logs with level greater or equal to <code>rate_limit_except_level</code> are not filtered. An empty string means that all levels are filtered.
<b>rate_limit_interval = 0</b>	integer value	Interval, number of seconds, of log rate limiting.
<b>syslog-log-facility =</b> <b>LOG_USER</b>	string value	Syslog facility to receive log lines. This option is ignored if <code>log_config_append</code> is set.
<b>use-journal = False</b>	boolean value	Enable journald for logging. If running in a systemd environment you may wish to enable journal support. Doing so will use the journal native protocol which includes structured metadata in addition to log messages. This option is ignored if <code>log_config_append</code> is set.
<b>use-json = False</b>	boolean value	Use JSON formatting for logging. This option is ignored if <code>log_config_append</code> is set.
<b>use-syslog = False</b>	boolean value	Use syslog for logging. Existing syslog format is DEPRECATED and will be changed later to honor RFC5424. This option is ignored if <code>log_config_append</code> is set.
<b>use_stderr = False</b>	boolean value	Log output to standard error. This option is ignored if <code>log_config_append</code> is set.
<b>watch-log-file = False</b>	boolean value	Uses logging handler designed to watch file system. When log file is moved or removed this handler will open a new log file with specified path instantaneously. It makes sense only if <code>log_file</code> option is specified and Linux platform is used. This option is ignored if <code>log_config_append</code> is set.

## 10.9.2. agent

The following table outlines the options available under the **[agent]** group in the **/etc/neutron/plugins/ml2/sriov\_agent.ini** file.

Table 10.52. agent

Configuration option = Default value	Type	Description
<b>extensions = []</b>	list value	Extensions list to use

## 10.9.3. sriov\_nic

The following table outlines the options available under the **[sriov\_nic]** group in the **/etc/neutron/plugins/ml2/sriov\_agent.ini** file.

Table 10.53. sriov\_nic

Configuration option = Default value	Type	Description
<b>exclude_devices = []</b>	list value	Comma-separated list of <network_device>: <vfs_to_exclude> tuples, mapping network_device to the agent's node-specific list of virtual functions that should not be used for virtual networking. vfs_to_exclude is a semicolon-separated list of virtual functions to exclude from network_device. The network_device in the mapping should appear in the physical_device_mappings list.
<b>physical_device_mappings = []</b>	list value	Comma-separated list of <physical_network>: <network_device> tuples mapping physical network names to the agent's node-specific physical network device interfaces of SR-IOV physical function to be used for VLAN networks. All physical networks listed in network_vlan_ranges on the server should have mappings to appropriate interfaces on each agent.



## CHAPTER 11. NOVA

The following chapter contains information about the configuration options in the **nova** service.

### 11.1. NOVA.CONF

This section contains options for the `/etc/nova/nova.conf` file.

#### 11.1.1. DEFAULT

The following table outlines the options available under the **[DEFAULT]** group in the `/etc/nova/nova.conf` file.

Table 11.1. DEFAULT


Configuration option = Default value	Type	Description
<b>allow_resize_to_same_host = False</b>	boolean value	Allow destination machine to match source for resize. Useful when testing in single-host environments. By default it is not allowed to resize to the same host. Setting this option to true will add the same host to the destination options. Also set to true if you allow the ServerGroupAffinityFilter and need to resize.

Configuration option = Default value	Type	Description
<b>allow_same_net_traffic = True</b>	boolean value	<p>Determine whether to allow network traffic from same network.</p> <p>When set to true, hosts on the same subnet are not filtered and are allowed to pass all types of traffic between them. On a flat network, this allows all instances from all projects unfiltered communication. With VLAN networking, this allows access between instances within the same project.</p> <p>This option only applies when using the <b>nova-network</b> service. When using another networking services, such as Neutron, security groups or other approaches should be used.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True: Network traffic should be allowed pass between all instances on the same network, regardless of their tenant and security policies</li> <li>● False: Network traffic should not be allowed pass between instances unless it is unblocked in a security group</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>use_neutron</b>: This must be set to <b>False</b> to enable <b>nova-network</b> networking</li> <li>● <b>firewall_driver</b>: This must be set to <b>nova.virt.libvirt.firewall.IptablesFirewallDriver</b> to ensure the libvirt firewall driver is enabled.</li> </ul>
<b>auto_assign_floating_ip = False</b>	boolean value	<p>Autoassigning floating IP to VM</p> <p>When set to True, floating IP is auto allocated and associated to the VM upon creation.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>use_neutron</b>: this options only works with nova-network.</li> </ul>



Configuration option = Default value	Type	Description
<b>backdoor_port = None</b>	string value	Enable eventlet backdoor. Acceptable values are 0, <port>, and <start>:<end>, where 0 results in listening on a random tcp port number; <port> results in listening on the specified port number (and not enabling backdoor if that port is in use); and <start>:<end> results in listening on the smallest unused port number within the specified range of port numbers. The chosen port is displayed in the service's log file.
<b>backdoor_socket = None</b>	string value	Enable eventlet backdoor, using the provided path as a unix socket that can receive connections. This option is mutually exclusive with <i>backdoor_port</i> in that only one should be provided. If both are provided then the existence of this option overrides the usage of that option.
<b>bandwidth_poll_interval = 600</b>	integer value	Interval to pull network bandwidth usage info.  Not supported on all hypervisors. If a hypervisor doesn't support bandwidth usage, it will not get the info in the usage events.  Possible values: <ul style="list-style-type: none"> <li>● 0: Will run at the default periodic interval.</li> <li>● Any value &lt; 0: Disables the option.</li> <li>● Any positive integer in seconds.</li> </ul>
<b>bindir = /usr/local/bin</b>	string value	The directory where the Nova binaries are installed.  This option is only relevant if the networking capabilities from Nova are used (see services below). Nova's networking capabilities are targeted to be fully replaced by Neutron in the future. It is very unlikely that you need to change this option from its default value.  Possible values: <ul style="list-style-type: none"> <li>● The full path to a directory.</li> </ul>

Configuration option = Default value	Type	Description
<b>block_device_allocate_retries = 60</b>	integer value	<p>Number of times to retry block device allocation on failures. Starting with Liberty, Cinder can use image volume cache. This may help with block device allocation performance. Look at the <code>cinder image_volume_cache_enabled</code> configuration option.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● 60 (default)</li> <li>● If value is 0, then one attempt is made.</li> <li>● Any negative value is treated as 0.</li> <li>● For any value &gt; 0, total attempts are (value + 1)</li> </ul>
<b>block_device_allocate_retries_interval = 3</b>	integer value	<p>Interval (in seconds) between block device allocation retries on failures.</p> <p>This option allows the user to specify the time interval between consecutive retries.</p> <p><i>block_device_allocate_retries</i> option specifies the maximum number of retries.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● 0: Disables the option.</li> <li>● Any positive integer in seconds enables the option.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>block_device_allocate_retries</b> in <code>compute_manager_opts</code> group.</li> </ul>
<b>cert = self.pem</b>	string value	Path to SSL certificate file.

Configuration option = Default value	Type	Description
<b>cnt_vpn_clients = 0</b>	integer value	<p>This option represents the number of IP addresses to reserve at the top of the address range for VPN clients. It also will be ignored if the configuration option for <b>network_manager</b> is not set to the default of <i>nova.network.manager.VlanManager</i>.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any integer, 0 or greater.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>use_neutron</b></li> <li>● <b>network_manager</b></li> </ul>
<b>compute_driver = None</b>	string value	<p>Defines which driver to use for controlling virtualization.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● <b>libvirt.LibvirtDriver</b></li> <li>● <b>xenapi.XenAPIDriver</b></li> <li>● <b>fake.FakeDriver</b></li> <li>● <b>ironic.IronicDriver</b></li> <li>● <b>vmwareapi.VMwareVCDriver</b></li> <li>● <b>hyperv.HyperVDriver</b></li> <li>● <b>powervm.PowerVMDriver</b></li> </ul>

Configuration option = Default value	Type	Description
<b>compute_monitors = []</b>	list value	<p>A comma-separated list of monitors that can be used for getting compute metrics. You can use the alias/name from the setuptools entry points for nova.compute.monitors.* namespaces. If no namespace is supplied, the "cpu." namespace is assumed for backwards-compatibility.</p> <div data-bbox="815 510 922 678" style="display: inline-block; vertical-align: top;">  </div> <p><b>NOTE</b></p> <p>Only one monitor per namespace (For example: cpu) can be loaded at a time.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● An empty list will disable the feature (Default).</li> <li>● An example value that would enable both the CPU and NUMA memory bandwidth monitors that use the virt driver variant:</li> </ul> <pre style="margin-left: 20px;">compute_monitors = cpu.virt_driver, numa_mem_bw.virt_driver</pre>

Configuration option = Default value	Type	Description
<b>config_drive_format = iso9660</b>	string value	<p>Configuration drive format</p> <p>Configuration drive format that will contain metadata attached to the instance when it boots.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● iso9660: A file system image standard that is widely supported across operating systems. NOTE: Mind the libvirt bug (<a href="https://bugs.launchpad.net/nova/+bug/1246201">https://bugs.launchpad.net/nova/+bug/1246201</a>) - If your hypervisor driver is libvirt, and you want live migrate to work without shared storage, then use VFAT.</li> <li>● vfat: For legacy reasons, you can configure the configuration drive to use VFAT format instead of ISO 9660.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● This option is meaningful when one of the following alternatives occur: <ol style="list-style-type: none"> <li>1. force_config_drive option set to <i>true</i></li> <li>2. the REST API call to create the instance contains an enable flag for config drive option</li> <li>3. the image used to create the instance requires a config drive, this is defined by img_config_drive property for that image.</li> </ol> </li> <li>● A compute node running Hyper-V hypervisor can be configured to attach configuration drive as a CD drive. To attach the configuration drive as a CD drive, set config_drive_cdrom option at hyperv section, to true.</li> </ul>
<b>conn_pool_min_size = 2</b>	integer value	The pool size limit for connections expiration policy
<b>conn_pool_ttl = 1200</b>	integer value	The time-to-live in sec of idle connections in the pool

Configuration option = Default value	Type	Description
<b>console_host = ccsosp-undercloud.usersys.redhat.com</b>	string value	<p>Console proxy host to be used to connect to instances on this host. It is the publicly visible name for the console host.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Current hostname (default) or any string representing hostname.</li> </ul>
<b>control_exchange = openstack</b>	string value	The default exchange under which topics are scoped. May be overridden by an exchange name specified in the transport_url option.
<b>cpu_allocation_ratio = 0.0</b>	floating point value	<p>This option helps you specify virtual CPU to physical CPU allocation ratio.</p> <p>From Ocata (15.0.0) this is used to influence the hosts selected by the Placement API. Note that when Placement is used, the CoreFilter is redundant, because the Placement API will have already filtered out hosts that would have failed the CoreFilter.</p> <p>This configuration specifies ratio for CoreFilter which can be set per compute node. For AggregateCoreFilter, it will fall back to this configuration value if no per-aggregate setting is found.</p> <p> <b>NOTE</b></p> <p>This can be set per-compute, or if set to 0.0, the value set on the scheduler node(s) or compute node(s) will be used and defaulted to 16.0. Once set to a non-default value, it is not possible to "unset" the config to get back to the default behavior. If you want to reset back to the default, explicitly specify 16.0.</p> <p> <b>NOTE</b></p> <p>As of the 16.0.0 Pike release, this configuration option is ignored for the ironic.IronicDriver compute driver and is hardcoded to 1.0.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Any valid positive integer or float value</li> </ul>





Configuration option = Default value	Type	Description
<b>create_unique_mac_address_attempts = 5</b>	integer value	<p>This option determines how many times nova-network will attempt to create a unique MAC address before giving up and raising a <b>VirtualInterfaceMacAddressException</b> error.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any positive integer. The default is 5.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li><b>use_neutron</b></li> </ul>
<b>daemon = False</b>	boolean value	Run as a background process.
<b>db_driver = nova.db</b>	string value	The driver to use for database access
<b>debug = False</b>	boolean value	If set to true, the logging level will be set to DEBUG instead of the default INFO level.
<b>default_access_ip_network_name = None</b>	string value	<p>Name of the network to be used to set access IPs for instances. If there are multiple IPs to choose from, an arbitrary one will be chosen.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>None (default)</li> <li>Any string representing network name.</li> </ul>
<b>default_availability_zone = nova</b>	string value	<p>Default availability zone for compute services.</p> <p>This option determines the default availability zone for <i>nova-compute</i> services, which will be used if the service(s) do not belong to aggregates with availability zone metadata.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any string representing an existing availability zone name.</li> </ul>

Configuration option = Default value	Type	Description
<b>default_ephemeral_format = None</b>	string value	<p>The default format an ephemeral_volume will be formatted with on creation.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● <b>ext2</b></li> <li>● <b>ext3</b></li> <li>● <b>ext4</b></li> <li>● <b>xfs</b></li> <li>● <b>ntfs</b> (only for Windows guests)</li> </ul>
<b>default_flavor = m1.small</b>	string value	<p>Default flavor to use for the EC2 API only. The Nova API does not support a default flavor.</p>
<b>default_floating_pool = nova</b>	string value	<p>Default pool for floating IPs.</p> <p>This option specifies the default floating IP pool for allocating floating IPs.</p> <p>While allocating a floating ip, users can optionally pass in the name of the pool they want to allocate from, otherwise it will be pulled from the default pool.</p> <p>If this option is not set, then <i>nova</i> is used as default floating pool.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any string representing a floating IP pool name</li> </ul>

Configuration option = Default value	Type	Description
<b>default_log_levels =</b> <b>[u'amqp=WARN',</b> <b>u'amqplib=WARN',</b> <b>u'boto=WARN',</b> <b>u'qpid=WARN',</b> <b>u'sqlalchemy=WARN',</b> <b>u'suds=INFO',</b> <b>u'oslo.messaging=INFO',</b> <b>u'oslo_messaging=INFO',</b> <b>u'iso8601=WARN',</b> <b>u'requests.packages.urlli</b> <b>b3.connectionpool=WAR</b> <b>N',</b> <b>u'urllib3.connectionpool=</b> <b>WARN',</b> <b>u'websocket=WARN',</b> <b>u'requests.packages.urlli</b> <b>b3.util.retry=WARN',</b> <b>u'urllib3.util.retry=WARN',</b> <b>u'keystonemiddleware=W</b> <b>ARN',</b> <b>u'routes.middleware=WA</b> <b>RN', u'stevedore=WARN',</b> <b>u'taskflow=WARN',</b> <b>u'keystoneauth=WARN',</b> <b>u'oslo.cache=INFO',</b> <b>u'dogpile.core.dogpile=IN</b> <b>FO']</b>	list value	List of package logging levels in logger=LEVEL pairs. This option is ignored if log_config_append is set.
<b>default_schedule_zone =</b> <b>None</b>	string value	<p>Default availability zone for instances.</p> <p>This option determines the default availability zone for instances, which will be used when a user does not specify one when creating an instance. The instance(s) will be bound to this availability zone for their lifetime.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any string representing an existing availability zone name.</li> <li>None, which means that the instance can move from one availability zone to another during its lifetime if it is moved from one compute node to another.</li> </ul>
<b>defer iptables_apply =</b> <b>False</b>	boolean value	Whether to batch up the application of IPTables rules during a host restart and apply all at the end of the init phase.

Configuration option = Default value	Type	Description
<b>dhcp_domain = novalocal</b>	string value	<p>This option allows you to specify the domain for the DHCP server.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any string that is a valid domain name.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li><b>use_neutron</b></li> </ul>
<b>dhcp_lease_time = 86400</b>	integer value	<p>The lifetime of a DHCP lease, in seconds. The default is 86400 (one day).</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any positive integer value.</li> </ul>
<b>dhcpbridge = \$bindir/nova-dhcpbridge</b>	string value	<p>The location of the binary nova-dhcpbridge. By default it is the binary named <i>nova-dhcpbridge</i> that is installed with all the other nova binaries.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any string representing the full path to the binary for dhcpbridge</li> </ul>
<b>dhcpbridge_flagfile = [u'/etc/nova/nova-dhcpbridge.conf']</b>	multi valued	<p>This option is a list of full paths to one or more configuration files for dhcpbridge. In most cases the default path of <i>/etc/nova/nova-dhcpbridge.conf</i> should be sufficient, but if you have special needs for configuring dhcpbridge, you can change or add to this list.</p> <p>Possible values</p> <ul style="list-style-type: none"> <li>A list of strings, where each string is the full path to a dhcpbridge configuration file.</li> </ul>

Configuration option = Default value	Type	Description
<b>disk_allocation_ratio = 0.0</b>	floating point value	<p>This option helps you specify virtual disk to physical disk allocation ratio.</p> <p>From Ocata (15.0.0) this is used to influence the hosts selected by the Placement API. Note that when Placement is used, the DiskFilter is redundant, because the Placement API will have already filtered out hosts that would have failed the DiskFilter.</p> <p>A ratio greater than 1.0 will result in over-subscription of the available physical disk, which can be useful for more efficiently packing instances created with images that do not use the entire virtual disk, such as sparse or compressed images. It can be set to a value between 0.0 and 1.0 in order to preserve a percentage of the disk for uses other than instances.</p> <p> <b>NOTE</b></p> <p>This can be set per-compute, or if set to 0.0, the value set on the scheduler node(s) or compute node(s) will be used and defaulted to 1.0. Once set to a non-default value, it is not possible to "unset" the config to get back to the default behavior. If you want to reset back to the default, explicitly specify 1.0.</p> <p> <b>NOTE</b></p> <p>As of the 16.0.0 Pike release, this configuration option is ignored for the ironic.IronicDriver compute driver and is hardcoded to 1.0.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any valid positive integer or float value</li> </ul>
<b>dmz_cidr = []</b>	list value	<p>This option is a list of zero or more IP address ranges in your network's DMZ that should be accepted.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>A list of strings, each of which should be a valid CIDR.</li> </ul>

Configuration option = Default value	Type	Description
<b>dns_server = []</b>	multi valued	<p>Despite the singular form of the name of this option, it is actually a list of zero or more server addresses that dnsmasq will use for DNS nameservers. If this is not empty, dnsmasq will not read <code>/etc/resolv.conf</code>, but will only use the servers specified in this option. If the option <code>use_network_dns_servers</code> is <code>True</code>, the <code>dns1</code> and <code>dns2</code> servers from the network will be appended to this list, and will be used as DNS servers, too.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A list of strings, where each string is either an IP address or a FQDN.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <b>use_network_dns_servers</b></li> </ul>
<b>dns_update_periodic_interval = -1</b>	integer value	<p>This option determines the time, in seconds, to wait between refreshing DNS entries for the network.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A positive integer</li> <li>• -1 to disable updates</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <b>use_neutron</b></li> </ul>
<code>`dnsmasq_config_file = `</code>	string value	<p>The path to the custom dnsmasq configuration file, if any.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• The full path to the configuration file, or an empty string if there is no custom dnsmasq configuration file.</li> </ul>

Configuration option = Default value	Type	Description
<b>ebtables_exec_attempts = 3</b>	integer value	<p>This option determines the number of times to retry ebtables commands before giving up. The minimum number of retries is 1.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any positive integer</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li><b>ebtables_retry_interval</b></li> </ul>
<b>ebtables_retry_interval = 1.0</b>	floating point value	<p>This option determines the time, in seconds, that the system will sleep in between ebtables retries. Note that each successive retry waits a multiple of this value, so for example, if this is set to the default of 1.0 seconds, and ebtables_exec_attempts is 4, after the first failure, the system will sleep for 1 * 1.0 seconds, after the second failure it will sleep 2 * 1.0 seconds, and after the third failure it will sleep 3 * 1.0 seconds.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any non-negative float or integer. Setting this to zero will result in no waiting between attempts.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>ebtables_exec_attempts</li> </ul>
<b>enable_network_quota = False</b>	boolean value	<p>This option is used to enable or disable quota checking for tenant networks.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>quota_networks</li> </ul>

Configuration option = Default value	Type	Description
<b>enable_new_services = True</b>	boolean value	<p>Enable new nova-compute services on this host automatically.</p> <p>When a new nova-compute service starts up, it gets registered in the database as an enabled service. Sometimes it can be useful to register new compute services in disabled state and then enabled them at a later point in time. This option only sets this behavior for nova-compute services, it does not auto-disable other services like nova-conductor, nova-scheduler, nova-consoleauth, or nova-osapi_compute.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● <b>True:</b> Each new compute service is enabled as soon as it registers itself.</li> <li>● <b>False:</b> Compute services must be enabled via an os-services REST API call or with the CLI with <b>nova service-enable &lt;hostname&gt; &lt;binary&gt;</b>, otherwise they are not ready to use.</li> </ul>
<b>enabled_apis = [u'osapi_compute', u'metadata']</b>	list value	List of APIs to be enabled by default.
<b>enabled_ssl_apis = []</b>	list value	<p>List of APIs with enabled SSL.</p> <p>Nova provides SSL support for the API servers. enabled_ssl_apis option allows configuring the SSL support.</p>
<b>executor_thread_pool_size = 64</b>	integer value	Size of executor thread pool when executor is threading or eventlet.
<b>fake_network = False</b>	boolean value	This option is used mainly in testing to avoid calls to the underlying network utilities.
<b>fatal_deprecations = False</b>	boolean value	Enables or disables fatal status of deprecations.



Configuration option = Default value	Type	Description
<b>firewall_driver = nova.virt.firewall.NoopFirewallDriver</b>	string value	<p>Firewall driver to use with <b>nova-network</b> service.</p> <p>This option only applies when using the <b>nova-network</b> service. When using another networking services, such as Neutron, this should be to set to the <b>nova.virt.firewall.NoopFirewallDriver</b>.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● <b>nova.virt.firewall.IptablesFirewallDriver</b></li> <li>● <b>nova.virt.firewall.NoopFirewallDriver</b></li> <li>● <b>nova.virt.libvirt.firewall.IptablesFirewallDriver</b></li> <li>● [...]</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>use_neutron</b>: This must be set to <b>False</b> to enable <b>nova-network</b> networking</li> </ul>
<b>fixed_ip_disassociate_timeout = 600</b>	integer value	<p>This is the number of seconds to wait before disassociating a deallocated fixed IP address. This is only used with the nova-network service, and has no effect when using neutron for networking.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any integer, zero or greater.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>use_neutron</b></li> </ul>
<b>fixed_range_v6 = fd00::/48</b>	string value	<p>This option determines the fixed IPv6 address block when creating a network.</p> <p>Please note that this option is only used when using nova-network instead of Neutron in your deployment.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any valid IPv6 CIDR</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>use_neutron</b></li> </ul>

Configuration option = Default value	Type	Description
<b>flat_injected = False</b>	boolean value	This option determines whether the network setup information is injected into the VM before it is booted. While it was originally designed to be used only by nova-network, it is also used by the vmware and xenapi virt drivers to control whether network information is injected into a VM. The libvirt virt driver also uses it when we use config_drive to configure network to control whether network information is injected into a VM.
<b>flat_interface = None</b>	string value	This option is the name of the virtual interface of the VM on which the bridge will be built. While it was originally designed to be used only by nova-network, it is also used by libvirt for the bridge interface name.  Possible values: <ul style="list-style-type: none"> <li>Any valid virtual interface name, such as <i>eth0</i></li> </ul>
<b>flat_network_bridge = None</b>	string value	This option determines the bridge used for simple network interfaces when no bridge is specified in the VM creation request.  Please note that this option is only used when using nova-network instead of Neutron in your deployment.  Possible values: <ul style="list-style-type: none"> <li>Any string representing a valid network bridge, such as <i>br100</i></li> </ul> Related options: <ul style="list-style-type: none"> <li><b>use_neutron</b></li> </ul>

Configuration option = Default value	Type	Description
<b>flat_network_dns = 8.8.4.4</b>	string value	<p>This is the address of the DNS server for a simple network. If this option is not specified, the default of 8.8.4.4 is used.</p> <p>Please note that this option is only used when using nova-network instead of Neutron in your deployment.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any valid IP address.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li><b>use_neutron</b></li> </ul>
<b>floating_ip_dns_manager = nova.network.noop_dns_driver.NoopDNSDriver</b>	string value	<p>Full class name for the DNS Manager for floating IPs.</p> <p>This option specifies the class of the driver that provides functionality to manage DNS entries associated with floating IPs.</p> <p>When a user adds a DNS entry for a specified domain to a floating IP, nova will add a DNS entry using the specified floating DNS driver. When a floating IP is deallocated, its DNS entry will automatically be deleted.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Full Python path to the class to be used</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>use_neutron: this options only works with nova-network.</li> </ul>

Configuration option = Default value	Type	Description
<b>force_config_drive = False</b>	boolean value	<p>Force injection to take place on a config drive</p> <p>When this option is set to true configuration drive functionality will be forced enabled by default, otherwise user can still enable configuration drives via the REST API or image metadata properties.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True: Force to use of configuration drive regardless the user's input in the REST API call.</li> <li>● False: Do not force use of configuration drive. Config drives can still be enabled via the REST API or image metadata properties.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● Use the <i>mkisofs_cmd</i> flag to set the path where you install the genisoimage program. If genisoimage is in same path as the nova-compute service, you do not need to set this flag.</li> <li>● To use configuration drive with Hyper-V, you must set the <i>mkisofs_cmd</i> value to the full path to an mkisofs.exe installation. Additionally, you must set the <i>qemu_img_cmd</i> value in the hyperv configuration section to the full path to an qemu-img command installation.</li> </ul>
<b>force_dhcp_release = True</b>	boolean value	<p>When this option is True, a call is made to release the DHCP for the instance when that instance is terminated.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>use_neutron</b></li> </ul>

Configuration option = Default value	Type	Description
<b>force_raw_images = True</b>	boolean value	<p>Force conversion of backing images to raw format.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• True: Backing image files will be converted to raw image format</li> <li>• False: Backing image files will not be converted</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <b>compute_driver</b>: Only the libvirt driver uses this option.</li> </ul>
<b>force_snat_range = []</b>	multi valued	<p>This is a list of zero or more IP ranges that traffic from the <b>routing_source_ip</b> will be SNATted to. If the list is empty, then no SNAT rules are created.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A list of strings, each of which should be a valid CIDR.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <b>routing_source_ip</b></li> </ul>
<b>forward_bridge_interface = [u'all']</b>	multi valued	<p>One or more interfaces that bridges can forward traffic to. If any of the items in this list is the special keyword <i>all</i>, then all traffic will be forwarded.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A list of zero or more interface names, or the word <i>all</i>.</li> </ul>

Configuration option = Default value	Type	Description
<b>gateway = None</b>	string value	<p>This is the default IPv4 gateway. It is used only in the testing suite.</p> <p>Please note that this option is only used when using nova-network instead of Neutron in your deployment.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Any valid IP address.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <b>use_neutron</b></li> <li>• <b>gateway_v6</b></li> </ul>
<b>gateway_v6 = None</b>	string value	<p>This is the default IPv6 gateway. It is used only in the testing suite.</p> <p>Please note that this option is only used when using nova-network instead of Neutron in your deployment.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Any valid IP address.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <b>use_neutron</b></li> <li>• <b>gateway</b></li> </ul>
<b>graceful_shutdown_timeout = 60</b>	integer value	Specify a timeout after which a gracefully shutdown server will exit. Zero value means endless wait.

Configuration option = Default value	Type	Description
<b>heal_instance_info_cache_interval = 60</b>	integer value	<p>Interval between instance network information cache updates.</p> <p>Number of seconds after which each compute node runs the task of querying Neutron for all of its instances networking information, then updates the Nova db with that information. Nova will never update it's cache if this option is set to 0. If we don't update the cache, the metadata service and nova-api endpoints will be proxying incorrect network data about the instance. So, it is not recommended to set this option to 0.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Any positive integer in seconds.</li> <li>• Any value <math>\neq 0</math> will disable the sync. This is not recommended.</li> </ul>
<b>host = ccosp-undercloud.usersys.redhat.com</b>	string value	<p>Hostname, FQDN or IP address of this host.</p> <p>Used as:</p> <ul style="list-style-type: none"> <li>• the oslo.messaging queue name for nova-compute worker</li> <li>• we use this value for the binding_host sent to neutron. This means if you use a neutron agent, it should have the same value for host.</li> <li>• cinder host attachment information</li> </ul> <p>Must be valid within AMQP key.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• String with hostname, FQDN or IP address. Default is hostname of this host.</li> </ul>
<b>image_cache_manager_interval = 2400</b>	integer value	<p>Number of seconds to wait between runs of the image cache manager.</p> <p>Possible values: * 0: run at the default rate. * -1: disable * Any other value</p>
<b>image_cache_subdirectory_name = _base</b>	string value	<p>Location of cached images.</p> <p>This is NOT the full path - just a folder name relative to <i>\$instances_path</i>. For per-compute-host cached images, set to <i>base\$my_ip</i></p>

Configuration option = Default value	Type	Description
<b>injected_network_template = \$pybasedir/nova/virt/interfaces.template</b>	string value	<p>Path to <code>/etc/network/interfaces</code> template.</p> <p>The path to a template file for the <code>/etc/network/interfaces</code>-style file, which will be populated by nova and subsequently used by cloudinit. This provides a method to configure network connectivity in environments without a DHCP server.</p> <p>The template will be rendered using Jinja2 template engine, and receive a top-level key called <b>interfaces</b>. This key will contain a list of dictionaries, one for each interface.</p> <p>Refer to the cloudinit documentaion for more information:</p> <p><a href="https://cloudinit.readthedocs.io/en/latest/topics/datasources.html">https://cloudinit.readthedocs.io/en/latest/topics/datasources.html</a></p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A path to a Jinja2-formatted template for a Debian <code>/etc/network/interfaces</code> file. This applies even if using a non Debian-derived guest.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <b>flat_inject</b>: This must be set to <b>True</b> to ensure nova embeds network configuration information in the metadata provided through the config drive.</li> </ul>
<b>instance_build_timeout = 0</b>	integer value	<p>Maximum time in seconds that an instance can take to build.</p> <p>If this timer expires, instance status will be changed to ERROR. Enabling this option will make sure an instance will not be stuck in BUILD state for a longer period.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• 0: Disables the option (default)</li> <li>• Any positive integer in seconds: Enables the option.</li> </ul>




Configuration option = Default value	Type	Description
<b>instance_delete_interval</b> = 300	integer value	<p>Interval for retrying failed instance file deletes.</p> <p>This option depends on <i>maximum_instance_delete_attempts</i>. This option specifies how often to retry deletes whereas <i>maximum_instance_delete_attempts</i> specifies the maximum number of retry attempts that can be made.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● 0: Will run at the default periodic interval.</li> <li>● Any value &lt; 0: Disables the option.</li> <li>● Any positive integer in seconds.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>maximum_instance_delete_attempts</b> from <i>instance_cleaning_opts</i> group.</li> </ul>
<code>instance_dns_domain = `</code>	string value	<p>If specified, Nova checks if the <i>availability_zone</i> of every instance matches what the database says the <i>availability_zone</i> should be for the specified <i>dns_domain</i>.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <i>use_neutron</i>: this options only works with <i>nova-network</i>.</li> </ul>
<b>instance_dns_manager = nova.network.noop_dns_driver.NoopDNSDriver</b>	string value	<p>Full class name for the DNS Manager for instance IPs.</p> <p>This option specifies the class of the driver that provides functionality to manage DNS entries for instances.</p> <p>On instance creation, nova will add DNS entries for the instance name and id, using the specified instance DNS driver and domain. On instance deletion, nova will remove the DNS entries.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Full Python path to the class to be used</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <i>use_neutron</i>: this options only works with <i>nova-network</i>.</li> </ul>

Configuration option = Default value	Type	Description
<code>`instance_format = [instance: %(uuid)s]`</code>	string value	The format for an instance that is passed with the log message.
<b>instance_name_template = instance-%08x</b>	string value	<p>Template string to be used to generate instance names.</p> <p>This template controls the creation of the database name of an instance. This is <b>not</b> the display name you enter when creating an instance (via Horizon or CLI). For a new deployment it is advisable to change the default value (which uses the database autoincrement) to another value which makes use of the attributes of an instance, like <b>instance-%(uuid)s</b>. If you already have instances in your deployment when you change this, your deployment will break.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A string which either uses the instance database ID (like the default)</li> <li>• A string with a list of named database columns, for example <b>%(id)d</b> or <b>%(uuid)s</b> or <b>%(hostname)s</b>.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• not to be confused with: <b>multi_instance_display_name_template</b></li> </ul>
<b>instance_usage_audit = False</b>	boolean value	This option enables periodic <code>compute.instance.exists</code> notifications. Each compute node must be configured to generate system usage data. These notifications are consumed by OpenStack Telemetry service.

Configuration option = Default value	Type	Description
<b>instance_usage_audit_period = month</b>	string value	<p>Time period to generate instance usages for. It is possible to define optional offset to given period by appending @ character followed by a number defining offset.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● period, example: <b>hour, day, month</b> or <b>year</b></li> <li>● period with offset, example: <b>month@15</b> will result in monthly audits starting on 15th day of month.</li> </ul>
<code>^instance_uuid_format = [instance:%(uuid)s]`</code>	string value	The format for an instance UUID that is passed with the log message.
<b>instances_path = \$state_path/instances</b>	string value	<p>Specifies where instances are stored on the hypervisor's disk. It can point to locally attached storage or a directory on NFS.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● <code>\$state_path/instances</code> where <code>state_path</code> is a config option that specifies the top-level directory for maintaining nova's state. (default) or Any string representing directory path.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>[workarounds]/ensure_libvirt_rbd_instance_dir_cleanup</b></li> </ul>
<b>internal_service_availability_zone = internal</b>	string value	<p>Availability zone for internal services.</p> <p>This option determines the availability zone for the various internal nova services, such as <i>nova-scheduler</i>, <i>nova-conductor</i>, etc.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any string representing an existing availability zone name.</li> </ul>

Configuration option = Default value	Type	Description
<code>`iptables_bottom_regex = `</code>	string value	<p>This expression, if defined, will select any matching iptables rules and place them at the bottom when applying metadata changes to the rules.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any string representing a valid regular expression, or an empty string</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li><code>iptables_top_regex</code></li> </ul>
<b><code>iptables_drop_action = DROP</code></b>	string value	<p>By default, packets that do not pass the firewall are DROPPed. In many cases, though, an operator may find it more useful to change this from DROP to REJECT, so that the user issuing those packets may have a better idea as to what's going on, or LOGDROP in order to record the blocked traffic before DROPPing.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>A string representing an iptables chain. The default is DROP.</li> </ul>
<code>`iptables_top_regex = `</code>	string value	<p>This expression, if defined, will select any matching iptables rules and place them at the top when applying metadata changes to the rules.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any string representing a valid regular expression, or an empty string</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li><b><code>iptables_bottom_regex</code></b></li> </ul>

Configuration option = Default value	Type	Description
<b>ipv6_backend = rfc2462</b>	string value	<p>Abstracts out IPv6 address generation to pluggable backends.</p> <p>nova-network can be put into dual-stack mode, so that it uses both IPv4 and IPv6 addresses. In dual-stack mode, by default, instances acquire IPv6 global unicast addresses with the help of stateless address auto-configuration mechanism.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>use_neutron</code>: this option only works with nova-network.</li> <li>● <code>use_ipv6</code>: this option only works if <code>ipv6</code> is enabled for nova-network.</li> </ul>
<b>key = None</b>	string value	SSL key file (if separate from cert).
<b>l3_lib = nova.network.l3.LinuxNetL3</b>	string value	<p>This option allows you to specify the L3 management library to be used.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any dot-separated string that represents the import path to an L3 networking library.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>use_neutron</b></li> </ul>
<b>ldap_dns_base_dn = ou=hosts,dc=example,dc=org</b>	string value	<p>Base distinguished name for the LDAP search query</p> <p>This option helps to decide where to look up the host in LDAP.</p>
<b>ldap_dns_password = password</b>	string value	Bind user's password for LDAP server
<b>ldap_dns_servers = [u'dns.example.org']</b>	multi valued	<p>DNS Servers for LDAP DNS driver</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A valid URL representing a DNS server</li> </ul>

Configuration option = Default value	Type	Description
<b>ldap_dns_soa_expiry = 86400</b>	integer value	<p>Expiry interval (in seconds) for LDAP DNS driver Start of Authority</p> <p>Time interval, a secondary/slave DNS server holds the information before it is no longer considered authoritative.</p>
<b>ldap_dns_soa_hostmaster = hostmaster@example.org</b>	string value	<p>Hostmaster for LDAP DNS driver Statement of Authority</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any valid string representing LDAP DNS hostmaster.</li> </ul>
<b>ldap_dns_soa_minimum = 7200</b>	integer value	<p>Minimum interval (in seconds) for LDAP DNS driver Start of Authority</p> <p>It is Minimum time-to-live applies for all resource records in the zone file. This value is supplied to other servers how long they should keep the data in cache.</p>
<b>ldap_dns_soa_refresh = 1800</b>	integer value	<p>Refresh interval (in seconds) for LDAP DNS driver Start of Authority</p> <p>Time interval, a secondary/slave DNS server waits before requesting for primary DNS server's current SOA record. If the records are different, secondary DNS server will request a zone transfer from primary.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p>Lower values would cause more traffic.</p> </div> </div>
<b>ldap_dns_soa_retry = 3600</b>	integer value	<p>Retry interval (in seconds) for LDAP DNS driver Start of Authority</p> <p>Time interval, a secondary/slave DNS server should wait, if an attempt to transfer zone failed during the previous refresh interval.</p>
<b>ldap_dns_url = ldap://ldap.example.com:389</b>	uri value	<p>URL for LDAP server which will store DNS entries</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>A valid LDAP URL representing the server</li> </ul>

Configuration option = Default value	Type	Description
<b>ldap_dns_user = uid=admin,ou=people,dc= example,dc=org</b>	string value	Bind user for LDAP server
<b>linuxnet_interface_driver = nova.network.linux_net.Li nuxBridgeInterfaceDriver</b>	string value	<p>This is the class used as the ethernet device driver for linuxnet bridge operations. The default value should be all you need for most cases, but if you wish to use a customized class, set this option to the full dot-separated import path for that class.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any string representing a dot-separated class path that Nova can import.</li> </ul>
<b>linuxnet_ovs_integration_ bridge = br-int</b>	string value	<p>The name of the Open vSwitch bridge that is used with linuxnet when connecting with Open vSwitch."</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any string representing a valid bridge name.</li> </ul>
<b>live_migration_retry_cou nt = 30</b>	integer value	<p>Maximum number of 1 second retries in live_migration. It specifies number of retries to iptables when it complains. It happens when an user continuously sends live-migration request to same host leading to concurrent request to iptables.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any positive integer representing retry count.</li> </ul>
<b>log-config-append = None</b>	string value	<p>The name of a logging configuration file. This file is appended to any existing logging configuration files. For details about logging configuration files, see the Python logging module documentation. Note that when logging configuration files are used then all logging configuration is set in the configuration file and other logging configuration options are ignored (for example, logging_context_format_string).</p>
<b>log-date-format = %Y-%m- %d %H:%M:%S</b>	string value	<p>Defines the format string for %(asctime)s in log records. Default: %(default)s . This option is ignored if log_config_append is set.</p>

Configuration option = Default value	Type	Description
<b>log-dir = None</b>	string value	(Optional) The base directory used for relative log_file paths. This option is ignored if log_config_append is set.
<b>log-file = None</b>	string value	(Optional) Name of log file to send logging output to. If no default is set, logging will go to stderr as defined by use_stderr. This option is ignored if log_config_append is set.
<b>log_options = True</b>	boolean value	Enables or disables logging values of all registered options when starting a service (at DEBUG level).
<b>logging_context_format_string = %(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [%(request_id)s %(user_identity)s] %(instance)s%(message)s</b>	string value	Format string to use for log messages with context.
<b>logging_debug_format_s uffix = %(funcName)s % (pathname)s:%(lineno)d</b>	string value	Additional data to append to log message when logging level for the message is DEBUG.
<b>logging_default_format_s tring = %(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [-] %(instance)s%(message)s</b>	string value	Format string to use for log messages when context is undefined.
<b>logging_exception_prefix = %(asctime)s.%(msecs)03d %(process)d ERROR %(name)s % (instance)s</b>	string value	Prefix each line of exception output with this format.
<b>logging_user_identity_for mat = %(user)s % (tenant)s %(domain)s % (user_domain)s % (project_domain)s</b>	string value	Defines the format string for %(user_identity)s that is used in logging_context_format_string.



Configuration option = Default value	Type	Description
<b>max_concurrent_builds = 10</b>	integer value	<p>Limits the maximum number of instance builds to run concurrently by nova-compute. Compute service can attempt to build an infinite number of instances, if asked to do so. This limit is enforced to avoid building unlimited instance concurrently on a compute node. This value can be set per compute node.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>● 0 : treated as unlimited.</li> <li>● Any positive integer representing maximum concurrent builds.</li> </ul>
<b>max_concurrent_live_migrations = 1</b>	integer value	<p>Maximum number of live migrations to run concurrently. This limit is enforced to avoid outbound live migrations overwhelming the host/network and causing failures. It is not recommended that you change this unless you are very sure that doing so is safe and stable in your environment.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● 0 : treated as unlimited.</li> <li>● Negative value defaults to 0.</li> <li>● Any positive integer representing maximum number of live migrations to run concurrently.</li> </ul>

Configuration option = Default value	Type	Description
<b>max_local_block_devices = 3</b>	integer value	<p>Maximum number of devices that will result in a local image being created on the hypervisor node.</p> <p>A negative number means unlimited. Setting <code>max_local_block_devices</code> to 0 means that any request that attempts to create a local disk will fail. This option is meant to limit the number of local discs (so root local disc that is the result of <code>--image</code> being used, and any other ephemeral and swap disks). 0 does not mean that images will be automatically converted to volumes and boot instances from volumes - it just means that all requests that attempt to create a local disk will fail.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● 0: Creating a local disk is not allowed.</li> <li>● Negative number: Allows unlimited number of local discs.</li> <li>● Positive number: Allows only these many number of local discs. (Default value is 3).</li> </ul>
<b>maximum_instance_delete_attempts = 5</b>	integer value	<p>The number of times to attempt to reap an instance's files.</p> <p>This option specifies the maximum number of retry attempts that can be made.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any positive integer defines how many attempts are made.</li> <li>● Any value <math>\leq 0</math> means no delete attempts occur, but you should use <b><code>instance_delete_interval</code></b> to disable the delete attempts.</li> </ul> <p>Related options: * <b><code>instance_delete_interval</code></b> in <code>interval_opts</code> group can be used to disable this option.</p>

Configuration option = Default value	Type	Description
<b>metadata_host = \$my_ip</b>	string value	<p>This option determines the IP address for the network metadata API server.</p> <p>This is really the client side of the metadata host equation that allows nova-network to find the metadata server when doing a default multi host networking.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any valid IP address. The default is the address of the Nova API server.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li><b>metadata_port</b></li> </ul>
<b>metadata_listen = 0.0.0.0</b>	string value	<p>IP address on which the metadata API will listen.</p> <p>The metadata API service listens on this IP address for incoming requests.</p>
<b>metadata_listen_port = 8775</b>	port value	<p>Port on which the metadata API will listen.</p> <p>The metadata API service listens on this port number for incoming requests.</p>
<b>metadata_port = 8775</b>	port value	<p>This option determines the port used for the metadata API server.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li><b>metadata_host</b></li> </ul>
<b>metadata_workers = None</b>	integer value	<p>Number of workers for metadata service. If not specified the number of available CPUs will be used.</p> <p>The metadata service can be configured to run as multi-process (workers). This overcomes the problem of reduction in throughput when API request concurrency increases. The metadata service will run in the specified number of processes.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>Any positive integer</li> <li>None (default value)</li> </ul>

Configuration option = Default value	Type	Description
<b>migrate_max_retries = -1</b>	integer value	<p>Number of times to retry live-migration before failing.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● If == -1, try until out of hosts (default)</li> <li>● If == 0, only try once, no retries</li> <li>● Integer greater than 0</li> </ul>
<b>mkisofs_cmd = genisoimage</b>	string value	<p>Name or path of the tool used for ISO image creation</p> <p>Use the <code>mkisofs_cmd</code> flag to set the path where you install the <code>genisoimage</code> program. If <code>genisoimage</code> is on the system path, you do not need to change the default value.</p> <p>To use configuration drive with Hyper-V, you must set the <code>mkisofs_cmd</code> value to the full path to an <code>mkisofs.exe</code> installation. Additionally, you must set the <code>qemu_img_cmd</code> value in the <code>hyperv</code> configuration section to the full path to an <code>qemu-img</code> command installation.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Name of the ISO image creator program, in case it is in the same directory as the <code>nova-compute</code> service</li> <li>● Path to ISO image creator program</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● This option is meaningful when config drives are enabled.</li> <li>● To use configuration drive with Hyper-V, you must set the <code>qemu_img_cmd</code> value in the <code>hyperv</code> configuration section to the full path to an <code>qemu-img</code> command installation.</li> </ul>
<b>monkey_patch = False</b>	boolean value	<p>Determine if monkey patching should be applied.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>monkey_patch_modules:</b> This must have values set for this option to have any effect</li> </ul>

Configuration option = Default value	Type	Description
<b>monkey_patch_modules</b> = [u'nova.compute.api:nova.notifications.notify_decorator']	list value	<p>List of modules/decorators to monkey patch.</p> <p>This option allows you to patch a decorator for all functions in specified modules.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• nova.compute.api:nova.notifications.notify_decorator</li> <li>• [...]</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <b>monkey_patch</b>: This must be set to <b>True</b> for this option to have any effect</li> </ul>
<b>multi_host</b> = False	boolean value	<p>Default value for multi_host in networks.</p> <p>nova-network service can operate in a multi-host or single-host mode. In multi-host mode each compute node runs a copy of nova-network and the instances on that compute node use the compute node as a gateway to the Internet. Where as in single-host mode, a central server runs the nova-network service. All compute nodes forward traffic from the instances to the cloud controller which then forwards traffic to the Internet.</p> <p>If this options is set to true, some rpc network calls will be sent directly to host.</p> <p>Note that this option is only used when using nova-network instead of Neutron in your deployment.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <b>use_neutron</b></li> </ul>
<b>multi_instance_display_name_template</b> = %(name)s-%(count)d	string value	<p>When creating multiple instances with a single request using the os-multiple-create API extension, this template will be used to build the display name for each instance. The benefit is that the instances end up with different hostnames. Example display names when creating two VM's: name-1, name-2.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Valid keys for the template are: name, uuid, count.</li> </ul>

Configuration option = Default value	Type	Description
<b>my_block_storage_ip = \$my_ip</b>	string value	<p>The IP address which is used to connect to the block storage network.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● String with valid IP address. Default is IP address of this host.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● my_ip - if my_block_storage_ip is not set, then my_ip value is used.</li> </ul>
<b>my_ip = 10.65.176.99</b>	string value	<p>The IP address which the host is using to connect to the management network.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● String with valid IP address. Default is IPv4 address of this host.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● metadata_host</li> <li>● my_block_storage_ip</li> <li>● routing_source_ip</li> <li>● vpn_ip</li> </ul>
<b>network_allocate_retries = 0</b>	integer value	<p>Number of times to retry network allocation. It is required to attempt network allocation retries if the virtual interface plug fails.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any positive integer representing retry count.</li> </ul>

Configuration option = Default value	Type	Description
<b>network_driver = nova.network.linux_net</b>	string value	<p>Driver to use for network creation.</p> <p>Network driver initializes (creates bridges and so on) only when the first VM lands on a host node. All network managers configure the network using network drivers. The driver is not tied to any particular network manager.</p> <p>The default Linux driver implements vlans, bridges, and iptables rules using linux utilities.</p> <p>Note that this option is only used when using nova-network instead of Neutron in your deployment.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>use_neutron</b></li> </ul>
<b>network_manager = nova.network.manager.VI anManager</b>	string value	Full class name for the Manager for network
<b>network_size = 256</b>	integer value	<p>This option determines the number of addresses in each private subnet.</p> <p>Please note that this option is only used when using nova-network instead of Neutron in your deployment.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any positive integer that is less than or equal to the available network size. Note that if you are creating multiple networks, they must all fit in the available IP address space. The default is 256.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>use_neutron</b></li> <li>● <b>num_networks</b></li> </ul>
<b>networks_path = \$state_path/networks</b>	string value	<p>The location where the network configuration files will be kept. The default is the <i>networks</i> directory off of the location where nova's Python module is installed.</p> <p>Possible values</p> <ul style="list-style-type: none"> <li>● A string containing the full path to the desired configuration directory</li> </ul>

Configuration option = Default value	Type	Description
<b>non_inheritable_image_properties =</b> [u'cache_in_nova', u'bittorrent', u'img_signature_hash_m ethod', u'img_signature', u'img_signature_key_typ e', u'img_signature_certificat e_uuid']	list value	<p>Image properties that should not be inherited from the instance when taking a snapshot.</p> <p>This option gives an opportunity to select which image-properties should not be inherited by newly created snapshots.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A comma-separated list whose item is an image property. Usually only the image properties that are only needed by base images can be included here, since the snapshots that are created from the base images don't need them.</li> <li>• Default list: cache_in_nova, bittorrent, img_signature_hash_method, img_signature, img_signature_key_type, img_signature_certificate_uuid</li> </ul>
<b>num_networks = 1</b>	integer value	<p>This option represents the number of networks to create if not explicitly specified when the network is created. The only time this is used is if a CIDR is specified, but an explicit network_size is not. In that case, the subnets are created by dividing the IP address space of the CIDR by num_networks. The resulting subnet sizes cannot be larger than the configuration option <b>network_size</b>; in that event, they are reduced to <b>network_size</b>, and a warning is logged.</p> <p>Please note that this option is only used when using nova-network instead of Neutron in your deployment.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Any positive integer is technically valid, although there are practical limits based upon available IP address space and virtual interfaces.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <b>use_neutron</b></li> <li>• <b>network_size</b></li> </ul>





Configuration option = Default value	Type	Description
<b>osapi_compute_listen = 0.0.0.0</b>	string value	<p>IP address on which the OpenStack API will listen.</p> <p>The OpenStack API service listens on this IP address for incoming requests.</p>
<b>osapi_compute_listen_port = 8774</b>	port value	<p>Port on which the OpenStack API will listen.</p> <p>The OpenStack API service listens on this port number for incoming requests.</p>
<code>osapi_compute_unique_server_name_scope = `</code>	string value	<p>Sets the scope of the check for unique instance names.</p> <p>The default doesn't check for unique names. If a scope for the name check is set, a launch of a new instance or an update of an existing instance with a duplicate name will result in an <i>'InstanceExists'</i> error. The uniqueness is case-insensitive. Setting this option can increase the usability for end users as they don't have to distinguish among instances with the same name by their IDs.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• "": An empty value means that no uniqueness check is done and duplicate names are possible.</li> <li>• "project": The instance name check is done only for instances within the same project.</li> <li>• "global": The instance name check is done for all instances regardless of the project.</li> </ul>
<b>osapi_compute_workers = None</b>	integer value	<p>Number of workers for OpenStack API service. The default will be the number of CPUs available.</p> <p>OpenStack API services can be configured to run as multi-process (workers). This overcomes the problem of reduction in throughput when API request concurrency increases. OpenStack API service will run in the specified number of processes.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>• Any positive integer</li> <li>• None (default value)</li> </ul>

Configuration option = Default value	Type	Description
<b>ovs_vsctl_timeout = 120</b>	integer value	<p>This option represents the period of time, in seconds, that the <code>ovs_vsctl</code> calls will wait for a response from the database before timing out. A setting of 0 means that the utility should wait forever for a response.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any positive integer if a limited timeout is desired, or zero if the calls should wait forever for a response.</li> </ul>
<b>password_length = 12</b>	integer value	Length of generated instance admin passwords.
<b>periodic_enable = True</b>	boolean value	<p>Enable periodic tasks.</p> <p>If set to true, this option allows services to periodically run tasks on the manager.</p> <p>In case of running multiple schedulers or conductors you may want to run periodic tasks on only one host - in this case disable this option for all hosts but one.</p>
<b>periodic_fuzzy_delay = 60</b>	integer value	<p>Number of seconds to randomly delay when starting the periodic task scheduler to reduce stampeding.</p> <p>When compute workers are restarted in unison across a cluster, they all end up running the periodic tasks at the same time causing problems for the external services. To mitigate this behavior, <code>periodic_fuzzy_delay</code> option allows you to introduce a random initial delay when starting the periodic task scheduler.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>Any positive integer (in seconds)</li> <li>0: disable the random delay</li> </ul>

Configuration option = Default value	Type	Description
<b>pointer_model = usbtablet</b>	string value	<p>Generic property to specify the pointer type.</p> <p>Input devices allow interaction with a graphical framebuffer. For example to provide a graphic tablet for absolute cursor movement.</p> <p>If set, the <i>hw_pointer_model</i> image property takes precedence over this configuration option.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● None: Uses default behavior provided by drivers (mouse on PS2 for libvirt x86)</li> <li>● ps2mouse: Uses relative movement. Mouse connected by PS2</li> <li>● usbtablet: Uses absolute movement. Tablet connect by USB</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● usbtablet must be configured with VNC enabled or SPICE enabled and SPICE agent disabled. When used with libvirt the instance mode should be configured as HVM.</li> </ul>
<b>preallocate_images = none</b>	string value	<p>The image preallocation mode to use.</p> <p>Image preallocation allows storage for instance images to be allocated up front when the instance is initially provisioned. This ensures immediate feedback is given if enough space isn't available. In addition, it should significantly improve performance on writes to new blocks and may even improve I/O performance to prewritten blocks due to reduced fragmentation.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● "none" ⇒ no storage provisioning is done up front</li> <li>● "space" ⇒ storage is fully allocated at instance start</li> </ul>
<b>public_interface = eth0</b>	string value	<p>This is the name of the network interface for public IP addresses. The default is <i>eth0</i>.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any string representing a network interface name</li> </ul>

Configuration option = Default value	Type	Description
<b>publish_errors = False</b>	boolean value	Enables or disables publication of error events.
<b>pybasedir = /usr/lib/python2.7/site- packages</b>	string value	<p>The directory where the Nova python modules are installed.</p> <p>This directory is used to store template files for networking and remote console access. It is also the default path for other config options which need to persist Nova internal data. It is very unlikely that you need to change this option from its default value.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• The full path to a directory.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <b>state_path</b></li> </ul>
<b>quota_networks = 3</b>	integer value	<p>This option controls the number of private networks that can be created per project (or per tenant).</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>• enable_network_quota</li> </ul>

Configuration option = Default value	Type	Description
<b>ram_allocation_ratio = 0.0</b>	floating point value	<p>This option helps you specify virtual RAM to physical RAM allocation ratio.</p> <p>From Ocata (15.0.0) this is used to influence the hosts selected by the Placement API. Note that when Placement is used, the RamFilter is redundant, because the Placement API will have already filtered out hosts that would have failed the RamFilter.</p> <p>This configuration specifies ratio for RamFilter which can be set per compute node. For AggregateRamFilter, it will fall back to this configuration value if no per-aggregate setting found.</p> <p> <b>NOTE</b></p> <p>This can be set per-compute, or if set to 0.0, the value set on the scheduler node(s) or compute node(s) will be used and defaulted to 1.5. Once set to a non-default value, it is not possible to "unset" the config to get back to the default behavior. If you want to reset back to the default, explicitly specify 1.5.</p> <p> <b>NOTE</b></p> <p>As of the 16.0.0 Pike release, this configuration option is ignored for the ironic.IronicDriver compute driver and is hardcoded to 1.0.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any valid positive integer or float value</li> </ul>
<b>rate_limit_burst = 0</b>	integer value	Maximum number of logged messages per rate_limit_interval.

Configuration option = Default value	Type	Description
<b>rate_limit_except_level = CRITICAL</b>	string value	Log level name used by rate limiting: CRITICAL, ERROR, INFO, WARNING, DEBUG or empty string. Logs with level greater or equal to rate_limit_except_level are not filtered. An empty string means that all levels are filtered.
<b>rate_limit_interval = 0</b>	integer value	Interval, number of seconds, of log rate limiting.
<b>reboot_timeout = 0</b>	integer value	<p>Time interval after which an instance is hard rebooted automatically.</p> <p>When doing a soft reboot, it is possible that a guest kernel is completely hung in a way that causes the soft reboot task to not ever finish. Setting this option to a time period in seconds will automatically hard reboot an instance if it has been stuck in a rebooting state longer than N seconds.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● 0: Disables the option (default).</li> <li>● Any positive integer in seconds: Enables the option.</li> </ul>

Configuration option = Default value	Type	Description
<b>reclaim_instance_interval</b> = 0	integer value	<p>Interval for reclaiming deleted instances.</p> <p>A value greater than 0 will enable SOFT_DELETE of instances. This option decides whether the server to be deleted will be put into the SOFT_DELETED state. If this value is greater than 0, the deleted server will not be deleted immediately, instead it will be put into a queue until it's too old (deleted time greater than the value of reclaim_instance_interval). The server can be recovered from the delete queue by using the restore action. If the deleted server remains longer than the value of reclaim_instance_interval, it will be deleted by a periodic task in the compute service automatically.</p> <p>Note that this option is read from both the API and compute nodes, and must be set globally otherwise servers could be put into a soft deleted state in the API and never actually reclaimed (deleted) on the compute node.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any positive integer(in seconds) greater than 0 will enable this option.</li> <li>Any value <math>\leq 0</math> will disable the option.</li> </ul>
<b>record</b> = None	string value	Filename that will be used for storing websocket frames received and sent by a proxy service (like VNC, spice, serial) running on this host. If this is not set, no recording will be done.
<b>remove_unused_base_images</b> = True	boolean value	Should unused base images be removed?
<b>remove_unused_original_minimum_age_seconds</b> = 86400	integer value	Unused unresized base images younger than this will not be removed.

Configuration option = Default value	Type	Description
<b>report_interval = 10</b>	integer value	<p>Number of seconds indicating how frequently the state of services on a given hypervisor is reported. Nova needs to know this to determine the overall health of the deployment.</p> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● <code>service_down_time report_interval</code> should be less than <code>service_down_time</code>. If <code>service_down_time</code> is less than <code>report_interval</code>, services will routinely be considered down, because they report in too rarely.</li> </ul>
<b>rescue_timeout = 0</b>	integer value	<p>Interval to wait before un-rescuing an instance stuck in RESCUE.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● 0: Disables the option (default)</li> <li>● Any positive integer in seconds: Enables the option.</li> </ul>
<b>reserved_host_cpus = 0</b>	integer value	<p>Number of physical CPUs to reserve for the host. The host resources usage is reported back to the scheduler continuously from nova-compute running on the compute node. To prevent the host CPU from being considered as available, this option is used to reserve random pCPU(s) for the host.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any positive integer representing number of physical CPUs to reserve for the host.</li> </ul>
<b>reserved_host_disk_mb = 0</b>	integer value	<p>Amount of disk resources in MB to make them always available to host. The disk usage gets reported back to the scheduler from nova-compute running on the compute nodes. To prevent the disk resources from being considered as available, this option can be used to reserve disk space for that host.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any positive integer representing amount of disk in MB to reserve for the host.</li> </ul>



Configuration option = Default value	Type	Description
<b>reserved_host_memory_mb = 512</b>	integer value	<p>Amount of memory in MB to reserve for the host so that it is always available to host processes. The host resources usage is reported back to the scheduler continuously from nova-compute running on the compute node. To prevent the host memory from being considered as available, this option is used to reserve memory for the host.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any positive integer representing amount of memory in MB to reserve for the host.</li> </ul>
<b>reserved_huge_pages = None</b>	dict value	<p>Number of huge/large memory pages to reserved per NUMA host cell.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>A list of valid key=value which reflect NUMA node ID, page size (Default unit is KiB) and number of pages to be reserved.</li> </ul> <pre>reserved_huge_pages = node:0,size:2048,count:64 reserved_huge_pages = node:1,size:1GB,count:1</pre> <p>In this example we are reserving on NUMA node 0 64 pages of 2MiB and on NUMA node 1 1 page of 1GiB.</p>
<b>resize_confirm_window = 0</b>	integer value	<p>Automatically confirm resizes after N seconds.</p> <p>Resize functionality will save the existing server before resizing. After the resize completes, user is requested to confirm the resize. The user has the opportunity to either confirm or revert all changes. Confirm resize removes the original server and changes server status from resized to active. Setting this option to a time period (in seconds) will automatically confirm the resize if the server is in resized state longer than that time.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>0: Disables the option (default)</li> <li>Any positive integer in seconds: Enables the option.</li> </ul>

Configuration option = Default value	Type	Description
<b>resize_fs_using_block_device = False</b>	boolean value	<p>Enable resizing of filesystems via a block device.</p> <p>If enabled, attempt to resize the filesystem by accessing the image over a block device. This is done by the host and may not be necessary if the image contains a recent version of cloud-init. Possible mechanisms require the nbd driver (for qcow and raw), or loop (for raw).</p>
<b>resume_guests_state_on_host_boot = False</b>	boolean value	<p>This option specifies whether to start guests that were running before the host rebooted. It ensures that all of the instances on a Nova compute node resume their state each time the compute node boots or restarts.</p>
<b>rootwrap_config = /etc/nova/rootwrap.conf</b>	string value	<p>Path to the rootwrap configuration file.</p> <p>Goal of the root wrapper is to allow a service-specific unprivileged user to run a number of actions as the root user in the safest manner possible. The configuration file used here must match the one defined in the sudoers entry.</p>
<b>routing_source_ip = \$my_ip</b>	string value	<p>The public IP address of the network host.</p> <p>This is used when creating an SNAT rule.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any valid IP address</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li><b>force_snat_range</b></li> </ul>
<b>rpc_ack_timeout_base = 15</b>	integer value	<p>Number of seconds to wait for an ack from a cast/call. After each retry attempt this timeout is multiplied by some specified multiplier.</p>
<b>rpc_ack_timeout_multiplier = 2</b>	integer value	<p>Number to multiply base ack timeout by after each retry attempt.</p>
<b>rpc_backend = rabbit</b>	string value	<p>The messaging driver to use, defaults to rabbit. Other drivers include amqp and zmq.</p>
<b>rpc_conn_pool_size = 30</b>	integer value	<p>Size of RPC connection pool.</p>

Configuration option = Default value	Type	Description
<b>rpc_message_ttl = 300</b>	integer value	Expiration timeout in seconds of a sent/received message after which it is not tracked anymore by a client/server.
<b>rpc_poll_timeout = 1</b>	integer value	The default number of seconds that poll should wait. Poll raises timeout exception when timeout expired.
<b>rpc_response_timeout = 60</b>	integer value	Seconds to wait for a response from a call.
<b>rpc_retry_attempts = 3</b>	integer value	Default number of message sending attempts in case of any problems occurred: positive value N means at most N retries, 0 means no retries, None or -1 (or any other negative values) mean to retry forever. This option is used only if acknowledgments are enabled.
<b>rpc_thread_pool_size = 100</b>	integer value	Maximum number of (green) threads to work concurrently.
<b>rpc_use_acks = False</b>	boolean value	Wait for message acknowledgements from receivers. This mechanism works only via proxy without PUB/SUB.
<b>rpc_zmq_bind_address = *</b>	string value	ZeroMQ bind address. Should be a wildcard (*), an ethernet interface, or IP. The "host" option should point or resolve to this address.
<b>rpc_zmq_bind_port_retries = 100</b>	integer value	Number of retries to find free port number before fail with ZMQBindError.
<b>rpc_zmq_contexts = 1</b>	integer value	Number of ZeroMQ contexts, defaults to 1.
<b>rpc_zmq_host = ccosp-undercloud.usersys.redhat.com</b>	string value	Name of this node. Must be a valid hostname, FQDN, or IP address. Must match "host" option, if running Nova.
<b>rpc_zmq_ipc_dir = /var/run/openstack</b>	string value	Directory for holding IPC sockets.
<b>rpc_zmq_matchmaker = redis</b>	string value	MatchMaker driver.
<b>rpc_zmq_max_port = 65536</b>	integer value	Maximal port number for random ports range.
<b>rpc_zmq_min_port = 49153</b>	port value	Minimal port number for random ports range.

Configuration option = Default value	Type	Description
<b>rpc_zmq_serialization = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>rpc_zmq_topic_backlog = None</b>	integer value	Maximum number of ingress messages to locally buffer per topic. Default is unlimited.
<b>run_external_periodic_tasks = True</b>	boolean value	Some periodic tasks can be run in a separate process. Should we run them here?
<b>running_deleted_instance_action = reap</b>	string value	<p>The compute service periodically checks for instances that have been deleted in the database but remain running on the compute node. The above option enables action to be taken when such instances are identified.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● reap: Powers down the instances and deletes them(default)</li> <li>● log: Logs warning message about deletion of the resource</li> <li>● shutdown: Powers down instances and marks them as non- bootable which can be later used for debugging/analysis</li> <li>● noop: Takes no action</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● running_deleted_instance_poll_interval</li> <li>● running_deleted_instance_timeout</li> </ul>
<b>running_deleted_instance_poll_interval = 1800</b>	integer value	<p>Time interval in seconds to wait between runs for the clean up action. If set to 0, above check will be disabled. If "running_deleted_instance_action" is set to "log" or "reap", a value greater than 0 must be set.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any positive integer in seconds enables the option.</li> <li>● 0: Disables the option.</li> <li>● 1800: Default value.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● running_deleted_instance_action</li> </ul>

Configuration option = Default value	Type	Description
<b>running_deleted_instance_timeout = 0</b>	integer value	<p>Time interval in seconds to wait for the instances that have been marked as deleted in database to be eligible for cleanup.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any positive integer in seconds(default is 0).</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>"running_deleted_instance_action"</li> </ul>
<b>scheduler_instance_sync_interval = 120</b>	integer value	<p>Interval between sending the scheduler a list of current instance UUIDs to verify that its view of instances is in sync with nova.</p> <p>If the CONF option <i>scheduler_tracks_instance_changes</i> is False, the sync calls will not be made. So, changing this option will have no effect.</p> <p>If the out of sync situations are not very common, this interval can be increased to lower the number of RPC messages being sent. Likewise, if sync issues turn out to be a problem, the interval can be lowered to check more frequently.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>0: Will run at the default periodic interval.</li> <li>Any value &lt; 0: Disables the option.</li> <li>Any positive integer in seconds.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>This option has no impact if <b>scheduler_tracks_instance_changes</b> is set to False.</li> </ul>
<b>send_arp_for_ha = False</b>	boolean value	<p>When True, when a device starts up, and upon binding floating IP addresses, arp messages will be sent to ensure that the arp caches on the compute hosts are up-to-date.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li><b>send_arp_for_ha_count</b></li> </ul>

Configuration option = Default value	Type	Description
<b>send_arp_for_ha_count = 3</b>	integer value	<p>When arp messages are configured to be sent, they will be sent with the count set to the value of this option. Of course, if this is set to zero, no arp messages will be sent.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any integer greater than or equal to 0</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li><b>send_arp_for_ha</b></li> </ul>
<b>service_down_time = 60</b>	integer value	<p>Maximum time in seconds since last check-in for up service</p> <p>Each compute node periodically updates their database status based on the specified report interval. If the compute node hasn't updated the status for more than <code>service_down_time</code>, then the compute node is considered down.</p> <p>Related Options:</p> <ul style="list-style-type: none"> <li><code>report_interval</code> (<code>service_down_time</code> should not be less than <code>report_interval</code>)</li> </ul>
<b>servicegroup_driver = db</b>	string value	<p>This option specifies the driver to be used for the servicegroup service.</p> <p>ServiceGroup API in nova enables checking status of a compute node. When a compute worker running the nova-compute daemon starts, it calls the join API to join the compute group. Services like nova scheduler can query the ServiceGroup API to check if a node is alive. Internally, the ServiceGroup client driver automatically updates the compute worker status. There are multiple backend implementations for this service: Database ServiceGroup driver and Memcache ServiceGroup driver.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li><code>db</code> : Database ServiceGroup driver</li> <li><code>mc</code> : Memcache ServiceGroup driver</li> </ul> <p>Related Options:</p> <ul style="list-style-type: none"> <li><code>service_down_time</code> (maximum time since last check-in for up service)</li> </ul>

Configuration option = Default value	Type	Description
<b>share_dhcp_address = False</b>	boolean value	<p>THIS VALUE SHOULD BE SET WHEN CREATING THE NETWORK.</p> <p>If True in multi_host mode, all compute hosts share the same dhcp address. The same IP address used for DHCP will be added on each nova-network node which is only visible to the VMs on the same host.</p> <p>The use of this configuration has been deprecated and may be removed in any release after Mitaka. It is recommended that instead of relying on this option, an explicit value should be passed to <i>create_networks()</i> as a keyword argument with the name <i>share_address</i>.</p>
<b>shelved_offload_time = 0</b>	integer value	<p>Time before a shelved instance is eligible for removal from a host.</p> <p>By default this option is set to 0 and the shelved instance will be removed from the hypervisor immediately after shelve operation. Otherwise, the instance will be kept for the value of <i>shelved_offload_time</i>(in seconds) so that during the time period the unshelve action will be faster, then the periodic task will remove the instance from hypervisor after <i>shelved_offload_time</i> passes.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● 0: Instance will be immediately offloaded after being shelved.</li> <li>● Any value &lt; 0: An instance will never offload.</li> <li>● Any positive integer in seconds: The instance will exist for the specified number of seconds before being offloaded.</li> </ul>

Configuration option = Default value	Type	Description
<b>shelved_poll_interval = 3600</b>	integer value	<p>Interval for polling shelved instances to offload.</p> <p>The periodic task runs for every <code>shelved_poll_interval</code> number of seconds and checks if there are any shelved instances. If it finds a shelved instance, based on the <code>shelved_offload_time</code> config value it offloads the shelved instances. Check <code>shelved_offload_time</code> config option description for details.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any value <math>\neq</math> 0: Disables the option.</li> <li>Any positive integer in seconds.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li><b>shelved_offload_time</b></li> </ul>
<b>shutdown_timeout = 60</b>	integer value	<p>Total time to wait in seconds for an instance to perform a clean shutdown.</p> <p>It determines the overall period (in seconds) a VM is allowed to perform a clean shutdown. While performing stop, rescue and shelve, rebuild operations, configuring this option gives the VM a chance to perform a controlled shutdown before the instance is powered off. The default timeout is 60 seconds.</p> <p>The timeout value can be overridden on a per image basis by means of <code>os_shutdown_timeout</code> that is an image metadata setting allowing different types of operating systems to specify how much time they need to shut down cleanly.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any positive integer in seconds (default value is 60).</li> </ul>
<b>source_is_ipv6 = False</b>	boolean value	Set to True if source host is addressed with IPv6.
<b>ssl_only = False</b>	boolean value	Disallow non-encrypted connections.



Configuration option = Default value	Type	Description
<b>state_path = \$pybasedir</b>	string value	<p>The top-level directory for maintaining Nova's state.</p> <p>This directory is used to store Nova's internal state. It is used by a variety of other config options which derive from this. In some scenarios (for example migrations) it makes sense to use a storage location which is shared between multiple compute hosts (for example via NFS). Unless the option <b>instances_path</b> gets overwritten, this directory can grow very large.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• The full path to a directory. Defaults to value provided in <b>pybasedir</b>.</li> </ul>
<b>subscribe_on = []</b>	list value	<p>List of publisher hosts SubConsumer can subscribe on. This option has higher priority then the default publishers list taken from the matchmaker.</p>
<b>sync_power_state_interval = 600</b>	integer value	<p>Interval to sync power states between the database and the hypervisor.</p> <p>The interval that Nova checks the actual virtual machine power state and the power state that Nova has in its database. If a user powers down their VM, Nova updates the API to report the VM has been powered down. Should something turn on the VM unexpectedly, Nova will turn the VM back off to keep the system in the expected state.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• 0: Will run at the default periodic interval.</li> <li>• Any value &lt; 0: Disables the option.</li> <li>• Any positive integer in seconds.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• If <b>handle_virt_lifecycle_events</b> in <code>workarounds_group</code> is false and this option is negative, then instances that get out of sync between the hypervisor and the Nova database will have to be synchronized manually.</li> </ul>

Configuration option = Default value	Type	Description
<b>sync_power_state_pool_size = 1000</b>	integer value	<p>Number of greenthreads available for use to sync power states.</p> <p>This option can be used to reduce the number of concurrent requests made to the hypervisor or system with real instance power states for performance reasons, for example, with Ironic.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any positive integer representing greenthreads count.</li> </ul>
<b>syslog-log-facility = LOG_USER</b>	string value	Syslog facility to receive log lines. This option is ignored if log_config_append is set.
<b>teardown_unused_network_gateway = False</b>	boolean value	<p>Determines whether unused gateway devices, both VLAN and bridge, are deleted if the network is in nova-network VLAN mode and is multi-hosted.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li><b>use_neutron</b></li> <li><b>vpn_ip</b></li> <li><b>fake_network</b></li> </ul>
<b>tempdir = None</b>	string value	Explicitly specify the temporary working directory.
<b>timeout_nbd = 10</b>	integer value	Amount of time, in seconds, to wait for NBD device start up.
<b>transport_url = None</b>	string value	<p>The network address and optional user credentials for connecting to the messaging backend, in URL format. The expected format is:</p> <pre>driver://[user:pass@]host:port[, [userN:passN@]hostN:portN]/virtual_host?query</pre> <p>Example:</p> <pre>rabbit://rabbitmq:password@127.0.0.1:5672//</pre> <p>For full details on the fields in the URL see the documentation of oslo_messaging.TransportURL at <a href="https://docs.openstack.org/oslo.messaging/latest/reference/transport.html">https://docs.openstack.org/oslo.messaging/latest/reference/transport.html</a></p>

Configuration option = Default value	Type	Description
<b>update_dns_entries = False</b>	boolean value	<p>When this option is True, whenever a DNS entry must be updated, a fanout cast message is sent to all network hosts to update their DNS entries in multi-host mode.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>use_neutron</b></li> </ul>
<b>update_resources_interval = 0</b>	integer value	<p>Interval for updating compute resources.</p> <p>This option specifies how often the <code>update_available_resources</code> periodic task should run. A number less than 0 means to disable the task completely. Leaving this at the default of 0 will cause this to run at the default periodic interval. Setting it to any positive value will cause it to run at approximately that number of seconds.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● 0: Will run at the default periodic interval.</li> <li>● Any value &lt; 0: Disables the option.</li> <li>● Any positive integer in seconds.</li> </ul>
<b>use-journal = False</b>	boolean value	<p>Enable journald for logging. If running in a systemd environment you may wish to enable journal support. Doing so will use the journal native protocol which includes structured metadata in addition to log messages. This option is ignored if <code>log_config_append</code> is set.</p>
<b>use-json = False</b>	boolean value	<p>Use JSON formatting for logging. This option is ignored if <code>log_config_append</code> is set.</p>
<b>use-syslog = False</b>	boolean value	<p>Use syslog for logging. Existing syslog format is DEPRECATED and will be changed later to honor RFC5424. This option is ignored if <code>log_config_append</code> is set.</p>
<b>use_cow_images = True</b>	boolean value	<p>Enable use of copy-on-write (cow) images.</p> <p>QEMU/KVM allow the use of qcow2 as backing files. By disabling this, backing files will not be used.</p>

Configuration option = Default value	Type	Description
<b>use_dynamic_connections = False</b>	boolean value	This option makes direct connections dynamic or static. It makes sense only with <code>use_router_proxy=False</code> which means to use direct connections for direct message types (ignored otherwise).
<b>use_ipv6 = False</b>	boolean value	Assign IPv6 and IPv4 addresses when creating instances.  Related options: <ul style="list-style-type: none"> <li>• <code>use_neutron</code>: this only works with nova-network.</li> </ul>
<b>use_network_dns_servers = False</b>	boolean value	When this option is set to True, the <code>dns1</code> and <code>dns2</code> servers for the network specified by the user on boot will be used for DNS, as well as any specified in the <b><code>dns_server</code></b> option.  Related options: <ul style="list-style-type: none"> <li>• <b><code>dns_server</code></b></li> </ul>
<b>use_neutron = True</b>	boolean value	Enable neutron as the backend for networking.  Determine whether to use Neutron or Nova Network as the back end. Set to true to use neutron.
<b>use_pub_sub = False</b>	boolean value	Use PUB/SUB pattern for fanout methods. PUB/SUB always uses proxy.
<b>use_rootwrap_daemon = False</b>	boolean value	Start and use a daemon that can run the commands that need to be run with root privileges. This option is usually enabled on nodes that run nova compute processes.
<b>use_router_proxy = False</b>	boolean value	Use ROUTER remote proxy.
<b>use_single_default_gateway = False</b>	boolean value	When set to True, only the first nic of a VM will get its default gateway from the DHCP server.
<b>use_stderr = False</b>	boolean value	Log output to standard error. This option is ignored if <code>log_config_append</code> is set.

Configuration option = Default value	Type	Description
<b>vcpu_pin_set = None</b>	string value	<p>Defines which physical CPUs (pCPUs) can be used by instance virtual CPUs (vCPUs).</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A comma-separated list of physical CPU numbers that virtual CPUs can be allocated to by default. Each element should be either a single CPU number, a range of CPU numbers, or a caret followed by a CPU number to be excluded from a previous range. For example:</li> </ul> <pre>vcpu_pin_set = "4-12,^8,15"</pre>
<b>vif_plugging_is_fatal = True</b>	boolean value	<p>Determine if instance should boot or fail on VIF plugging timeout.</p> <p>Nova sends a port update to Neutron after an instance has been scheduled, providing Neutron with the necessary information to finish setup of the port. Once completed, Neutron notifies Nova that it has finished setting up the port, at which point Nova resumes the boot of the instance since network connectivity is now supposed to be present. A timeout will occur if the reply is not received after a given interval.</p> <p>This option determines what Nova does when the VIF plugging timeout event happens. When enabled, the instance will error out. When disabled, the instance will continue to boot on the assumption that the port is ready.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• True: Instances should fail after VIF plugging timeout</li> <li>• False: Instances should continue booting after VIF plugging timeout</li> </ul>

Configuration option = Default value	Type	Description
<b>vif_plugging_timeout = 300</b>	integer value	<p>Timeout for Neutron VIF plugging event message arrival.</p> <p>Number of seconds to wait for Neutron vif plugging events to arrive before continuing or failing (see <i>vif_plugging_is_fatal</i>).</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <i>vif_plugging_is_fatal</i> - If <b>vif_plugging_timeout</b> is set to zero and <b>vif_plugging_is_fatal</b> is False, events should not be expected to arrive at all.</li> </ul>
<b>virt_mkfs = []</b>	multi valued	<p>Name of the mkfs commands for ephemeral device.</p> <p>The format is &lt;os_type&gt;=&lt;mkfs command&gt;</p>
<b>vlan_interface = None</b>	string value	<p>This option is the name of the virtual interface of the VM on which the VLAN bridge will be built. While it was originally designed to be used only by nova-network, it is also used by libvirt and xenapi for the bridge interface name.</p> <p>Please note that this setting will be ignored in nova-network if the configuration option for <b>network_manager</b> is not set to the default of <i>nova.network.manager.VlanManager</i>.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Any valid virtual interface name, such as <i>eth0</i></li> </ul>

Configuration option = Default value	Type	Description
<b>vlan_start = 100</b>	integer value	<p>This is the VLAN number used for private networks. Note that the when creating the networks, if the specified number has already been assigned, nova-network will increment this number until it finds an available VLAN.</p> <p>Please note that this option is only used when using nova-network instead of Neutron in your deployment. It also will be ignored if the configuration option for <b>network_manager</b> is not set to the default of <i>nova.network.manager.VlanManager</i>.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any integer between 1 and 4094. Values outside of that range will raise a ValueError exception.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li><b>network_manager</b></li> <li><b>use_neutron</b></li> </ul>
<b>volume_usage_poll_interval = 0</b>	integer value	<p>Interval for gathering volume usages.</p> <p>This option updates the volume usage cache for every volume_usage_poll_interval number of seconds.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any positive integer(in seconds) greater than 0 will enable this option.</li> <li>Any value <math>\leq 0</math> will disable the option.</li> </ul>

Configuration option = Default value	Type	Description
<b>vpn_ip = \$my_ip</b>	string value	<p>This option is no longer used since the /os-cloudpipe API was removed in the 16.0.0 Pike release. This is the public IP address for the cloudpipe VPN servers. It defaults to the IP address of the host.</p> <p>Please note that this option is only used when using nova-network instead of Neutron in your deployment. It also will be ignored if the configuration option for <b>network_manager</b> is not set to the default of <i>nova.network.manager.VlanManager</i>.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any valid IP address. The default is <b>\$my_ip</b>, the IP address of the VM.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li><b>network_manager</b></li> <li><b>use_neutron</b></li> <li><b>vpn_start</b></li> </ul>
<b>vpn_start = 1000</b>	port value	<p>This is the port number to use as the first VPN port for private networks.</p> <p>Please note that this option is only used when using nova-network instead of Neutron in your deployment. It also will be ignored if the configuration option for <b>network_manager</b> is not set to the default of <i>nova.network.manager.VlanManager</i>, or if you specify a value the <i>vpn_start</i> parameter when creating a network.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any integer representing a valid port number. The default is 1000.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li><b>use_neutron</b></li> <li><b>vpn_ip</b></li> <li><b>network_manager</b></li> </ul>



Configuration option = Default value	Type	Description
<b>watch-log-file = False</b>	boolean value	Uses logging handler designed to watch file system. When log file is moved or removed this handler will open a new log file with specified path instantaneously. It makes sense only if log_file option is specified and Linux platform is used. This option is ignored if log_config_append is set.
<b>web = /usr/share/spice- html5</b>	string value	Path to directory with content which will be served by a web server.
<b>zmq_failover_connection s = 2</b>	integer value	How many additional connections to a host will be made for failover reasons. This option is actual only in dynamic connections mode.
<b>zmq_immediate = True</b>	boolean value	This option configures round-robin mode in zmq socket. True means not keeping a queue when server side disconnects. False means to keep queue and messages even if server is disconnected, when the server appears we send all accumulated messages to it.
<b>zmq_linger = -1</b>	integer value	Number of seconds to wait before all pending messages will be sent after closing a socket. The default value of -1 specifies an infinite linger period. The value of 0 specifies no linger period. Pending messages shall be discarded immediately when the socket is closed. Positive values specify an upper bound for the linger period.
<b>zmq_target_expire = 300</b>	integer value	Expiration timeout in seconds of a name service record about existing target (< 0 means no timeout).
<b>zmq_target_update = 180</b>	integer value	Update period in seconds of a name service record about existing target.
<b>zmq_tcp_keepalive = -1</b>	integer value	Enable/disable TCP keepalive (KA) mechanism. The default value of -1 (or any other negative value) means to skip any overrides and leave it to OS default; 0 and 1 (or any other positive value) mean to disable and enable the option respectively.
<b>zmq_tcp_keepalive_cnt = -1</b>	integer value	The number of retransmissions to be carried out before declaring that remote end is not available. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.

Configuration option = Default value	Type	Description
<b>zmq_tcp_keepalive_idle = -1</b>	integer value	The duration between two keepalive transmissions in idle condition. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_intvl = -1</b>	integer value	The duration between two successive keepalive retransmissions, if acknowledgement to the previous keepalive transmission is not received. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.

### 11.1.2. api

The following table outlines the options available under the **[api]** group in the `/etc/nova/nova.conf` file.

Table 11.2. api

Configuration option = Default value	Type	Description
<b>allow_instance_snapshots = True</b>	boolean value	Operators can turn off the ability for a user to take snapshots of their instances by setting this option to False. When disabled, any attempt to take a snapshot will result in a HTTP 400 response ("Bad Request").
<b>auth_strategy = keystone</b>	string value	This determines the strategy to use for authentication: keystone or noauth2. <i>noauth2</i> is designed for testing only, as it does no actual credential checking. <i>noauth2</i> provides administrative credentials only if <i>admin</i> is specified as the username.
<b>compute_link_prefix = None</b>	string value	This string is prepended to the normal URL that is returned in links to the OpenStack Compute API. If it is empty (the default), the URLs are returned unchanged.  Possible values: <ul style="list-style-type: none"> <li>Any string, including an empty string (the default).</li> </ul>

Configuration option = Default value	Type	Description
<b>config_drive_skip_versions = 1.0 2007-01-19 2007-03-01 2007-08-29 2007-10-10 2007-12-15 2008-02-01 2008-09-01</b>	string value	<p>When gathering the existing metadata for a config drive, the EC2-style metadata is returned for all versions that don't appear in this option. As of the Liberty release, the available versions are:</p> <ul style="list-style-type: none"> <li>● 1.0</li> <li>● 2007-01-19</li> <li>● 2007-03-01</li> <li>● 2007-08-29</li> <li>● 2007-10-10</li> <li>● 2007-12-15</li> <li>● 2008-02-01</li> <li>● 2008-09-01</li> <li>● 2009-04-04</li> </ul> <p>The option is in the format of a single string, with each version separated by a space.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any string that represents zero or more versions, separated by spaces.</li> </ul>
<b>enable_instance_password = True</b>	boolean value	<p>Enables returning of the instance password by the relevant server API calls such as create, rebuild, evacuate, or rescue. If the hypervisor does not support password injection, then the password returned will not be correct, so if your hypervisor does not support password injection, set this to False.</p>
<b>fping_path = /usr/sbin/fping</b>	string value	<p>The full path to the fping binary.</p>
<b>glance_link_prefix = None</b>	string value	<p>This string is prepended to the normal URL that is returned in links to Glance resources. If it is empty (the default), the URLs are returned unchanged.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any string, including an empty string (the default).</li> </ul>

Configuration option = Default value	Type	Description
<b>hide_server_address_states = [u'building']</b>	list value	<p>This option is a list of all instance states for which network address information should not be returned from the API.</p> <p>Possible values:</p> <p>A list of strings, where each string is a valid VM state, as defined in <code>nova/compute/vm_states.py</code>. As of the Newton release, they are:</p> <ul style="list-style-type: none"> <li>● "active"</li> <li>● "building"</li> <li>● "paused"</li> <li>● "suspended"</li> <li>● "stopped"</li> <li>● "rescued"</li> <li>● "resized"</li> <li>● "soft-delete"</li> <li>● "deleted"</li> <li>● "error"</li> <li>● "shelved"</li> <li>● "shelved_offloaded"</li> </ul>
<b>max_limit = 1000</b>	integer value	As a query can potentially return many thousands of items, you can limit the maximum number of items in a single response by setting this option.
<b>metadata_cache_expiration = 15</b>	integer value	This option is the time (in seconds) to cache metadata. When set to 0, metadata caching is disabled entirely; this is generally not recommended for performance reasons. Increasing this setting should improve response times of the metadata API when under heavy load. Higher values may increase memory usage, and result in longer times for host metadata changes to take effect.

Configuration option = Default value	Type	Description
<b>neutron_default_tenant_id = default</b>	string value	<p>Tenant ID for getting the default network from Neutron API (also referred in some places as the <i>project ID</i>) to use.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● use_neutron_default_nets</li> </ul>
<b>use_forwarded_for = False</b>	boolean value	<p>When True, the <i>X-Forwarded-For</i> header is treated as the canonical remote address. When False (the default), the <i>remote_address</i> header is used.</p> <p>You should only enable this if you have an HTML sanitizing proxy.</p>
<b>use_neutron_default_nets = False</b>	boolean value	<p>When True, the TenantNetworkController will query the Neutron API to get the default networks to use.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● neutron_default_tenant_id</li> </ul>
<b>vendordata_dynamic_connect_timeout = 5</b>	integer value	<p>Maximum wait time for an external REST service to connect.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any integer with a value greater than three (the TCP packet retransmission timeout). Note that instance start may be blocked during this wait time, so this value should be kept small.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● vendordata_providers</li> <li>● vendordata_dynamic_targets</li> <li>● vendordata_dynamic_ssl_certfile</li> <li>● vendordata_dynamic_read_timeout</li> <li>● vendordata_dynamic_failure_fatal</li> </ul>

Configuration option = Default value	Type	Description
<b>vendordata_dynamic_failure_fatal = False</b>	boolean value	<p>Should failures to fetch dynamic vendordata be fatal to instance boot?</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● vendordata_providers</li> <li>● vendordata_dynamic_targets</li> <li>● vendordata_dynamic_ssl_certfile</li> <li>● vendordata_dynamic_connect_timeout</li> <li>● vendordata_dynamic_read_timeout</li> </ul>
<b>vendordata_dynamic_read_timeout = 5</b>	integer value	<p>Maximum wait time for an external REST service to return data once connected.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any integer. Note that instance start is blocked during this wait time, so this value should be kept small.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● vendordata_providers</li> <li>● vendordata_dynamic_targets</li> <li>● vendordata_dynamic_ssl_certfile</li> <li>● vendordata_dynamic_connect_timeout</li> <li>● vendordata_dynamic_failure_fatal</li> </ul>

Configuration option = Default value	Type	Description
<code>`vendordata_dynamic_ssl_certificatefile = `</code>	string value	<p>Path to an optional certificate file or CA bundle to verify dynamic vendordata REST services ssl certificates against.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● An empty string, or a path to a valid certificate file</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>vendordata_providers</code></li> <li>● <code>vendordata_dynamic_targets</code></li> <li>● <code>vendordata_dynamic_connect_timeout</code></li> <li>● <code>vendordata_dynamic_read_timeout</code></li> <li>● <code>vendordata_dynamic_failure_fatal</code></li> </ul>
<code><b>vendordata_dynamic_targets = []</b></code>	list value	<p>A list of targets for the dynamic vendordata provider. These targets are of the form <code>&lt;name&gt;@&lt;url&gt;</code>.</p> <p>The dynamic vendordata provider collects metadata by contacting external REST services and querying them for information about the instance. This behaviour is documented in the <code>vendordata.rst</code> file in the nova developer reference.</p>
<code><b>vendordata_jsonfile_path = None</b></code>	string value	<p>Cloud providers may store custom data in vendor data file that will then be available to the instances via the metadata service, and to the rendering of config-drive. The default class for this, <code>JsonFileVendorData</code>, loads this information from a JSON file, whose path is configured by this option. If there is no path set by this option, the class returns an empty dictionary.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any string representing the path to the data file, or an empty string (default).</li> </ul>

Configuration option = Default value	Type	Description
<b>vendordata_providers = [u'StaticJSON']</b>	list value	<p>A list of vendordata providers.</p> <p>vendordata providers are how deployers can provide metadata via configdrive and metadata that is specific to their deployment. There are currently two supported providers: StaticJSON and DynamicJSON.</p> <p>StaticJSON reads a JSON file configured by the flag <code>vendordata_jsonfile_path</code> and places the JSON from that file into <code>vendor_data.json</code> and <code>vendor_data2.json</code>.</p> <p>DynamicJSON is configured via the <code>vendordata_dynamic_targets</code> flag, which is documented separately. For each of the endpoints specified in that flag, a section is added to the <code>vendor_data2.json</code>.</p> <p>For more information on the requirements for implementing a vendordata dynamic endpoint, please see the <code>vendordata.rst</code> file in the nova developer reference.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A list of vendordata providers, with StaticJSON and DynamicJSON being current options.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>vendordata_dynamic_targets</code></li> <li>● <code>vendordata_dynamic_ssl_certfile</code></li> <li>● <code>vendordata_dynamic_connect_timeout</code></li> <li>● <code>vendordata_dynamic_read_timeout</code></li> <li>● <code>vendordata_dynamic_failure_fatal</code></li> </ul>

### 11.1.3. api\_database

The following table outlines the options available under the **[api\_database]** group in the `/etc/nova/nova.conf` file.

Table 11.3. api\_database



Configuration option = Default value	Type	Description
<b>connection = None</b>	string value	The SQLAlchemy connection string to use to connect to the database.
<b>connection_debug = 0</b>	integer value	Verbosity of SQL debugging information: 0=None, 100=Everything.
<b>connection_recycle_time = 3600</b>	integer value	Connections which have been present in the connection pool longer than this number of seconds will be replaced with a new one the next time they are checked out from the pool.
<b>connection_trace = False</b>	boolean value	Add Python stack traces to SQL as comment strings.
<b>max_overflow = None</b>	integer value	If set, use this value for max_overflow with SQLAlchemy.
<b>max_pool_size = None</b>	integer value	Maximum number of SQL connections to keep open in a pool. Setting a value of 0 indicates no limit.
<b>max_retries = 10</b>	integer value	Maximum number of database connection retries during startup. Set to -1 to specify an infinite retry count.
<b>mysql_sql_mode = TRADITIONAL</b>	string value	The SQL mode to be used for MySQL sessions. This option, including the default, overrides any server-set SQL mode. To use whatever SQL mode is set by the server configuration, set this to no value. Example: mysql_sql_mode=
<b>pool_timeout = None</b>	integer value	If set, use this value for pool_timeout with SQLAlchemy.
<b>retry_interval = 10</b>	integer value	Interval between retries of opening a SQL connection.
<b>slave_connection = None</b>	string value	The SQLAlchemy connection string to use to connect to the slave database.
<b>sqlite_synchronous = True</b>	boolean value	If True, SQLite uses synchronous mode.

#### 11.1.4. barbican

The following table outlines the options available under the **[barbican]** group in the `/etc/nova/nova.conf` file.

Table 11.4. barbican

Configuration option = Default value	Type	Description
<b>auth_endpoint = http://localhost/identity/v3</b>	string value	Use this endpoint to connect to Keystone
<b>barbican_api_version = None</b>	string value	Version of the Barbican API, for example: "v1"
<b>barbican_endpoint = None</b>	string value	Use this endpoint to connect to Barbican, for example: "http://localhost:9311/"
<b>number_of_retries = 60</b>	integer value	Number of times to retry poll for key creation completion
<b>retry_delay = 1</b>	integer value	Number of seconds to wait before retrying poll for key creation completion
<b>verify_ssl = True</b>	boolean value	Specifies if insecure TLS (https) requests. If False, the server's certificate will not be validated

### 11.1.5. cache

The following table outlines the options available under the **[cache]** group in the `/etc/nova/nova.conf` file.

Table 11.5. cache

Configuration option = Default value	Type	Description
<b>backend = dogpile.cache.null</b>	string value	Cache backend module. For eventlet-based or environments with hundreds of threaded servers, Memcache with pooling ( <code>oslo_cache.memcache_pool</code> ) is recommended. For environments with less than 100 threaded servers, Memcached ( <code>dogpile.cache.memcached</code> ) or Redis ( <code>dogpile.cache.redis</code> ) is recommended. Test environments with a single instance of the server can use the <code>dogpile.cache.memory</code> backend.
<b>backend_argument = []</b>	multi valued	Arguments supplied to the backend module. Specify this option once per argument to be passed to the <code>dogpile.cache</code> backend. Example format: " <code>&lt;argname&gt;: &lt;value&gt;</code> ".
<b>config_prefix = cache.oslo</b>	string value	Prefix for building the configuration dictionary for the cache region. This should not need to be changed unless there is another <code>dogpile.cache</code> region with the same configuration name.

Configuration option = Default value	Type	Description
<b>debug_cache_backend = False</b>	boolean value	Extra debugging from the cache backend (cache keys, get/set/delete/etc calls). This is only really useful if you need to see the specific cache-backend get/set/delete calls with the keys/values. Typically this should be left set to false.
<b>enabled = False</b>	boolean value	Global toggle for caching.
<b>expiration_time = 600</b>	integer value	Default TTL, in seconds, for any cached item in the dogpile.cache region. This applies to any cached method that doesn't have an explicit cache expiration time defined for it.
<b>memcache_dead_retry = 300</b>	integer value	Number of seconds memcached server is considered dead before it is tried again. (dogpile.cache.memcache and oslo_cache.memcache_pool backends only).
<b>memcache_pool_connection_get_timeout = 10</b>	integer value	Number of seconds that an operation will wait to get a memcache client connection.
<b>memcache_pool_maxsize = 10</b>	integer value	Max total number of open connections to every memcached server. (oslo_cache.memcache_pool backend only).
<b>memcache_pool_unused_timeout = 60</b>	integer value	Number of seconds a connection to memcached is held unused in the pool before it is closed. (oslo_cache.memcache_pool backend only).
<b>memcache_servers = [u'localhost:11211']</b>	list value	Memcache servers in the format of "host:port". (dogpile.cache.memcache and oslo_cache.memcache_pool backends only).
<b>memcache_socket_timeout = 3</b>	integer value	Timeout in seconds for every call to a server. (dogpile.cache.memcache and oslo_cache.memcache_pool backends only).
<b>proxies = []</b>	list value	Proxy classes to import that will affect the way the dogpile.cache backend functions. See the dogpile.cache documentation on changing-backend-behavior.

### 11.1.6. cells

The following table outlines the options available under the **[cells]** group in the **/etc/nova/nova.conf** file.

Table 11.6. cells

Configuration option = Default value	Type	Description
<b>bandwidth_update_interval = 600</b>	integer value	<p>Bandwidth update interval.</p> <p>Seconds between bandwidth usage cache updates for cells.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>An integer, corresponding to the interval time in seconds.</li> </ul>
<b>call_timeout = 60</b>	integer value	<p>Call timeout.</p> <p>Cell messaging module waits for response(s) to be put into the eventlet queue. This option defines the seconds waited for response from a call to a cell.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>An integer, corresponding to the interval time in seconds.</li> </ul>
<b>capabilities = [u'hypervisor=xenserver;kvm', u'os=linux;windows']</b>	list value	<p>Cell capabilities.</p> <p>List of arbitrary key=value pairs defining capabilities of the current cell to be sent to the parent cells. These capabilities are intended to be used in cells scheduler filters/weighers.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>key=value pairs list for example; <b>hypervisor=xenserver;kvm,os=linux;windows</b></li> </ul>
<b>cell_type = compute</b>	string value	<p>Type of cell.</p> <p>When cells feature is enabled the hosts in the OpenStack Compute cloud are partitioned into groups. Cells are configured as a tree. The top-level cell's cell_type must be set to <b>api</b>. All other cells are defined as a <b>compute cell</b> by default.</p> <p>Related option:</p> <ul style="list-style-type: none"> <li>quota_driver: Disable quota checking for the child cells. (nova.quota.NoopQuotaDriver)</li> </ul>
<b>cells_config = None</b>	string value	Optional cells configuration.

Configuration option = Default value	Type	Description
		<p>Configuration file from which to read cells configuration. If given, overrides reading cells from the database.</p> <p>Cells store all inter-cell communication data, including user names and passwords, in the database. Because the cells data is not updated very frequently, use this option to specify a JSON file to store cells data. With this configuration, the database is no longer consulted when reloading the cells data. The file must have columns present in the Cell model (excluding common database fields and the id column). You must specify the queue connection information through a <code>transport_url</code> field, instead of username, password, and so on.</p> <p>The <code>transport_url</code> has the following form:  <code>rabbit://USERNAME:PASSWORD@HOSTNAME:PORT/VIRTUAL_HOST</code></p> <p>Possible values:</p> <p>The scheme can be either <code>qpid</code> or <code>rabbit</code>, the following sample shows this optional configuration:</p> <pre data-bbox="817 972 1450 2033"> {   "parent": {     "name": "parent",     "api_url": "http://api.example.com:8774",     "transport_url": "rabbit://rabbit.example.com",     "weight_offset": 0.0,     "weight_scale": 1.0,     "is_parent": true   },   "cell1": {     "name": "cell1",     "api_url": "http://api.example.com:8774",     "transport_url": "rabbit://rabbit1.example.com",     "weight_offset": 0.0,     "weight_scale": 1.0,     "is_parent": false   },   "cell2": {     "name": "cell2",     "api_url": "http://api.example.com:8774",     "transport_url": "rabbit://rabbit2.example.com",     "weight_offset": 0.0,     "weight_scale": 1.0,     "is_parent": false   } } </pre>

Configuration option = Default value	Type	Description
<b>db_check_interval = 60</b>	integer value	<p>DB check interval.</p> <p>Cell state manager updates cell status for all cells from the DB only after this particular interval time is passed. Otherwise cached status are used. If this value is 0 or negative all cell status are updated from the DB whenever a state is needed.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Interval time, in seconds.</li> </ul>
<b>enable = False</b>	boolean value	<p>Enable cell v1 functionality.</p> <p>Note that cells v1 is considered experimental and not recommended for new Nova deployments. Cells v1 is being replaced by cells v2 which starting in the 15.0.0 Ocata release, all Nova deployments are at least a cells v2 cell of one. Setting this option, or any other options in the [cells] group, is not required for cells v2.</p> <p>When this functionality is enabled, it lets you to scale an OpenStack Compute cloud in a more distributed fashion without having to use complicated technologies like database and message queue clustering. Cells are configured as a tree. The top-level cell should have a host that runs a nova-api service, but no nova-compute services. Each child cell should run all of the typical nova-* services in a regular Compute cloud except for nova-api. You can think of cells as a normal Compute deployment in that each cell has its own database server and message queue broker.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● name: A unique cell name must be given when this functionality is enabled.</li> <li>● cell_type: Cell type should be defined for all cells.</li> </ul>

Configuration option = Default value	Type	Description
<b>instance_update_num_instances = 1</b>	integer value	<p>Instance update num instances</p> <p>On every run of the periodic task, nova cells manager will attempt to sync instance_updated_at_threshold number of instances. When the manager gets the list of instances, it shuffles them so that multiple nova-cells services do not attempt to sync the same instances in lockstep.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Positive integer number</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● This value is used with the <b>instance_updated_at_threshold</b> value in a periodic task run.</li> </ul>
<b>instance_update_sync_database_limit = 100</b>	integer value	<p>Instance update sync database limit.</p> <p>Number of instances to pull from the database at one time for a sync. If there are more instances to update the results will be paged through.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● An integer, corresponding to a number of instances.</li> </ul>
<b>instance_updated_at_threshold = 3600</b>	integer value	<p>Instance updated at threshold</p> <p>Number of seconds after an instance was updated or deleted to continue to update cells. This option lets cells manager to only attempt to sync instances that have been updated recently. i.e., a threshold of 3600 means to only update instances that have modified in the last hour.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Threshold in seconds</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● This value is used with the <b>instance_update_num_instances</b> value in a periodic task run.</li> </ul>

Configuration option = Default value	Type	Description
<b>max_hop_count = 10</b>	integer value	<p>Maximum hop count</p> <p>When processing a targeted message, if the local cell is not the target, a route is defined between neighbouring cells. And the message is processed across the whole routing path. This option defines the maximum hop counts until reaching the target.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Positive integer value</li> </ul>
<b>mute_child_interval = 300</b>	integer value	<p>Mute child interval.</p> <p>Number of seconds after which a lack of capability and capacity update the child cell is to be treated as a mute cell. Then the child cell will be weighed as recommend highly that it be skipped.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• An integer, corresponding to the interval time in seconds.</li> </ul>
<b>mute_weight_multiplier = -10000.0</b>	floating point value	<p>Mute weight multiplier.</p> <p>Multiplier used to weigh mute children. Mute children cells are recommended to be skipped so their weight is multiplied by this negative value.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Negative numeric number</li> </ul>
<b>name = nova</b>	string value	<p>Name of the current cell.</p> <p>This value must be unique for each cell. Name of a cell is used as its id, leaving this option unset or setting the same name for two or more cells may cause unexpected behaviour.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>• enabled: This option is meaningful only when cells service is enabled</li> </ul>



Configuration option = Default value	Type	Description
<b>offset_weight_multiplier = 1.0</b>	floating point value	<p>Offset weight multiplier</p> <p>Multiplier used to weigh offset weigher. Cells with higher weight_offsets in the DB will be preferred. The weight_offset is a property of a cell stored in the database. It can be used by a deployer to have scheduling decisions favor or disfavor cells based on the setting.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Numeric multiplier</li> </ul>
<b>ram_weight_multiplier = 10.0</b>	floating point value	<p>Ram weight multiplier.</p> <p>Multiplier used for weighing ram. Negative numbers indicate that Compute should stack VMs on one host instead of spreading out new VMs to more hosts in the cell.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Numeric multiplier</li> </ul>
<b>reserve_percent = 10.0</b>	floating point value	<p>Reserve percentage</p> <p>Percentage of cell capacity to hold in reserve, so the minimum amount of free resource is considered to be;</p> <p><b>█</b> <math>\text{min\_free} = \text{total} * (\text{reserve\_percent} / 100.0)</math></p> <p>This option affects both memory and disk utilization.</p> <p>The primary purpose of this reserve is to ensure some space is available for users who want to resize their instance to be larger. Note that currently once the capacity expands into this reserve space this option is ignored.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• An integer or float, corresponding to the percentage of cell capacity to be held in reserve.</li> </ul>

Configuration option = Default value	Type	Description
<b>rpc_driver_queue_base = cells.intercell</b>	string value	<p>RPC driver queue base.</p> <p>When sending a message to another cell by JSON-ifying the message and making an RPC cast to <i>process_message</i>, a base queue is used. This option defines the base queue name to be used when communicating between cells. Various topics by message type will be appended to this.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• The base queue name to be used when communicating between cells.</li> </ul>
<b>scheduler = nova.cells.scheduler.CellsScheduler</b>	string value	<p>Cells scheduler.</p> <p>The class of the driver used by the cells scheduler. This should be the full Python path to the class to be used. If nothing is specified in this option, the CellsScheduler is used.</p>
<b>scheduler_filter_classes = [u'nova.cells.filters.all_filters']</b>	list value	<p>Scheduler filter classes.</p> <p>Filter classes the cells scheduler should use. An entry of "nova.cells.filters.all_filters" maps to all cells filters included with nova. As of the Mitaka release the following filter classes are available:</p> <p>Different cell filter: A scheduler hint of <i>different_cell</i> with a value of a full cell name may be specified to route a build away from a particular cell.</p> <p>Image properties filter: Image metadata named <i>hypervisor_version_requires</i> with a version specification may be specified to ensure the build goes to a cell which has hypervisors of the required version. If either the version requirement on the image or the hypervisor capability of the cell is not present, this filter returns without filtering out the cells.</p> <p>Target cell filter: A scheduler hint of <i>target_cell</i> with a value of a full cell name may be specified to route a build to a particular cell. No error handling is done as there's no way to know whether the full path is a valid.</p> <p>As an admin user, you can also add a filter that directs builds to a particular cell.</p>

Configuration option = Default value	Type	Description
<b>scheduler_retries = 10</b>	integer value	<p>Scheduler retries.</p> <p>How many retries when no cells are available. Specifies how many times the scheduler tries to launch a new instance when no cells are available.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Positive integer value</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● This value is used with the <b>scheduler_retry_delay</b> value while retrying to find a suitable cell.</li> </ul>
<b>scheduler_retry_delay = 2</b>	integer value	<p>Scheduler retry delay.</p> <p>Specifies the delay (in seconds) between scheduling retries when no cell can be found to place the new instance on. When the instance could not be scheduled to a cell after <b>scheduler_retries</b> in combination with <b>scheduler_retry_delay</b>, then the scheduling of the instance failed.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Time in seconds.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● This value is used with the <b>scheduler_retries</b> value while retrying to find a suitable cell.</li> </ul>

Configuration option = Default value	Type	Description
<b>scheduler_weight_classes =</b> <b>[u'nova.cells.weights.all_weighers']</b>	list value	<p>Scheduler weight classes.</p> <p>Weigher classes the cells scheduler should use. An entry of "nova.cells.weights.all_weighers" maps to all cell weighers included with nova. As of the Mitaka release the following weight classes are available:</p> <p><b>mute_child:</b> Downgrades the likelihood of child cells being chosen for scheduling requests, which haven't sent capacity or capability updates in a while. Options include <code>mute_weight_multiplier</code> (multiplier for mute children; value should be negative).</p> <p><b>ram_by_instance_type:</b> Select cells with the most RAM capacity for the instance type being requested. Because higher weights win, Compute returns the number of available units for the instance type requested. The <code>ram_weight_multiplier</code> option defaults to 10.0 that adds to the weight by a factor of 10. Use a negative number to stack VMs on one host instead of spreading out new VMs to more hosts in the cell.</p> <p><b>weight_offset:</b> Allows modifying the database to weight a particular cell. The highest weight will be the first cell to be scheduled for launching an instance. When the <code>weight_offset</code> of a cell is set to 0, it is unlikely to be picked but it could be picked if other cells have a lower weight, like if they're full. And when the <code>weight_offset</code> is set to a very high value (for example, 9999999999999999), it is likely to be picked if another cell do not have a higher weight.</p>

### 11.1.7. cinder

The following table outlines the options available under the **[cinder]** group in the `/etc/nova/nova.conf` file.

Table 11.7. cinder

Configuration option = Default value	Type	Description
<b>auth-url = None</b>	string value	Authentication URL
<b>auth_section = None</b>	string value	Config Section from which to load plugin specific options
<b>auth_type = None</b>	string value	Authentication type to load

Configuration option = Default value	Type	Description
<b>cafile = None</b>	string value	PEM encoded Certificate Authority to use when verifying HTTPs connections.
<b>catalog_info = volumev3:cinderv3:public URL</b>	string value	<p>Info to match when looking for cinder in the service catalog.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Format is separated values of the form: &lt;service_type&gt;&lt;service_name&gt;: &lt;endpoint_type&gt;</li> </ul> <p>Note: Nova does not support the Cinder v2 API since the Nova 17.0.0 Queens release.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● endpoint_template - Setting this option will override catalog_info</li> </ul>
<b>certfile = None</b>	string value	PEM encoded client certificate cert file
<b>cross_az_attach = True</b>	boolean value	<p>Allow attach between instance and volume in different availability zones.</p> <p>If False, volumes attached to an instance must be in the same availability zone in Cinder as the instance availability zone in Nova. This also means care should be taken when booting an instance from a volume where source is not "volume" because Nova will attempt to create a volume using the same availability zone as what is assigned to the instance. If that AZ is not in Cinder (or allow_availability_zone_fallback=False in cinder.conf), the volume create request will fail and the instance will fail the build request. By default there is no availability zone restriction on volume attach.</p>
<b>default-domain-id = None</b>	string value	Optional domain ID to use with v3 and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.
<b>default-domain-name = None</b>	string value	Optional domain name to use with v3 API and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.
<b>domain-id = None</b>	string value	Domain ID to scope to

Configuration option = Default value	Type	Description
<b>domain-name = None</b>	string value	Domain name to scope to
<b>endpoint_template = None</b>	string value	<p>If this option is set then it will override service catalog lookup with this template for cinder endpoint</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● URL for cinder endpoint API e.g. <a href="http://localhost:8776/v3/%(project_id)s">http://localhost:8776/v3/%(project_id)s</a></li> </ul> <p>Note: Nova does not support the Cinder v2 API since the Nova 17.0.0 Queens release.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>catalog_info</code> - If <code>endpoint_template</code> is not set, <code>catalog_info</code> will be used.</li> </ul>
<b>http_retries = 3</b>	integer value	<p>Number of times cinderclient should retry on any failed http call. 0 means connection is attempted only once. Setting it to any positive integer means that on failure connection is retried that many times e.g. setting it to 3 means total attempts to connect will be 4.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any integer value. 0 means connection is attempted only once</li> </ul>
<b>insecure = False</b>	boolean value	Verify HTTPS connections.
<b>keyfile = None</b>	string value	PEM encoded client certificate key file
<b>os_region_name = None</b>	string value	<p>Region name of this node. This is used when picking the URL in the service catalog.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any string representing region name</li> </ul>
<b>password = None</b>	string value	User's password
<b>project-domain-id = None</b>	string value	Domain ID containing project
<b>project-domain-name = None</b>	string value	Domain name containing project

Configuration option = Default value	Type	Description
<b>project-id = None</b>	string value	Project ID to scope to
<b>project-name = None</b>	string value	Project name to scope to
<b>system-scope = None</b>	string value	Scope for system operations
<b>tenant-id = None</b>	string value	Tenant ID
<b>tenant-name = None</b>	string value	Tenant Name
<b>timeout = None</b>	integer value	Timeout value for http requests
<b>trust-id = None</b>	string value	Trust ID
<b>user-domain-id = None</b>	string value	User's domain id
<b>user-domain-name = None</b>	string value	User's domain name
<b>user-id = None</b>	string value	User ID
<b>username = None</b>	string value	Username

### 11.1.8. compute

The following table outlines the options available under the **[compute]** group in the `/etc/nova/nova.conf` file.

Table 11.8. compute

Configuration option = Default value	Type	Description
---	------	-------------

Configuration option = Default value	Type	Description
<b>consecutive_build_service_disable_threshold = 10</b>	integer value	<p>Enables reporting of build failures to the scheduler.</p> <p>Any nonzero value will enable sending build failure statistics to the scheduler for use by the BuildFailureWeigher.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any positive integer enables reporting build failures.</li> <li>Zero to disable reporting build failures.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>[filter_scheduler]/build_failure_weight_multiplier</li> </ul>
<b>cpu_shared_set = None</b>	string value	<p>Defines which physical CPUs (pCPUs) will be used for best-effort guest vCPU resources.</p> <p>Currently only used by libvirt driver to place guest emulator threads when hw:emulator_threads_policy:share.</p> <pre>cpu_shared_set = "4-12,^8,15"</pre>
<b>live_migration_wait_for_vif_plug = False</b>	boolean value	<p>Determine if the source compute host should wait for a <b>network-vif-plugged</b> event from the (neutron) networking service before starting the actual transfer of the guest to the destination compute host.</p> <p>If you set this option the same on all of your compute hosts, which you should do if you use the same networking backend universally, you do not have to worry about this.</p> <p>Before starting the transfer of the guest, some setup occurs on the destination compute host, including plugging virtual interfaces. Depending on the networking backend <b>on the destination host</b>, a <b>network-vif-plugged</b> event may be triggered and then received on the source compute host and the source compute can wait for that event to ensure networking is set up on the destination host before starting the guest transfer in the hypervisor.</p> <p>By default, this is False for two reasons:</p> <ol style="list-style-type: none"> <li>Backward compatibility: deployments should test this out and ensure it works for them before enabling it.</li> </ol>



Configuration option = Default value	Type	Description
		<p>2. The compute service cannot reliably determine which types of virtual interfaces (<b>port.binding:vif_type</b>) will send <b>network-vif-plugged</b> events without an accompanying port <b>binding:host_id</b> change. Open vSwitch and linuxbridge should be OK, but OpenDaylight is at least one known backend that will not currently work in this case, see bug <a href="https://launchpad.net/bugs/1755890">https://launchpad.net/bugs/1755890</a> for more details.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• True: wait for <b>network-vif-plugged</b> events before starting guest transfer</li> <li>• False: do not wait for <b>network-vif-plugged</b> events before starting guest transfer (this is how things have always worked before this option was introduced)</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• [DEFAULT]/vif_plugging_is_fatal: if <b>live_migration_wait_for_vif_plug</b> is True and <b>vif_plugging_timeout</b> is greater than 0, and a timeout is reached, the live migration process will fail with an error but the guest transfer will not have started to the destination host</li> <li>• [DEFAULT]/vif_plugging_timeout: if <b>live_migration_wait_for_vif_plug</b> is True, this controls the amount of time to wait before timing out and either failing if <b>vif_plugging_is_fatal</b> is True, or simply continuing with the live migration</li> </ul>
<b>resource_provider_association_refresh = 300</b>	integer value	<p>Interval for updating nova-compute-side cache of the compute node resource provider's aggregates and traits info.</p> <p>This option specifies the number of seconds between attempts to update a provider's aggregates and traits information in the local cache of the compute node.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Any positive integer in seconds.</li> </ul>

### 11.1.9. conductor

The following table outlines the options available under the **[conductor]** group in the **/etc/nova/nova.conf** file.

Table 11.9. conductor

Configuration option = Default value	Type	Description
<b>topic = conductor</b>	string value	Topic exchange name on which conductor nodes listen.
<b>workers = None</b>	integer value	Number of workers for OpenStack Conductor service. The default will be the number of CPUs available.

### 11.1.10. console

The following table outlines the options available under the **[console]** group in the `/etc/nova/nova.conf` file.

Table 11.10. console

Configuration option = Default value	Type	Description
<b>allowed_origins = []</b>	list value	<p>Adds list of allowed origins to the console websocket proxy to allow connections from other origin hostnames. Websocket proxy matches the host header with the origin header to prevent cross-site requests. This list specifies if any there are values other than host are allowed in the origin header.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A list where each element is an allowed origin hostnames, else an empty list</li> </ul>

### 11.1.11. consoleauth

The following table outlines the options available under the **[consoleauth]** group in the `/etc/nova/nova.conf` file.

Table 11.11. consoleauth

Configuration option = Default value	Type	Description
<b>token_ttl = 600</b>	integer value	<p>The lifetime of a console auth token (in seconds).</p> <p>A console auth token is used in authorizing console access for a user. Once the auth token time to live count has elapsed, the token is considered expired. Expired tokens are then deleted.</p>

### 11.1.12. cors

The following table outlines the options available under the **[cors]** group in the `/etc/nova/nova.conf` file.

Table 11.12. cors

Configuration option = Default value	Type	Description
<b>allow_credentials = True</b>	boolean value	Indicate that the actual request can include user credentials
<b>allow_headers = [u'X-Auth-Token', u'X-Openstack-Request-Id', u'X-Identity-Status', u'X-Roles', u'X-Service-Catalog', u'X-User-Id', u'X-Tenant-Id']</b>	list value	Indicate which header field names may be used during the actual request.
<b>allow_methods = [u'GET', u'PUT', u'POST', u'DELETE', u'PATCH']</b>	list value	Indicate which methods can be used during the actual request.
<b>allowed_origin = None</b>	list value	Indicate whether this resource may be shared with the domain received in the requests "origin" header. Format: "<protocol>://<host>[:<port>]", no trailing slash. Example: <a href="https://horizon.example.com">https://horizon.example.com</a>
<b>expose_headers = [u'X-Auth-Token', u'X-Openstack-Request-Id', u'X-Subject-Token', u'X-Service-Token']</b>	list value	Indicate which headers are safe to expose to the API. Defaults to HTTP Simple Headers.
<b>max_age = 3600</b>	integer value	Maximum cache age of CORS preflight requests.

### 11.1.13. crypto

The following table outlines the options available under the **[crypto]** group in the `/etc/nova/nova.conf` file.

Table 11.13. crypto

Configuration option = Default value	Type	Description
---	------	-------------

Configuration option = Default value	Type	Description
<b>ca_file = cacert.pem</b>	string value	<p>Filename of root CA (Certificate Authority). This is a container format and includes root certificates.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any file name containing root CA, cacert.pem is default</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>ca_path</li> </ul>
<b>ca_path = \$state_path/CA</b>	string value	<p>Directory path where root CA is located.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>ca_file</li> </ul>
<b>crl_file = crl.pem</b>	string value	<p>Filename of root Certificate Revocation List (CRL). This is a list of certificates that have been revoked, and therefore, entities presenting those (revoked) certificates should no longer be trusted.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>ca_path</li> </ul>
<b>key_file = private/cakey.pem</b>	string value	<p>Filename of a private key.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>keys_path</li> </ul>
<b>keys_path = \$state_path/keys</b>	string value	<p>Directory path where keys are located.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>key_file</li> </ul>
<b>project_cert_subject = /C=US/ST=California/O=OpenStack/OU=NovaDev/CN=project-ca-%.16s-%s</b>	string value	<p>Subject for certificate for projects, %s for project, timestamp</p>
<b>use_project_ca = False</b>	boolean value	<p>Option to enable/disable use of CA for each project.</p>

Configuration option = Default value	Type	Description
<b>user_cert_subject =</b> <b>/C=US/ST=California/O=OpenStack/OU=NovaDev/CN=%s</b>	string value	Subject for certificate for users, %s for project, user, timestamp

#### 11.1.14. database

The following table outlines the options available under the **[database]** group in the **/etc/nova/nova.conf** file.

Table 11.14. database

Configuration option = Default value	Type	Description
<b>backend = sqlalchemy</b>	string value	The back end to use for the database.
<b>connection = None</b>	string value	The SQLAlchemy connection string to use to connect to the database.
<b>connection_debug = 0</b>	integer value	Verbosity of SQL debugging information: 0=None, 100=Everything.
<b>connection_recycle_time = 3600</b>	integer value	Connections which have been present in the connection pool longer than this number of seconds will be replaced with a new one the next time they are checked out from the pool.
<b>connection_trace = False</b>	boolean value	Add Python stack traces to SQL as comment strings.
<b>db_inc_retry_interval = True</b>	boolean value	If True, increases the interval between retries of a database operation up to db_max_retry_interval.
<b>db_max_retries = 20</b>	integer value	Maximum retries in case of connection error or deadlock error before error is raised. Set to -1 to specify an infinite retry count.
<b>db_max_retry_interval = 10</b>	integer value	If db_inc_retry_interval is set, the maximum seconds between retries of a database operation.
<b>db_retry_interval = 1</b>	integer value	Seconds between retries of a database transaction.
<b>max_overflow = 50</b>	integer value	If set, use this value for max_overflow with SQLAlchemy.

Configuration option = Default value	Type	Description
<b>max_pool_size = 5</b>	integer value	Maximum number of SQL connections to keep open in a pool. Setting a value of 0 indicates no limit.
<b>max_retries = 10</b>	integer value	Maximum number of database connection retries during startup. Set to -1 to specify an infinite retry count.
<b>min_pool_size = 1</b>	integer value	Minimum number of SQL connections to keep open in a pool.
<b>mysql_enable_ndb = False</b>	boolean value	If True, transparently enables support for handling MySQL Cluster (NDB).
<b>mysql_sql_mode = TRADITIONAL</b>	string value	The SQL mode to be used for MySQL sessions. This option, including the default, overrides any server-set SQL mode. To use whatever SQL mode is set by the server configuration, set this to no value. Example: <code>mysql_sql_mode=</code>
<b>pool_timeout = None</b>	integer value	If set, use this value for <code>pool_timeout</code> with SQLAlchemy.
<b>retry_interval = 10</b>	integer value	Interval between retries of opening a SQL connection.
<b>slave_connection = None</b>	string value	The SQLAlchemy connection string to use to connect to the slave database.
<b>sqlite_synchronous = True</b>	boolean value	If True, SQLite uses synchronous mode.
<b>use_db_reconnect = False</b>	boolean value	Enable the experimental use of database reconnect on connection lost.
<b>use_tpool = False</b>	boolean value	Enable the experimental use of thread pooling for all DB API calls

### 11.1.15. devices

The following table outlines the options available under the **[devices]** group in the `/etc/nova/nova.conf` file.

Table 11.15. devices

Configuration option = Default value	Type	Description
<b>enabled_vgpu_types = []</b>	list value	<p>A list of the vGPU types enabled in the compute node.</p> <p>Some pGPUs (e.g. NVIDIA GRID K1) support different vGPU types. User can use this option to specify a list of enabled vGPU types that may be assigned to a guest instance. But please note that Nova only supports a single type in the Queens release. If more than one vGPU type is specified (as a comma-separated list), only the first one will be used. An example is as the following: [devices] enabled_vgpu_types = GRID K100,Intel GVT-g,MxGPU.2,nvidia-11</p>

### 11.1.16. ephemeral\_storage\_encryption

The following table outlines the options available under the **[ephemeral\_storage\_encryption]** group in the **/etc/nova/nova.conf** file.

Table 11.16. ephemeral\_storage\_encryption

Configuration option = Default value	Type	Description
<b>cipher = aes-xts-plain64</b>	string value	<p>Cipher-mode string to be used.</p> <p>The cipher and mode to be used to encrypt ephemeral storage. The set of cipher-mode combinations available depends on kernel support. According to the dm-crypt documentation, the cipher is expected to be in the format: "&lt;cipher&gt;-&lt;chainmode&gt;-&lt;ivmode&gt;".</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any crypto option listed in <b>/proc/crypto</b>.</li> </ul>
<b>enabled = False</b>	boolean value	Enables/disables LVM ephemeral storage encryption.
<b>key_size = 512</b>	integer value	<p>Encryption key length in bits.</p> <p>The bit length of the encryption key to be used to encrypt ephemeral storage. In XTS mode only half of the bits are used for encryption key.</p>

### 11.1.17. filter\_scheduler

The following table outlines the options available under the **[filter\_scheduler]** group in the **/etc/nova/nova.conf** file.

Table 11.17. filter\_scheduler

Configuration option = Default value	Type	Description
<b>aggregate_image_properties_isolation_namespace = None</b>	string value	<p>Image property namespace for use in the host aggregate.</p> <p>Images and hosts can be configured so that certain images can only be scheduled to hosts in a particular aggregate. This is done with metadata values set on the host aggregate that are identified by beginning with the value of this option. If the host is part of an aggregate with such a metadata key, the image in the request spec must have the value of that metadata in its properties in order for the scheduler to consider the host as acceptable.</p> <p>This option is only used by the FilterScheduler and its subclasses; if you use a different scheduler, this option has no effect. Also note that this setting only affects scheduling if the <i>aggregate_image_properties_isolation</i> filter is enabled.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A string, where the string corresponds to an image property namespace</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>aggregate_image_properties_isolation_separator</code></li> </ul>



Configuration option = Default value	Type	Description
<b>aggregate_image_properties_isolation_separator =</b> .	string value	<p>Separator character(s) for image property namespace and name.</p> <p>When using the <code>aggregate_image_properties_isolation</code> filter, the relevant metadata keys are prefixed with the namespace defined in the <code>aggregate_image_properties_isolation_namespace</code> configuration option plus a separator. This option defines the separator to be used.</p> <p>This option is only used by the <code>FilterScheduler</code> and its subclasses; if you use a different scheduler, this option has no effect. Also note that this setting only affects scheduling if the <code>aggregate_image_properties_isolation</code> filter is enabled.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A string, where the string corresponds to an image property namespace separator character</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <code>aggregate_image_properties_isolation_namespace</code></li> </ul>

Configuration option = Default value	Type	Description
<b>available_filters =</b> <b>[u'nova.scheduler.filters.a</b> <b>ll_filters']</b>	multi valued	<p>Filters that the scheduler can use.</p> <p>An unordered list of the filter classes the nova scheduler may apply. Only the filters specified in the <i>enabled_filters</i> option will be used, but any filter appearing in that option must also be included in this list.</p> <p>By default, this is set to all filters that are included with nova.</p> <p>This option is only used by the FilterScheduler and its subclasses; if you use a different scheduler, this option has no effect.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A list of zero or more strings, where each string corresponds to the name of a filter that may be used for selecting a host</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <code>enabled_filters</code></li> </ul>
<b>baremetal_enabled_filters =</b> <b>[u'RetryFilter',</b> <b>u'AvailabilityZoneFilter',</b> <b>u'ComputeFilter',</b> <b>u'ComputeCapabilitiesFilter',</b> <b>u'ImagePropertiesFilter',</b> <b>u'ExactRamFilter',</b> <b>u'ExactDiskFilter',</b> <b>u'ExactCoreFilter']</b>	list value	<p>Filters used for filtering baremetal hosts.</p> <p>Filters are applied in order, so place your most restrictive filters first to make the filtering process more efficient.</p> <p>This option is only used by the FilterScheduler and its subclasses; if you use a different scheduler, this option has no effect.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A list of zero or more strings, where each string corresponds to the name of a filter to be used for selecting a baremetal host</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• If the <i>scheduler_use_baremetal_filters</i> option is False, this option has no effect.</li> </ul>

Configuration option = Default value	Type	Description
<b>build_failure_weight_multiplier = 1000000.0</b>	floating point value	<p>Multiplier used for weighing hosts that have had recent build failures.</p> <p>This option determines how much weight is placed on a compute node with recent build failures. Build failures may indicate a failing, misconfigured, or otherwise ailing compute node, and avoiding it during scheduling may be beneficial. The weight is inversely proportional to the number of recent build failures the compute node has experienced. This value should be set to some high value to offset weight given by other enabled weighers due to available resources. To disable weighing compute hosts by the number of recent failures, set this to zero.</p> <p>This option is only used by the FilterScheduler and its subclasses; if you use a different scheduler, this option has no effect.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● An integer or float value, where the value corresponds to the multiplier ratio for this weigher.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● [compute]/consecutive_build_service_disable_threshold - Must be nonzero for a compute to report data considered by this weigher.</li> </ul>

Configuration option = Default value	Type	Description
<b>cpu_weight_multiplier = 1.0</b>	floating point value	<p>CPU weight multiplier ratio.</p> <p>Multiplier used for weighting free vCPUs. Negative numbers indicate stacking rather than spreading.</p> <p>This option is only used by the FilterScheduler and its subclasses; if you use a different scheduler, this option has no effect. Also note that this setting only affects scheduling if the <i>cpu</i> weigher is enabled.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● An integer or float value, where the value corresponds to the multiplier ratio for this weigher.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>filter_scheduler.weight_classes</b>: This weigher must be added to list of enabled weight classes if the <b>weight_classes</b> setting is set to a non-default value.</li> </ul>
<b>disk_weight_multiplier = 1.0</b>	floating point value	<p>Disk weight multiplier ratio.</p> <p>Multiplier used for weighing free disk space. Negative numbers mean to stack vs spread.</p> <p>This option is only used by the FilterScheduler and its subclasses; if you use a different scheduler, this option has no effect. Also note that this setting only affects scheduling if the <i>disk</i> weigher is enabled.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● An integer or float value, where the value corresponds to the multiplier ratio for this weigher.</li> </ul>

Configuration option = Default value	Type	Description
<b>enabled_filters =</b> <b>[u'RetryFilter',</b> <b>u'AvailabilityZoneFilter',</b> <b>u'ComputeFilter',</b> <b>u'ComputeCapabilitiesFilter',</b> <b>u'ImagePropertiesFilter',</b> <b>u'ServerGroupAntiAffinityFilter',</b> <b>u'ServerGroupAffinityFilter']</b>	list value	<p>Filters that the scheduler will use.</p> <p>An ordered list of filter class names that will be used for filtering hosts. These filters will be applied in the order they are listed so place your most restrictive filters first to make the filtering process more efficient.</p> <p>This option is only used by the FilterScheduler and its subclasses; if you use a different scheduler, this option has no effect.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A list of zero or more strings, where each string corresponds to the name of a filter to be used for selecting a host</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• All of the filters in this option <b>must</b> be present in the <i>scheduler_available_filters</i> option, or a SchedulerHostFilterNotFound exception will be raised.</li> </ul>
<b>host_subset_size = 1</b>	integer value	<p>Size of subset of best hosts selected by scheduler.</p> <p>New instances will be scheduled on a host chosen randomly from a subset of the N best hosts, where N is the value set by this option.</p> <p>Setting this to a value greater than 1 will reduce the chance that multiple scheduler processes handling similar requests will select the same host, creating a potential race condition. By selecting a host randomly from the N hosts that best fit the request, the chance of a conflict is reduced. However, the higher you set this value, the less optimal the chosen host may be for a given request.</p> <p>This option is only used by the FilterScheduler and its subclasses; if you use a different scheduler, this option has no effect.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• An integer, where the integer corresponds to the size of a host subset. Any integer is valid, although any value less than 1 will be treated as 1</li> </ul>

Configuration option = Default value	Type	Description
<b>image_properties_default_architecture = None</b>	string value	<p>The default architecture to be used when using the image properties filter.</p> <p>When using the ImagePropertiesFilter, it is possible that you want to define a default architecture to make the user experience easier and avoid having something like x86_64 images landing on aarch64 compute nodes because the user did not specify the <i>hw_architecture</i> property in Glance.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● CPU Architectures such as x86_64, aarch64, s390x.</li> </ul>
<b>io_ops_weight_multiplier = -1.0</b>	floating point value	<p>IO operations weight multiplier ratio.</p> <p>This option determines how hosts with differing workloads are weighed. Negative values, such as the default, will result in the scheduler preferring hosts with lighter workloads whereas positive values will prefer hosts with heavier workloads. Another way to look at it is that positive values for this option will tend to schedule instances onto hosts that are already busy, while negative values will tend to distribute the workload across more hosts. The absolute value, whether positive or negative, controls how strong the <i>io_ops</i> weigher is relative to other weighers.</p> <p>This option is only used by the FilterScheduler and its subclasses; if you use a different scheduler, this option has no effect. Also note that this setting only affects scheduling if the <i>io_ops</i> weigher is enabled.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● An integer or float value, where the value corresponds to the multiplier ratio for this weigher.</li> </ul>

Configuration option = Default value	Type	Description
<b>isolated_hosts = []</b>	list value	<p>List of hosts that can only run certain images.</p> <p>If there is a need to restrict some images to only run on certain designated hosts, list those host names here.</p> <p>This option is only used by the <i>FilterScheduler</i> and its subclasses; if you use a different scheduler, this option has no effect. Also note that this setting only affects scheduling if the <i>IsolatedHostsFilter</i> filter is enabled.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A list of strings, where each string corresponds to the name of a host</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● scheduler/isolated_images</li> <li>● scheduler/restrict_isolated_hosts_to_isolated_images</li> </ul>
<b>isolated_images = []</b>	list value	<p>List of UUIDs for images that can only be run on certain hosts.</p> <p>If there is a need to restrict some images to only run on certain designated hosts, list those image UUIDs here.</p> <p>This option is only used by the <i>FilterScheduler</i> and its subclasses; if you use a different scheduler, this option has no effect. Also note that this setting only affects scheduling if the <i>IsolatedHostsFilter</i> filter is enabled.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A list of UUID strings, where each string corresponds to the UUID of an image</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● scheduler/isolated_hosts</li> <li>● scheduler/restrict_isolated_hosts_to_isolated_images</li> </ul>

Configuration option = Default value	Type	Description
<b>max_instances_per_host = 50</b>	integer value	<p>Maximum number of instances that be active on a host.</p> <p>If you need to limit the number of instances on any given host, set this option to the maximum number of instances you want to allow. The <code>num_instances_filter</code> will reject any host that has at least as many instances as this option's value.</p> <p>This option is only used by the FilterScheduler and its subclasses; if you use a different scheduler, this option has no effect. Also note that this setting only affects scheduling if the <code>num_instances_filter</code> filter is enabled.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● An integer, where the integer corresponds to the max instances that can be scheduled on a host.</li> </ul>
<b>max_io_ops_per_host = 8</b>	integer value	<p>The number of instances that can be actively performing IO on a host.</p> <p>Instances performing IO includes those in the following states: build, resize, snapshot, migrate, rescue, unshelve.</p> <p>This option is only used by the FilterScheduler and its subclasses; if you use a different scheduler, this option has no effect. Also note that this setting only affects scheduling if the <code>io_ops_filter</code> filter is enabled.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● An integer, where the integer corresponds to the max number of instances that can be actively performing IO on any given host.</li> </ul>



Configuration option = Default value	Type	Description
<p><b>pci_weight_multiplier = 1.0</b></p>	<p>floating point value</p>	<p>PCI device affinity weight multiplier.</p> <p>The PCI device affinity weighter computes a weighting based on the number of PCI devices on the host and the number of PCI devices requested by the instance. The <b>NUMATopologyFilter</b> filter must be enabled for this to have any significance. For more information, refer to the filter documentation:</p> <p><a href="https://docs.openstack.org/nova/latest/user/filter-scheduler.html">https://docs.openstack.org/nova/latest/user/filter-scheduler.html</a></p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A positive integer or float value, where the value corresponds to the multiplier ratio for this weigher.</li> </ul>
<p><b>ram_weight_multiplier = 1.0</b></p>	<p>floating point value</p>	<p>RAM weight multiplier ratio.</p> <p>This option determines how hosts with more or less available RAM are weighed. A positive value will result in the scheduler preferring hosts with more available RAM, and a negative number will result in the scheduler preferring hosts with less available RAM. Another way to look at it is that positive values for this option will tend to spread instances across many hosts, while negative values will tend to fill up (stack) hosts as much as possible before scheduling to a less-used host. The absolute value, whether positive or negative, controls how strong the RAM weigher is relative to other weighers.</p> <p>This option is only used by the FilterScheduler and its subclasses; if you use a different scheduler, this option has no effect. Also note that this setting only affects scheduling if the <i>ram</i> weigher is enabled.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• An integer or float value, where the value corresponds to the multiplier ratio for this weigher.</li> </ul>

Configuration option = Default value	Type	Description
<b>restrict_isolated_hosts_to_isolated_images = True</b>	boolean value	<p>Prevent non-isolated images from being built on isolated hosts.</p> <p>This option is only used by the <code>FilterScheduler</code> and its subclasses; if you use a different scheduler, this option has no effect. Also note that this setting only affects scheduling if the <code>IsolatedHostsFilter</code> filter is enabled. Even then, this option doesn't affect the behavior of requests for isolated images, which will <b>always</b> be restricted to isolated hosts.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <code>scheduler/isolated_images</code></li> <li>• <code>scheduler/isolated_hosts</code></li> </ul>
<b>shuffle_best_same_weighted_hosts = False</b>	boolean value	<p>Enable spreading the instances between hosts with the same best weight.</p> <p>Enabling it is beneficial for cases when <code>host_subset_size</code> is 1 (default), but there is a large number of hosts with same maximal weight. This scenario is common in Ironic deployments where there are typically many baremetal nodes with identical weights returned to the scheduler. In such case enabling this option will reduce contention and chances for rescheduling events. At the same time it will make the instance packing (even in unweighed case) less dense.</p>
<b>soft_affinity_weight_multiplier = 1.0</b>	floating point value	<p>Multiplier used for weighing hosts for group soft-affinity.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• An integer or float value, where the value corresponds to weight multiplier for hosts with group soft affinity. Only a positive value are meaningful, as negative values would make this behave as a soft anti-affinity weigher.</li> </ul>


Configuration option = Default value	Type	Description
<b>soft_anti_affinity_weight_multiplier = 1.0</b>	floating point value	<p>Multiplier used for weighing hosts for group soft-anti-affinity.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● An integer or float value, where the value corresponds to weight multiplier for hosts with group soft anti-affinity. Only a positive value are meaningful, as negative values would make this behave as a soft affinity weigher.</li> </ul>
<b>track_instance_changes = True</b>	boolean value	<p>Enable querying of individual hosts for instance information.</p> <p>The scheduler may need information about the instances on a host in order to evaluate its filters and weighers. The most common need for this information is for the (anti-)affinity filters, which need to choose a host based on the instances already running on a host.</p> <p>If the configured filters and weighers do not need this information, disabling this option will improve performance. It may also be disabled when the tracking overhead proves too heavy, although this will cause classes requiring host usage data to query the database on each request instead.</p> <p>This option is only used by the FilterScheduler and its subclasses; if you use a different scheduler, this option has no effect.</p> <div data-bbox="815 1447 922 1765" style="border: 1px solid black; padding: 5px; width: 60px; height: 140px; margin: 10px 0;">  </div> <p><b>NOTE</b></p> <p>In a multi-cell (v2) setup where the cell MQ is separated from the top-level, computes cannot directly communicate with the scheduler. Thus, this option cannot be enabled in that scenario. See also the [workarounds]/disable_group_policy_check_upcall option.</p>

Configuration option = Default value	Type	Description
<b>use_baremetal_filters = False</b>	boolean value	<p>Enable baremetal filters.</p> <p>Set this to True to tell the nova scheduler that it should use the filters specified in the <i>baremetal_enabled_filters</i> option. If you are not scheduling baremetal nodes, leave this at the default setting of False.</p> <p>This option is only used by the FilterScheduler and its subclasses; if you use a different scheduler, this option has no effect.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>• If this option is set to True, then the filters specified in the <i>baremetal_enabled_filters</i> are used instead of the filters specified in <i>enabled_filters</i>.</li> </ul>
<b>weight_classes = [u'nova.scheduler.weights.all_weighers']</b>	list value	<p>Weighers that the scheduler will use.</p> <p>Only hosts which pass the filters are weighed. The weight for any host starts at 0, and the weighers order these hosts by adding to or subtracting from the weight assigned by the previous weigher. Weights may become negative. An instance will be scheduled to one of the N most-weighted hosts, where N is <i>scheduler_host_subset_size</i>.</p> <p>By default, this is set to all weighers that are included with Nova.</p> <p>This option is only used by the FilterScheduler and its subclasses; if you use a different scheduler, this option has no effect.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A list of zero or more strings, where each string corresponds to the name of a weigher that will be used for selecting a host</li> </ul>

### 11.1.18. glance

The following table outlines the options available under the **[glance]** group in the */etc/nova/nova.conf* file.

Table 11.18. glance

Configuration option = Default value	Type	Description
<b>allowed_direct_url_schemes = []</b>	list value	<p>List of url schemes that can be directly accessed.</p> <p>This option specifies a list of url schemes that can be downloaded directly via the <code>direct_url</code>. This <code>direct_URL</code> can be fetched from Image metadata which can be used by nova to get the image more efficiently. nova-compute could benefit from this by invoking a copy when it has access to the same file system as glance.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• [file], Empty list (default)</li> </ul>
<b>api_servers = None</b>	list value	<p>List of glance api servers endpoints available to nova.</p> <p>https is used for ssl-based glance api servers.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p>The preferred mechanism for endpoint discovery is via keystoneauth1 loading options. Only use <code>api_servers</code> if you need multiple endpoints and are unable to use a load balancer for some reason.</p> </div> </div> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A list of any fully qualified url of the form "scheme://hostname:port[/path]" (i.e. "http://10.0.1.0:9292" or "https://my.glance.server/image").</li> </ul>
<b>cafile = None</b>	string value	PEM encoded Certificate Authority to use when verifying HTTPs connections.
<b>certfile = None</b>	string value	PEM encoded client certificate cert file
<b>debug = False</b>	boolean value	Enable or disable debug logging with glanceclient.

Configuration option = Default value	Type	Description
<b>default_trusted_certificate_ids = []</b>	list value	<p>List of certificate IDs for certificates that should be trusted.</p> <p>May be used as a default list of trusted certificate IDs for certificate validation. The value of this option will be ignored if the user provides a list of trusted certificate IDs with an instance API request. The value of this option will be persisted with the instance data if signature verification and certificate validation are enabled and if the user did not provide an alternative list. If left empty when certificate validation is enabled the user must provide a list of trusted certificate IDs otherwise certificate validation will fail.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>• The value of this option may be used if both <code>verify_glance_signatures</code> and <code>enable_certificate_validation</code> are enabled.</li> </ul>
<b>enable_certificate_validation = False</b>	boolean value	<p>Enable certificate validation for image signature verification.</p> <p>During image signature verification nova will first verify the validity of the image's signing certificate using the set of trusted certificates associated with the instance. If certificate validation fails, signature verification will not be performed and the image will be placed into an error state. This provides end users with stronger assurances that the image data is unmodified and trustworthy. If left disabled, image signature verification can still occur but the end user will not have any assurance that the signing certificate used to generate the image signature is still trustworthy.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>• This option only takes effect if <code>verify_glance_signatures</code> is enabled.</li> <li>• The value of <code>default_trusted_certificate_ids</code> may be used when this option is enabled.</li> </ul>
<b>endpoint-override = None</b>	string value	<p>Always use this endpoint URL for requests for this client. NOTE: The unversioned endpoint should be specified here; to request a particular API version, use the <b>version</b>, <b>min-version</b>, and/or <b>max-version</b> options.</p>

Configuration option = Default value	Type	Description
<b>insecure = False</b>	boolean value	Verify HTTPS connections.
<b>keyfile = None</b>	string value	PEM encoded client certificate key file
<b>num_retries = 0</b>	integer value	Enable glance operation retries.  Specifies the number of retries when uploading / downloading an image to / from glance. 0 means no retries.
<b>region-name = None</b>	string value	The default region_name for endpoint URL discovery.
<b>service-name = None</b>	string value	The default service_name for endpoint URL discovery.
<b>service-type = image</b>	string value	The default service_type for endpoint URL discovery.
<b>timeout = None</b>	integer value	Timeout value for http requests
<b>valid-interfaces = [u'internal', u'public']</b>	list value	List of interfaces, in order of preference, for endpoint URL.
<b>verify_glance_signatures = False</b>	boolean value	Enable image signature verification.  nova uses the image signature metadata from glance and verifies the signature of a signed image while downloading that image. If the image signature cannot be verified or if the image signature metadata is either incomplete or unavailable, then nova will not boot the image and instead will place the instance into an error state. This provides end users with stronger assurances of the integrity of the image data they are using to create servers.  Related options: <ul style="list-style-type: none"> <li>• The options in the <b>key_manager</b> group, as the key_manager is used for the signature validation.</li> <li>• Both enable_certificate_validation and default_trusted_certificate_ids below depend on this option being enabled.</li> </ul>

### 11.1.19. guestfs

The following table outlines the options available under the **[guestfs]** group in the `/etc/nova/nova.conf` file.

Table 11.19. guestfs

Configuration option = Default value	Type	Description
<b>debug = False</b>	boolean value	<p>Enable/disables guestfs logging.</p> <p>This configures guestfs to debug messages and push them to OpenStack logging system. When set to True, it traces libguestfs API calls and enable verbose debug messages. In order to use the above feature, "libguestfs" package must be installed.</p> <p>Related options: Since libguestfs access and modifies VM's managed by libvirt, below options should be set to give access to those VM's.</p> <ul style="list-style-type: none"> <li>● libvirt.inject_key</li> <li>● libvirt.inject_partition</li> <li>● libvirt.inject_password</li> </ul>

### 11.1.20. healthcheck

The following table outlines the options available under the **[healthcheck]** group in the **/etc/nova/nova.conf** file.

Table 11.20. healthcheck

Configuration option = Default value	Type	Description
<b>backends = []</b>	list value	Additional backends that can perform health checks and report that information back as part of a request.
<b>detailed = False</b>	boolean value	Show more detailed information as part of the response
<b>disable_by_file_path = None</b>	string value	Check the presence of a file to determine if an application is running on a port. Used by DisableByFileHealthcheck plugin.
<b>disable_by_file_paths = []</b>	list value	Check the presence of a file based on a port to determine if an application is running on a port. Expects a "port:path" list of strings. Used by DisableByFilesPortsHealthcheck plugin.
<b>path = /healthcheck</b>	string value	The path to respond to healthcheck requests on.

### 11.1.21. hyperv



The following table outlines the options available under the **[hyperv]** group in the **/etc/nova/nova.conf** file.

**Table 11.21. hyperv**

Configuration option = Default value	Type	Description
<b>config_drive_cdrom = False</b>	boolean value	<p>Configuration drive cdrom</p> <p>OpenStack can be configured to write instance metadata to a configuration drive, which is then attached to the instance before it boots. The configuration drive can be attached as a disk drive (default) or as a CD drive.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True: Attach the configuration drive image as a CD drive.</li> <li>● False: Attach the configuration drive image as a disk drive (Default).</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● This option is meaningful with <code>force_config_drive</code> option set to <code>True</code> or when the REST API call to create an instance will have <code>--config-drive=True</code> flag.</li> <li>● <code>config_drive_format</code> option must be set to <code>iso9660</code> in order to use CD drive as the configuration drive image.</li> <li>● To use configuration drive with Hyper-V, you must set the <code>mkisofs_cmd</code> value to the full path to an <code>mkisofs.exe</code> installation. Additionally, you must set the <code>qemu_img_cmd</code> value to the full path to an <code>qemu-img</code> command installation.</li> <li>● You can configure the Compute service to always create a configuration drive by setting the <code>force_config_drive</code> option to <code>True</code>.</li> </ul>

Configuration option = Default value	Type	Description
<b>config_drive_inject_password = False</b>	boolean value	<p>Configuration drive inject password</p> <p>Enables setting the admin password in the configuration drive image.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>• This option is meaningful when used with other options that enable configuration drive usage with Hyper-V, such as <code>force_config_drive</code>.</li> <li>• Currently, the only accepted <code>config_drive_format</code> is <code>iso9660</code>.</li> </ul>
<b>dynamic_memory_ratio = 1.0</b>	floating point value	<p>Dynamic memory ratio</p> <p>Enables dynamic memory allocation (ballooning) when set to a value greater than 1. The value expresses the ratio between the total RAM assigned to an instance and its startup RAM amount. For example a ratio of 2.0 for an instance with 1024MB of RAM implies 512MB of RAM allocated at startup.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• 1.0: Disables dynamic memory allocation (Default).</li> <li>• Float values greater than 1.0: Enables allocation of total implied RAM divided by this value for startup.</li> </ul>
<b>enable_instance_metrics_collection = False</b>	boolean value	<p>Enable instance metrics collection</p> <p>Enables metrics collections for an instance by using Hyper-V's metric APIs. Collected data can be retrieved by other apps and services, e.g.: Ceilometer.</p>

Configuration option = Default value	Type	Description
<b>enable_remotefx = False</b>	boolean value	<p>Enable RemoteFX feature</p> <p>This requires at least one DirectX 11 capable graphics adapter for Windows / Hyper-V Server 2012 R2 or newer and RDS-Virtualization feature has to be enabled.</p> <p>Instances with RemoteFX can be requested with the following flavor extra specs:</p> <p><b>os:resolution. Guest VM screen resolution size.</b>  <b>Acceptable values</b>  1024x768, 1280x1024, 1600x1200, 1920x1200, 2560x1600, 3840x2160</p> <p><b>3840x2160</b> is only available on Windows / Hyper-V Server 2016.</p> <p><b>os:monitors. Guest VM number of monitors.</b>  <b>Acceptable values</b>  [1, 4] - Windows / Hyper-V Server 2012 R2 [1, 8] - Windows / Hyper-V Server 2016</p> <p><b>os:vram. Guest VM VRAM amount. Only available on Windows / Hyper-V Server 2016. Acceptable values::</b></p> <p><b>64, 128, 256, 512, 1024</b></p>
`instances_path_share = `	string value	<p>Instances path share</p> <p>The name of a Windows share mapped to the "instances_path" dir and used by the resize feature to copy files to the target host. If left blank, an administrative share (hidden network share) will be used, looking for the same "instances_path" used locally.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● "": An administrative share will be used (Default).</li> <li>● Name of a Windows share.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● "instances_path": The directory which will be used if this option here is left blank.</li> </ul>

Configuration option = Default value	Type	Description
<b>iscsi_initiator_list = []</b>	list value	<p>List of iSCSI initiators that will be used for establishing iSCSI sessions.</p> <p>If none are specified, the Microsoft iSCSI initiator service will choose the initiator.</p>
<b>limit_cpu_features = False</b>	boolean value	<p>Limit CPU features</p> <p>This flag is needed to support live migration to hosts with different CPU features and checked during instance creation in order to limit the CPU features used by the instance.</p>
<b>mounted_disk_query_retry_count = 10</b>	integer value	<p>Mounted disk query retry count</p> <p>The number of times to retry checking for a mounted disk. The query runs until the device can be found or the retry count is reached.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Positive integer values. Values greater than 1 is recommended (Default: 10).</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>Time interval between disk mount retries is declared with "mounted_disk_query_retry_interval" option.</li> </ul>
<b>mounted_disk_query_retry_interval = 5</b>	integer value	<p>Mounted disk query retry interval</p> <p>Interval between checks for a mounted disk, in seconds.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Time in seconds (Default: 5).</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>This option is meaningful when the mounted_disk_query_retry_count is greater than 1.</li> <li>The retry loop runs with mounted_disk_query_retry_count and mounted_disk_query_retry_interval configuration options.</li> </ul>

Configuration option = Default value	Type	Description
<b>power_state_check_timeframe = 60</b>	integer value	<p>Power state check timeframe</p> <p>The timeframe to be checked for instance power state changes. This option is used to fetch the state of the instance from Hyper-V through the WMI interface, within the specified timeframe.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Timeframe in seconds (Default: 60).</li> </ul>
<b>power_state_event_polling_interval = 2</b>	integer value	<p>Power state event polling interval</p> <p>Instance power state change event polling frequency. Sets the listener interval for power state events to the given value. This option enhances the internal lifecycle notifications of instances that reboot themselves. It is unlikely that an operator has to change this value.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Time in seconds (Default: 2).</li> </ul>

Configuration option = Default value	Type	Description
<b>qemu_img_cmd = qemu-img.exe</b>	string value	<p>qemu-img command</p> <p>qemu-img is required for some of the image related operations like converting between different image types. You can get it from here: (<a href="http://qemu.weilnetz.de/">http://qemu.weilnetz.de/</a>) or you can install the Cloudbase OpenStack Hyper-V Compute Driver (<a href="https://cloudbase.it/openstack-hyperv-driver/">https://cloudbase.it/openstack-hyperv-driver/</a>) which automatically sets the proper path for this config option. You can either give the full path of qemu-img.exe or set its path in the PATH environment variable and leave this option to the default value.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Name of the qemu-img executable, in case it is in the same directory as the nova-compute service or its path is in the PATH environment variable (Default).</li> <li>● Path of qemu-img command (DRIVELETTER:\PATH\TO\QEMU-IMG\COMMAND).</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● If the config_drive_cdrom option is False, qemu-img will be used to convert the ISO to a VHD, otherwise the configuration drive will remain an ISO. To use configuration drive with Hyper-V, you must set the mkisofs_cmd value to the full path to an mkisofs.exe installation.</li> </ul>
<b>use_multipath_io = False</b>	boolean value	<p>Use multipath connections when attaching iSCSI or FC disks.</p> <p>This requires the Multipath IO Windows feature to be enabled. MPIO must be configured to claim such devices.</p>

Configuration option = Default value	Type	Description
<b>volume_attach_retry_count = 10</b>	integer value	<p>Volume attach retry count</p> <p>The number of times to retry attaching a volume. Volume attachment is retried until success or the given retry count is reached.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Positive integer values (Default: 10).</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● Time interval between attachment attempts is declared with <code>volume_attach_retry_interval</code> option.</li> </ul>
<b>volume_attach_retry_interval = 5</b>	integer value	<p>Volume attach retry interval</p> <p>Interval between volume attachment attempts, in seconds.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Time in seconds (Default: 5).</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● This options is meaningful when <code>volume_attach_retry_count</code> is greater than 1.</li> <li>● The retry loop runs with <code>volume_attach_retry_count</code> and <code>volume_attach_retry_interval</code> configuration options.</li> </ul>

Configuration option = Default value	Type	Description
<b>vswitch_name = None</b>	string value	<p>External virtual switch name</p> <p>The Hyper-V Virtual Switch is a software-based layer-2 Ethernet network switch that is available with the installation of the Hyper-V server role. The switch includes programmatically managed and extensible capabilities to connect virtual machines to both virtual networks and the physical network. In addition, Hyper-V Virtual Switch provides policy enforcement for security, isolation, and service levels. The vSwitch represented by this config option must be an external one (not internal or private).</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● If not provided, the first of a list of available vswitches is used. This list is queried using WQL.</li> <li>● Virtual switch name.</li> </ul>
<b>wait_soft_reboot_seconds = 60</b>	integer value	<p>Wait soft reboot seconds</p> <p>Number of seconds to wait for instance to shut down after soft reboot request is made. We fall back to hard reboot if instance does not shutdown within this window.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Time in seconds (Default: 60).</li> </ul>

### 11.1.22. ironic

The following table outlines the options available under the **[ironic]** group in the `/etc/nova/nova.conf` file.

Table 11.22. ironic

Configuration option = Default value	Type	Description
<b>api_endpoint = None</b>	uri value	URL override for the Ironic API endpoint.



Configuration option = Default value	Type	Description
<b>api_max_retries = 60</b>	integer value	The number of times to retry when a request conflicts. If set to 0, only try once, no retries.  Related options: <ul style="list-style-type: none"><li>• api_retry_interval</li></ul>
<b>api_retry_interval = 2</b>	integer value	The number of seconds to wait before retrying the request.  Related options: <ul style="list-style-type: none"><li>• api_max_retries</li></ul>
<b>auth-url = None</b>	string value	Authentication URL
<b>auth_section = None</b>	string value	Config Section from which to load plugin specific options
<b>auth_type = None</b>	string value	Authentication type to load
<b>cafile = None</b>	string value	PEM encoded Certificate Authority to use when verifying HTTPs connections.
<b>certfile = None</b>	string value	PEM encoded client certificate cert file
<b>domain-id = None</b>	string value	Domain ID to scope to
<b>domain-name = None</b>	string value	Domain name to scope to
<b>endpoint-override = None</b>	string value	Always use this endpoint URL for requests for this client. NOTE: The unversioned endpoint should be specified here; to request a particular API version, use the <b>version</b> , <b>min-version</b> , and/or <b>max-version</b> options.
<b>insecure = False</b>	boolean value	Verify HTTPS connections.
<b>keyfile = None</b>	string value	PEM encoded client certificate key file
<b>password = None</b>	string value	User's password
<b>project-domain-id = None</b>	string value	Domain ID containing project
<b>project-domain-name = None</b>	string value	Domain name containing project

Configuration option = Default value	Type	Description
<b>project-id = None</b>	string value	Project ID to scope to
<b>project-name = None</b>	string value	Project name to scope to
<b>region-name = None</b>	string value	The default region_name for endpoint URL discovery.
<b>serial_console_state_timeout = 10</b>	integer value	Timeout (seconds) to wait for node serial console state changed. Set to 0 to disable timeout.
<b>service-name = None</b>	string value	The default service_name for endpoint URL discovery.
<b>service-type = baremetal</b>	string value	The default service_type for endpoint URL discovery.
<b>system-scope = None</b>	string value	Scope for system operations
<b>timeout = None</b>	integer value	Timeout value for http requests
<b>trust-id = None</b>	string value	Trust ID
<b>user-domain-id = None</b>	string value	User's domain id
<b>user-domain-name = None</b>	string value	User's domain name
<b>user-id = None</b>	string value	User ID
<b>username = None</b>	string value	Username
<b>valid-interfaces = [u'internal', u'public']</b>	list value	List of interfaces, in order of preference, for endpoint URL.

### 11.1.23. key\_manager

The following table outlines the options available under the **[key\_manager]** group in the `/etc/nova/nova.conf` file.

Table 11.23. key\_manager

Configuration option = Default value	Type	Description
---	------	-------------

Configuration option = Default value	Type	Description
<b>auth_type = None</b>	string value	The type of authentication credential to create. Possible values are <i>token</i> , <i>password</i> , <i>keystone_token</i> , and <i>keystone_password</i> . Required if no context is passed to the credential factory.
<b>auth_url = None</b>	string value	Use this endpoint to connect to Keystone.
<b>backend = barbican</b>	string value	Specify the key manager implementation. Options are "barbican" and "vault". Default is "barbican". Will support the values earlier set using [key_manager]/api_class for some time.
<b>domain_id = None</b>	string value	Domain ID for domain scoping. Optional for <i>keystone_token</i> and <i>keystone_password</i> auth_type.
<b>domain_name = None</b>	string value	Domain name for domain scoping. Optional for <i>keystone_token</i> and <i>keystone_password</i> auth_type.
<b>fixed_key = None</b>	string value	Fixed key returned by key manager, specified in hex.  Possible values: <ul style="list-style-type: none"> <li>• Empty string or a key in hex value</li> </ul>
<b>password = None</b>	string value	Password for authentication. Required for <i>password</i> and <i>keystone_password</i> auth_type.
<b>project_domain_id = None</b>	string value	Project's domain ID for project. Optional for <i>keystone_token</i> and <i>keystone_password</i> auth_type.
<b>project_domain_name = None</b>	string value	Project's domain name for project. Optional for <i>keystone_token</i> and <i>keystone_password</i> auth_type.
<b>project_id = None</b>	string value	Project ID for project scoping. Optional for <i>keystone_token</i> and <i>keystone_password</i> auth_type.
<b>project_name = None</b>	string value	Project name for project scoping. Optional for <i>keystone_token</i> and <i>keystone_password</i> auth_type.
<b>reauthenticate = True</b>	boolean value	Allow fetching a new token if the current one is going to expire. Optional for <i>keystone_token</i> and <i>keystone_password</i> auth_type.
<b>token = None</b>	string value	Token for authentication. Required for <i>token</i> and <i>keystone_token</i> auth_type if no context is passed to the credential factory.

Configuration option = Default value	Type	Description
<b>trust_id = None</b>	string value	Trust ID for trust scoping. Optional for <i>keystone_token</i> and <i>keystone_password</i> auth_type.
<b>user_domain_id = None</b>	string value	User's domain ID for authentication. Optional for <i>keystone_token</i> and <i>keystone_password</i> auth_type.
<b>user_domain_name = None</b>	string value	User's domain name for authentication. Optional for <i>keystone_token</i> and <i>keystone_password</i> auth_type.
<b>user_id = None</b>	string value	User ID for authentication. Optional for <i>keystone_token</i> and <i>keystone_password</i> auth_type.
<b>username = None</b>	string value	Username for authentication. Required for <i>password</i> auth_type. Optional for the <i>keystone_password</i> auth_type.

#### 11.1.24. keystone

The following table outlines the options available under the **[keystone]** group in the `/etc/nova/nova.conf` file.

Table 11.24. keystone

Configuration option = Default value	Type	Description
<b>cafile = None</b>	string value	PEM encoded Certificate Authority to use when verifying HTTPs connections.
<b>certfile = None</b>	string value	PEM encoded client certificate cert file
<b>endpoint-override = None</b>	string value	Always use this endpoint URL for requests for this client. NOTE: The unversioned endpoint should be specified here; to request a particular API version, use the <b>version</b> , <b>min-version</b> , and/or <b>max-version</b> options.
<b>insecure = False</b>	boolean value	Verify HTTPS connections.
<b>keyfile = None</b>	string value	PEM encoded client certificate key file
<b>region-name = None</b>	string value	The default region_name for endpoint URL discovery.
<b>service-name = None</b>	string value	The default service_name for endpoint URL discovery.

Configuration option = Default value	Type	Description
<b>service-type = identity</b>	string value	The default service_type for endpoint URL discovery.
<b>timeout = None</b>	integer value	Timeout value for http requests
<b>valid-interfaces = [u'internal', u'public']</b>	list value	List of interfaces, in order of preference, for endpoint URL.

### 11.1.25. keystone\_auth token

The following table outlines the options available under the **[keystone\_auth token]** group in the **/etc/nova/nova.conf** file.

Table 11.25. keystone\_auth token

Configuration option = Default value	Type	Description
<b>auth_section = None</b>	string value	Config Section from which to load plugin specific options
<b>auth_type = None</b>	string value	Authentication type to load
<b>auth_uri = None</b>	string value	Complete "public" Identity API endpoint. This endpoint should not be an "admin" endpoint, as it should be accessible by all end users. Unauthenticated clients are redirected to this endpoint to authenticate. Although this endpoint should ideally be unversioned, client support in the wild varies. If you're using a versioned v2 endpoint here, then this should <b>not</b> be the same endpoint the service user utilizes for validating tokens, because normal end users may not be able to reach that endpoint. This option is deprecated in favor of <code>www_authenticate_uri</code> and will be removed in the S release.
<b>auth_version = None</b>	string value	API version of the admin Identity API endpoint.
<b>cache = None</b>	string value	Request environment key where the Swift cache object is stored. When <code>auth_token</code> middleware is deployed with a Swift cache, use this option to have the middleware share a caching backend with swift. Otherwise, use the <b>memcached_servers</b> option instead.

Configuration option = Default value	Type	Description
<b>cafile = None</b>	string value	A PEM encoded Certificate Authority to use when verifying HTTPs connections. Defaults to system CAs.
<b>certfile = None</b>	string value	Required if identity server requires client certificate
<b>check_revocations_for_cached = False</b>	boolean value	If true, the revocation list will be checked for cached tokens. This requires that PKI tokens are configured on the identity server.
<b>delay_auth_decision = False</b>	boolean value	Do not handle authorization requests within the middleware, but delegate the authorization decision to downstream WSGI components.
<b>enforce_token_bind = permissive</b>	string value	Used to control the use and type of token binding. Can be set to: "disabled" to not check token binding. "permissive" (default) to validate binding information if the bind type is of a form known to the server and ignore it if not. "strict" like "permissive" but if the bind type is unknown the token will be rejected. "required" any form of token binding is needed to be allowed. Finally the name of a binding method that must be present in tokens.
<b>hash_algorithms = [u'md5']</b>	list value	Hash algorithms to use for hashing PKI tokens. This may be a single algorithm or multiple. The algorithms are those supported by Python standard hashlib.new(). The hashes will be tried in the order given, so put the preferred one first for performance. The result of the first hash will be stored in the cache. This will typically be set to multiple values only while migrating from a less secure algorithm to a more secure one. Once all the old tokens are expired this option should be set to a single value for better performance.
<b>http_connect_timeout = None</b>	integer value	Request timeout value for communicating with Identity API server.
<b>http_request_max_retries = 3</b>	integer value	How many times are we trying to reconnect when communicating with Identity API Server.
<b>include_service_catalog = True</b>	boolean value	(Optional) Indicate whether to set the X-Service-Catalog header. If False, middleware will not ask for service catalog on token validation and will not set the X-Service-Catalog header.
<b>insecure = False</b>	boolean value	Verify HTTPS connections.

Configuration option = Default value	Type	Description
<b>keyfile = None</b>	string value	Required if identity server requires client certificate
<b>memcache_pool_conn_get_timeout = 10</b>	integer value	(Optional) Number of seconds that an operation will wait to get a memcached client connection from the pool.
<b>memcache_pool_dead_retry = 300</b>	integer value	(Optional) Number of seconds memcached server is considered dead before it is tried again.
<b>memcache_pool_maxsize = 10</b>	integer value	(Optional) Maximum total number of open connections to every memcached server.
<b>memcache_pool_socket_timeout = 3</b>	integer value	(Optional) Socket timeout in seconds for communicating with a memcached server.
<b>memcache_pool_unused_timeout = 60</b>	integer value	(Optional) Number of seconds a connection to memcached is held unused in the pool before it is closed.
<b>memcache_secret_key = None</b>	string value	(Optional, mandatory if memcache_security_strategy is defined) This string is used for key derivation.
<b>memcache_security_strategy = None</b>	string value	(Optional) If defined, indicate whether token data should be authenticated or authenticated and encrypted. If MAC, token data is authenticated (with HMAC) in the cache. If ENCRYPT, token data is encrypted and authenticated in the cache. If the value is not one of these options or empty, auth_token will raise an exception on initialization.
<b>memcache_use_advanced_pool = False</b>	boolean value	(Optional) Use the advanced (eventlet safe) memcached client pool. The advanced pool will only work under python 2.x.
<b>memcached_servers = None</b>	list value	Optionally specify a list of memcached server(s) to use for caching. If left undefined, tokens will instead be cached in-process.
<b>region_name = None</b>	string value	The region in which the identity server can be found.

Configuration option = Default value	Type	Description
<b>revocation_cache_time = 10</b>	integer value	Determines the frequency at which the list of revoked tokens is retrieved from the Identity service (in seconds). A high number of revocation events combined with a low cache duration may significantly reduce performance. Only valid for PKI tokens. This option has been deprecated in the Ocata release and will be removed in the P release.
<b>service_token_roles = [u'service']</b>	list value	A choice of roles that must be present in a service token. Service tokens are allowed to request that an expired token can be used and so this check should tightly control that only actual services should be sending this token. Roles here are applied as an ANY check so any role in this list must be present. For backwards compatibility reasons this currently only affects the allow_expired check.
<b>service_token_roles_required = False</b>	boolean value	For backwards compatibility reasons we must let valid service tokens pass that don't pass the service_token_roles check as valid. Setting this true will become the default in a future release and should be enabled if possible.
<b>signing_dir = None</b>	string value	Directory used to cache files related to PKI tokens. This option has been deprecated in the Ocata release and will be removed in the P release.
<b>token_cache_time = 300</b>	integer value	In order to prevent excessive effort spent validating tokens, the middleware caches previously-seen tokens for a configurable duration (in seconds). Set to -1 to disable caching completely.
<b>www_authenticate_uri = None</b>	string value	Complete "public" Identity API endpoint. This endpoint should not be an "admin" endpoint, as it should be accessible by all end users. Unauthenticated clients are redirected to this endpoint to authenticate. Although this endpoint should ideally be unversioned, client support in the wild varies. If you're using a versioned v2 endpoint here, then this should <b>not</b> be the same endpoint the service user utilizes for validating tokens, because normal end users may not be able to reach that endpoint.

### 11.1.26. libvirt

The following table outlines the options available under the **[libvirt]** group in the `/etc/nova/nova.conf` file.



Table 11.26. libvirt

Configuration option = Default value	Type	Description
<b>checksum_base_images</b> = <b>False</b>	boolean value	Write a checksum for files in <code>_base</code> to disk
<b>checksum_interval_seconds</b> = <b>3600</b>	integer value	How frequently to checksum base images
<code>`connection_uri = `</code>	string value	<p>Overrides the default libvirt URI of the chosen virtualization type.</p> <p>If set, Nova will use this URI to connect to libvirt.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• An URI like <b>qemu:///system</b> or <b>xen+ssh://oirste/</b> for example. This is only necessary if the URI differs to the commonly known URIs for the chosen virtualization type.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <b>virt_type</b>: Influences what is used as default value here.</li> </ul>

Configuration option = Default value	Type	Description
<b>cpu_mode = None</b>	string value	<p>Is used to set the CPU mode an instance should have.</p> <p>If <code>virt_type="kvm qemu"</code>, it will default to "host-model", otherwise it will default to "none".</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● <b>host-model</b>: Clones the host CPU feature flags</li> <li>● <b>host-passthrough</b>: Use the host CPU model exactly</li> <li>● <b>custom</b>: Use a named CPU model</li> <li>● <b>none</b>: Don't set a specific CPU model. For instances with <b>virt_type</b> as KVM/QEMU, the default CPU model from QEMU will be used, which provides a basic set of CPU features that are compatible with most hosts.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>cpu_model</b>: This should be set ONLY when <b>cpu_mode</b> is set to <b>custom</b>. Otherwise, it would result in an error and the instance launch will fail.</li> </ul>
<b>cpu_model = None</b>	string value	<p>Set the name of the libvirt CPU model the instance should use.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● The named CPU models listed in <b><code>/usr/share/libvirt/cpu_map.xml</code></b></li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>cpu_mode</b>: This should be set to <b>custom</b> ONLY when you want to configure (via <b>cpu_model</b>) a specific named CPU model. Otherwise, it would result in an error and the instance launch will fail.</li> <li>● <b>virt_type</b>: Only the virtualization types <b>kvm</b> and <b>qemu</b> use this.</li> </ul>
<b>cpu_model_extra_flags = []</b>	list value	<p>This allows specifying granular CPU feature flags when configuring CPU models. For example, to explicitly specify the <b>pcid</b> (Process-Context ID, an Intel processor feature – which is now required to address the guest performance degradation as a</p>

Configuration option = Default value	Type	Description
		<p>result of applying the "Meltdown" CVE fixes to certain Intel CPU models) flag to the "IvyBridge" virtual CPU model::</p> <pre data-bbox="818 257 1238 398">[libvirt] cpu_mode = custom cpu_model = IvyBridge cpu_model_extra_flags = pcid</pre> <p>To specify multiple CPU flags (e.g. the Intel <b>VMX</b> to expose the virtualization extensions to the guest, or <b>pdpe1gb</b> to configure 1GB huge pages for CPU models that do not provide it):</p> <pre data-bbox="818 629 1337 801">[libvirt] cpu_mode = custom cpu_model = Haswell-noTSX-IBRS cpu_model_extra_flags = PCID, VMX, pdpe1gb</pre> <p>As it can be noticed from above, the <b>cpu_model_extra_flags</b> config attribute is case insensitive. And specifying extra flags is valid in combination with all the three possible values for <b>cpu_mode: custom</b> (this also requires an explicit <b>cpu_model</b> to be specified), <b>host-model</b>, or <b>host-passthrough</b>. A valid example for allowing extra CPU flags even for <b>host-passthrough</b> mode is that sometimes QEMU may disable certain CPU features – e.g. Intel's "invtsc", Invariable Time Stamp Counter, CPU flag. And if you need to expose that CPU flag to the Nova instance, the you need to explicitly ask for it.</p> <p>The possible values for <b>cpu_model_extra_flags</b> depends on the CPU model in use. Refer to <b>/usr/share/libvirt/cpu_map.xml</b> possible CPU feature flags for a given CPU model.</p> <p>Note that when using this config attribute to set the <i>PCID</i> CPU flag with the <b>custom</b> CPU mode, not all virtual (i.e. libvirt / QEMU) CPU models need it:</p> <ul data-bbox="882 1653 1422 1989" style="list-style-type: none"> <li>• The only virtual CPU models that include the <i>PCID</i> capability are Intel "Haswell", "Broadwell", and "Skylake" variants.</li> <li>• The libvirt / QEMU CPU models "Nehalem", "Westmere", "SandyBridge", and "IvyBridge" will <i>not</i> expose the <i>PCID</i> capability by default, even if the host CPUs by the same name include it. I.e. <i>PCID</i> needs to be explicitly specified when using the said virtual CPU models.</li> </ul> <p>For more information about <b>ssbd</b> and related options, please refer to the following security updates:</p>

Configuration option = Default value	Type	Description
		<p><a href="https://www.us-cert.gov/ncas/alerts/TA18-141A">https://www.us-cert.gov/ncas/alerts/TA18-141A</a></p> <p><a href="https://www.redhat.com/archives/libvir-list/2018-May/msg01562.html">https://www.redhat.com/archives/libvir-list/2018-May/msg01562.html</a></p> <p><a href="https://www.redhat.com/archives/libvir-list/2018-June/msg01111.html">https://www.redhat.com/archives/libvir-list/2018-June/msg01111.html</a></p> <p>Besides <b>custom</b>, the libvirt driver has two other CPU modes: The default, <b>host-model</b>, tells it to do the right thing with respect to handling <i>PCID</i> CPU flag for the guest – <b>assuming</b> you are running updated processor microcode, host and guest kernel, libvirt, and QEMU. The other mode, <b>host-passthrough</b>, checks if <i>PCID</i> is available in the hardware, and if so directly passes it through to the Nova guests. Thus, in context of <i>PCID</i>, with either of these CPU modes (<b>host-model</b> or <b>host-passthrough</b>), there is no need to use the <b>cpu_model_extra_flags</b>.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>cpu_mode</code></li> <li>● <code>cpu_model</code></li> </ul>
<b>disk_cachemodes = []</b>	list value	<p>Specific cache modes to use for different disk types.</p> <p>For example: file=directsync,block=none,network=writeback</p> <p>For local or direct-attached storage, it is recommended that you use writethrough (default) mode, as it ensures data integrity and has acceptable I/O performance for applications running in the guest, especially for read operations. However, caching mode none is recommended for remote NFS storage, because direct I/O operations (<code>O_DIRECT</code>) perform better than synchronous I/O operations (with <code>O_SYNC</code>). Caching mode none effectively turns all guest I/O operations into direct I/O operations on the host, which is the NFS client in this environment.</p> <p>Possible cache modes:</p> <ul style="list-style-type: none"> <li>● default: Same as writethrough.</li> <li>● none: With caching mode set to none, the host page cache is disabled, but the disk write cache is enabled for the guest. In this mode, the write performance in the guest is optimal because write operations bypass the host page cache and go directly to the disk write cache. If the disk write cache is battery-backed, or if the applications or storage stack in the guest transfer data properly (either through <code>fsync</code> operations or file system barriers), then data integrity can be ensured. However, because the host page cache is disabled, the read</li> </ul>

Configuration option = Default value	Type	Description
		<p>performance in the guest would not be as good as in the modes where the host page cache is enabled, such as writethrough mode. Shareable disk devices, like for a multi-attachable block storage volume, will have their cache mode set to <i>none</i> regardless of configuration.</p> <ul style="list-style-type: none"> <li>● <b>writethrough:</b> writethrough mode is the default caching mode. With caching set to writethrough mode, the host page cache is enabled, but the disk write cache is disabled for the guest. Consequently, this caching mode ensures data integrity even if the applications and storage stack in the guest do not transfer data to permanent storage properly (either through fsync operations or file system barriers). Because the host page cache is enabled in this mode, the read performance for applications running in the guest is generally better. However, the write performance might be reduced because the disk write cache is disabled.</li> <li>● <b>writeback:</b> With caching set to writeback mode, both the host page cache and the disk write cache are enabled for the guest. Because of this, the I/O performance for applications running in the guest is good, but the data is not protected in a power failure. As a result, this caching mode is recommended only for temporary data where potential data loss is not a concern.</li> <li>● <b>directsync:</b> Like "writethrough", but it bypasses the host page cache.</li> <li>● <b>unsafe:</b> Caching mode of unsafe ignores cache transfer operations completely. As its name implies, this caching mode should be used only for temporary data where data loss is not a concern. This mode can be useful for speeding up guest installations, but you should switch to another caching mode in production environments.</li> </ul>

Configuration option = Default value	Type	Description
<b>disk_prefix = None</b>	string value	<p>Override the default disk prefix for the devices attached to an instance.</p> <p>If set, this is used to identify a free disk device name for a bus.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any prefix which will result in a valid disk device name like <i>sda</i> or <i>hda</i> for example. This is only necessary if the device names differ to the commonly known device name prefixes for a virtualization type such as: <i>sd</i>, <i>xvd</i>, <i>uvd</i>, <i>vd</i>.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li><b>virt_type</b>: Influences which device type is used, which determines the default disk prefix.</li> </ul>
<b>enabled_perf_events = []</b>	list value	<p>This is a performance event list which could be used as monitor. These events will be passed to libvirt domain xml while creating a new instances. Then event statistics data can be collected from libvirt. The minimum libvirt version is 2.0.0. For more information about <b>Performance monitoring events</b>, refer <a href="https://libvirt.org/formatdomain.html#elementsPerf">https://libvirt.org/formatdomain.html#elementsPerf</a>.</p> <p>Possible values: * A string list. For example: <b>enabled_perf_events = cmt, mbml, mbmt</b> The supported events list can be found in <a href="https://libvirt.org/html/libvirt-libvirt-domain.html">https://libvirt.org/html/libvirt-libvirt-domain.html</a>, which you may need to search key words <b>VIR_PERF_PARAM_*</b></p>
<b>gid_maps = []</b>	list value	List of guid targets and ranges.Syntax is guest-gid:host-gid:countMaximum of 5 allowed.
<b>hw_disk_discard = None</b>	string value	<p>Discard option for nova managed disks.</p> <p>Requires:</p> <ul style="list-style-type: none"> <li>Libvirt &gt;= 1.0.6</li> <li>Qemu &gt;= 1.5 (raw format)</li> <li>Qemu &gt;= 1.6 (qcow2 format)</li> </ul>

Configuration option = Default value	Type	Description
<b>hw_machine_type = None</b>	list value	For qemu or KVM guests, set this option to specify a default machine type per host architecture. You can find a list of supported machine types in your environment by checking the output of the "virsh capabilities" command. The format of the value for this config option is host-arch=machine-type. For example: x86_64=machinetype1,armv7l=machinetype2
<b>image_info_filename_pattern = \$instances_path/\$image_cache_subdirectory_name/%(image)s.info</b>	string value	Allows image information files to be stored in non-standard locations
<code>images_rbd_ceph_conf = `</code>	string value	Path to the ceph configuration file to use
<b>images_rbd_pool = rbd</b>	string value	The RADOS pool in which rbd volumes are stored
<b>images_type = default</b>	string value	VM Images format.  If default is specified, then use_cow_images flag is used instead of this one.  Related options: <ul style="list-style-type: none"> <li>● virt.use_cow_images</li> <li>● images_volume_group</li> <li>● [workarounds]/ensure_libvirt_rbd_instance_dir_cleanup</li> </ul>
<b>images_volume_group = None</b>	string value	LVM Volume Group that is used for VM images, when you specify images_type=lv  Related options: <ul style="list-style-type: none"> <li>● images_type</li> </ul>

Configuration option = Default value	Type	Description
<b>inject_key = False</b>	boolean value	<p>Allow the injection of an SSH key at boot time.</p> <p>There is no agent needed within the image to do this. If <b>libguestfs</b> is available on the host, it will be used. Otherwise <b>nbd</b> is used. The file system of the image will be mounted and the SSH key, which is provided in the REST API call will be injected as SSH key for the root user and appended to the <b>authorized_keys</b> of that user. The SELinux context will be set if necessary. Be aware that the injection is <b>not</b> possible when the instance gets launched from a volume.</p> <p>This config option will enable directly modifying the instance disk and does not affect what cloud-init may do using data from config_drive option or the metadata service.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>inject_partition</b>: That option will decide about the discovery and usage of the file system. It also can disable the injection at all.</li> </ul>



Configuration option = Default value	Type	Description
<b>inject_partition = -2</b>	integer value	<p>Determines the way how the file system is chosen to inject data into it.</p> <p><b>libguestfs</b> will be used a first solution to inject data. If that's not available on the host, the image will be locally mounted on the host as a fallback solution. If libguestfs is not able to determine the root partition (because there are more or less than one root partition) or cannot mount the file system it will result in an error and the instance won't be boot.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● -2 ⇒ disable the injection of data.</li> <li>● -1 ⇒ find the root partition with the file system to mount with libguestfs</li> <li>● 0 ⇒ The image is not partitioned</li> <li>● &gt;0 ⇒ The number of the partition to use for the injection</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>inject_key</b>: If this option allows the injection of a SSH key it depends on value greater or equal to -1 for <b>inject_partition</b>.</li> <li>● <b>inject_password</b>: If this option allows the injection of an admin password it depends on value greater or equal to -1 for <b>inject_partition</b>.</li> <li>● <b>guestfs</b> You can enable the debug log level of libguestfs with this config option. A more verbose output will help in debugging issues.</li> <li>● <b>virt_type</b>: If you use <b>lxc</b> as virt_type it will be treated as a single partition image</li> </ul>

Configuration option = Default value	Type	Description
<b>inject_password = False</b>	boolean value	<p>Allow the injection of an admin password for instance only at <b>create</b> and <b>rebuild</b> process.</p> <p>There is no agent needed within the image to do this. If <b>libguestfs</b> is available on the host, it will be used. Otherwise <b>nbd</b> is used. The file system of the image will be mounted and the admin password, which is provided in the REST API call will be injected as password for the root user. If no root user is available, the instance won't be launched and an error is thrown. Be aware that the injection is <b>not</b> possible when the instance gets launched from a volume.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• True: Allows the injection.</li> <li>• False (default): Disallows the injection. Any via the REST API provided admin password will be silently ignored.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <b>inject_partition</b>: That option will decide about the discovery and usage of the file system. It also can disable the injection at all.</li> </ul>
<b>iscsi_iface = None</b>	string value	<p>The iSCSI transport iface to use to connect to target in case offload support is desired.</p> <p>Default format is of the form &lt;transport_name&gt;. &lt;hwaddress&gt; where &lt;transport_name&gt; is one of (be2iscsi, bnx2i, cxgb3i, cxgb4i, qla4xxx, ocs) and &lt;hwaddress&gt; is the MAC address of the interface and can be generated via the <code>iscsiadm -m iface</code> command. Do not confuse the <code>iscsi_iface</code> parameter to be provided here with the actual transport name.</p>
<b>iser_use_multipath = False</b>	boolean value	<p>Use multipath connection of the iSER volume.</p> <p>iSER volumes can be connected as multipath devices. This will provide high availability and fault tolerance.</p>
<b>live_migration_bandwidth = 0</b>	integer value	<p>Maximum bandwidth(in MiB/s) to be used during migration.</p> <p>If set to 0, the hypervisor will choose a suitable default. Some hypervisors do not support this feature and will return an error if bandwidth is not 0. Please refer to the libvirt documentation for further details.</p>

Configuration option = Default value	Type	Description
<b>live_migration_completion_timeout = 800</b>	integer value	<p>Time to wait, in seconds, for migration to successfully complete transferring data before aborting the operation.</p> <p>Value is per GiB of guest RAM + disk to be transferred, with lower bound of a minimum of 2 GiB. Should usually be larger than downtime delay * downtime steps. Set to 0 to disable timeouts.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● live_migration_downtime</li> <li>● live_migration_downtime_steps</li> <li>● live_migration_downtime_delay</li> </ul>
<b>live_migration_downtime = 500</b>	integer value	<p>Maximum permitted downtime, in milliseconds, for live migration switchover.</p> <p>Will be rounded up to a minimum of 100ms. You can increase this value if you want to allow live-migrations to complete faster, or avoid live-migration timeout errors by allowing the guest to be paused for longer during the live-migration switchover.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● live_migration_completion_timeout</li> </ul>
<b>live_migration_downtime_delay = 75</b>	integer value	<p>Time to wait, in seconds, between each step increase of the migration downtime.</p> <p>Minimum delay is 3 seconds. Value is per GiB of guest RAM + disk to be transferred, with lower bound of a minimum of 2 GiB per device.</p>
<b>live_migration_downtime_steps = 10</b>	integer value	<p>Number of incremental steps to reach max downtime value.</p> <p>Will be rounded up to a minimum of 3 steps.</p>

Configuration option = Default value	Type	Description
<b>live_migration_inbound_addr = None</b>	string value	<p>The IP address or hostname to be used as the target for live migration traffic.</p> <p>If this option is set to None, the hostname of the migration target compute node will be used.</p> <p>This option is useful in environments where the live-migration traffic can impact the network plane significantly. A separate network for live-migration traffic can then use this config option and avoids the impact on the management network.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A valid IP address or hostname, else None.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <b>live_migration_tunneled:</b> The live_migration_inbound_addr value is ignored if tunneling is enabled.</li> </ul>
<b>live_migration_permit_auto_converge = False</b>	boolean value	<p>This option allows nova to start live migration with auto converge on.</p> <p>Auto converge throttles down CPU if a progress of on-going live migration is slow. Auto converge will only be used if this flag is set to True and post copy is not permitted or post copy is unavailable due to the version of libvirt and QEMU in use.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>• live_migration_permit_post_copy</li> </ul>

Configuration option = Default value	Type	Description
<b>live_migration_permit_post_copy = False</b>	boolean value	<p>This option allows nova to switch an on-going live migration to post-copy mode, i.e., switch the active VM to the one on the destination node before the migration is complete, therefore ensuring an upper bound on the memory that needs to be transferred. Post-copy requires libvirt&gt;=1.3.3 and QEMU&gt;=2.5.0.</p> <p>When permitted, post-copy mode will be automatically activated if a live-migration memory copy iteration does not make percentage increase of at least 10% over the last iteration.</p> <p>The live-migration force complete API also uses post-copy when permitted. If post-copy mode is not available, force complete falls back to pausing the VM to ensure the live-migration operation will complete.</p> <p>When using post-copy mode, if the source and destination hosts loose network connectivity, the VM being live-migrated will need to be rebooted. For more details, please see the Administration guide.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● live_migration_permit_auto_converge</li> </ul>
<b>live_migration_progress_timeout = 0</b>	integer value	<p>Time to wait, in seconds, for migration to make forward progress in transferring data before aborting the operation.</p> <p>Set to 0 to disable timeouts.</p> <p>This is deprecated, and now disabled by default because we have found serious bugs in this feature that caused false live-migration timeout failures. This feature will be removed or replaced in a future release.</p>

Configuration option = Default value	Type	Description
<b>live_migration_scheme = None</b>	string value	<p>URI scheme used for live migration.</p> <p>Override the default libvirt live migration scheme (which is dependent on <code>virt_type</code>). If this option is set to <code>None</code>, nova will automatically choose a sensible default based on the hypervisor. It is not recommended that you change this unless you are very sure that hypervisor supports a particular scheme.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>virt_type</b>: This option is meaningful only when <b>virt_type</b> is set to <b>kvm</b> or <b>qemu</b>.</li> <li>● <b>live_migration_uri</b>: If <b>live_migration_uri</b> value is not <code>None</code>, the scheme used for live migration is taken from <b>live_migration_uri</b> instead.</li> </ul>
<b>live_migration_tunnelled = False</b>	boolean value	<p>Enable tunnelled migration.</p> <p>This option enables the tunnelled migration feature, where migration data is transported over the libvirtd connection. If enabled, we use the <code>VIR_MIGRATE_TUNNELLED</code> migration flag, avoiding the need to configure the network to allow direct hypervisor to hypervisor communication. If <code>False</code>, use the native transport. If not set, Nova will choose a sensible default based on, for example the availability of native encryption support in the hypervisor. Enabling this option will definitely impact performance massively.</p> <p>Note that this option is NOT compatible with use of block migration.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>live_migration_inbound_addr</b>: The <code>live_migration_inbound_addr</code> value is ignored if tunneling is enabled.</li> </ul>

Configuration option = Default value	Type	Description
<b>live_migration_uri = None</b>	string value	<p>Live migration target URI to use.</p> <p>Override the default libvirt live migration target URI (which is dependent on <code>virt_type</code>). Any included "%s" is replaced with the migration target hostname.</p> <p>If this option is set to None (which is the default), Nova will automatically generate the <b>live_migration_uri</b> value based on only 4 supported <b>virt_type</b> in following list:</p> <ul style="list-style-type: none"> <li>● <i>kvm: qemu+tcp://%s/system</i></li> <li>● <i>qemu: qemu+tcp://%s/system</i></li> <li>● <i>xen: xenmigr://%s/system</i></li> <li>● <i>parallels: parallels+tcp://%s/system</i></li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>live_migration_inbound_addr</b>: If <b>live_migration_inbound_addr</b> value is not None and <b>live_migration_tunnelled</b> is False, the ip/hostname address of target compute node is used instead of <b>live_migration_uri</b> as the uri for live migration.</li> <li>● <b>live_migration_scheme</b>: If <b>live_migration_uri</b> is not set, the scheme used for live migration is taken from <b>live_migration_scheme</b> instead.</li> </ul>
<b>mem_stats_period_seconds = 10</b>	integer value	<p>A number of seconds to memory usage statistics period. Zero or negative value mean to disable memory usage statistics.</p>
<b>nfs_mount_options = None</b>	string value	<p>Mount options passed to the NFS client. See section of the nfs man page for details.</p> <p>Mount options controls the way the filesystem is mounted and how the NFS client behaves when accessing files on this mount point.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any string representing mount options separated by commas.</li> <li>● Example string: <code>vers=3,lookupcache=pos</code></li> </ul>

Configuration option = Default value	Type	Description
<b>nfs_mount_point_base = \$state_path/mnt</b>	string value	<p>Directory where the NFS volume is mounted on the compute node. The default is <i>mnt</i> directory of the location where nova's Python module is installed.</p> <p>NFS provides shared storage for the OpenStack Block Storage service.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A string representing absolute path of mount point.</li> </ul>
<b>num_aoe_discover_tries = 3</b>	integer value	<p>Number of times to rediscover AoE target to find volume.</p> <p>Nova provides support for block storage attaching to hosts via AOE (ATA over Ethernet). This option allows the user to specify the maximum number of retry attempts that can be made to discover the AoE device.</p>
<b>num_iser_scan_tries = 5</b>	integer value	<p>Number of times to scan iSER target to find volume.</p> <p>iSER is a server network protocol that extends iSCSI protocol to use Remote Direct Memory Access (RDMA). This option allows the user to specify the maximum number of scan attempts that can be made to find iSER volume.</p>
<b>num_volume_scan_tries = 5</b>	integer value	Number of times to scan given storage protocol to find volume.
<b>quobyte_client_cfg = None</b>	string value	Path to a Quobyte Client configuration file.
<b>quobyte_mount_point_base = \$state_path/mnt</b>	string value	<p>Directory where the Quobyte volume is mounted on the compute node.</p> <p>Nova supports Quobyte volume driver that enables storing Block Storage service volumes on a Quobyte storage back end. This Option specifies the path of the directory where Quobyte volume is mounted.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A string representing absolute path of mount point.</li> </ul>



Configuration option = Default value	Type	Description
<b>rbd_secret_uuid = None</b>	string value	The libvirt UUID of the secret for the rbd_user volumes.
<b>rbd_user = None</b>	string value	The RADOS client name for accessing rbd(RADOS Block Devices) volumes.  Libvirt will refer to this user when connecting and authenticating with the Ceph RBD server.
<b>realtime_scheduler_priority = 1</b>	integer value	In a realtime host context vCPUs for guest will run in that scheduling priority. Priority depends on the host kernel (usually 1-99)
<b>remote_filesystem_transport = ssh</b>	string value	libvirt's transport method for remote file operations.  Because libvirt cannot use RPC to copy files over network to/from other compute nodes, other method must be used for: <ul style="list-style-type: none"> <li>● creating directory on remote host</li> <li>● creating file on remote host</li> <li>● removing file from remote host</li> <li>● copying file to remote host</li> </ul>
<b>remove_unused_resized_minimum_age_seconds = 3600</b>	integer value	Unused resized base images younger than this will not be removed

Configuration option = Default value	Type	Description
<b>rescue_image_id = None</b>	string value	<p>The ID of the image to boot from to rescue data from a corrupted instance.</p> <p>If the rescue REST API operation doesn't provide an ID of an image to use, the image which is referenced by this ID is used. If this option is not set, the image from the instance is used.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>An ID of an image or nothing. If it points to an <b>Amazon Machine Image (AMI)</b>, consider to set the config options <b>rescue_kernel_id</b> and <b>rescue_ramdisk_id</b> too. If nothing is set, the image of the instance is used.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li><b>rescue_kernel_id</b>: If the chosen rescue image allows the separate definition of its kernel disk, the value of this option is used, if specified. This is the case when <b>Amazon's AMI/AKI/ARI</b> image format is used for the rescue image.</li> <li><b>rescue_ramdisk_id</b>: If the chosen rescue image allows the separate definition of its RAM disk, the value of this option is used if, specified. This is the case when <b>Amazon's AMI/AKI/ARI</b> image format is used for the rescue image.</li> </ul>
<b>rescue_kernel_id = None</b>	string value	<p>The ID of the kernel (AKI) image to use with the rescue image.</p> <p>If the chosen rescue image allows the separate definition of its kernel disk, the value of this option is used, if specified. This is the case when <b>Amazon's AMI/AKI/ARI</b> image format is used for the rescue image.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>An ID of an kernel image or nothing. If nothing is specified, the kernel disk from the instance is used if it was launched with one.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li><b>rescue_image_id</b>: If that option points to an image in <b>Amazon's AMI/AKI/ARI</b> image format, it's useful to use <b>rescue_kernel_id</b> too.</li> </ul>

Configuration option = Default value	Type	Description
<b>rescue_ramdisk_id = None</b>	string value	<p>The ID of the RAM disk (ARI) image to use with the rescue image.</p> <p>If the chosen rescue image allows the separate definition of its RAM disk, the value of this option is used, if specified. This is the case when <b>Amazon's</b> AMI/AKI/ARI image format is used for the rescue image.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● An ID of a RAM disk image or nothing. If nothing is specified, the RAM disk from the instance is used if it was launched with one.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>rescue_image_id</b>: If that option points to an image in <b>Amazon's</b> AMI/AKI/ARI image format, it's useful to use <b>rescue_ramdisk_id</b> too.</li> </ul>
<b>rng_dev_path = /dev/urandom</b>	string value	<p>The path to an RNG (Random Number Generator) device that will be used as the source of entropy on the host. Since libvirt 1.3.4, any path (that returns random numbers when read) is accepted. The recommended source of entropy is <b>/dev/urandom</b> – it is non-blocking, therefore relatively fast; and avoids the limitations of <b>/dev/random</b>, which is a legacy interface. For more details (and comparison between different RNG sources), refer to the "Usage" section in the Linux kernel API documentation for <b>[u]random</b>:  <a href="http://man7.org/linux/man-pages/man4/urandom.4.html">http://man7.org/linux/man-pages/man4/urandom.4.html</a> and  <a href="http://man7.org/linux/man-pages/man7/random.7.html">http://man7.org/linux/man-pages/man7/random.7.html</a>.</p>
<b>rx_queue_size = None</b>	integer value	<p>Configure virtio rx queue size.</p> <p>This option is only usable for virtio-net device with vhost and vhost-user backend. Available only with QEMU/KVM. Requires libvirt v2.3 QEMU v2.7.</p>
<code>`smbfs_mount_options = `</code>	string value	<p>Mount options passed to the SMBFS client.</p> <p>Provide SMBFS options as a single string containing all parameters. See mount.cifs man page for details. Note that the libvirt-qemu <b>uid</b> and <b>gid</b> must be specified.</p>

Configuration option = Default value	Type	Description
<b>smbfs_mount_point_base = \$state_path/mnt</b>	string value	Directory where the SMBFS shares are mounted on the compute node.
<b>snapshot_compression = False</b>	boolean value	<p>Enable snapshot compression for <b>qcow2</b> images.</p> <p>Note: you can set <b>snapshot_image_format</b> to <b>qcow2</b> to force all snapshots to be in <b>qcow2</b> format, independently from their original image type.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>snapshot_image_format</b></li> </ul>
<b>snapshot_image_format = None</b>	string value	<p>Determine the snapshot image format when sending to the image service.</p> <p>If set, this decides what format is used when sending the snapshot to the image service. If not set, defaults to same type as source image.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● <b>raw</b>: RAW disk format</li> <li>● <b>qcow2</b>: KVM default disk format</li> <li>● <b>vmdk</b>: VMWare default disk format</li> <li>● <b>vdi</b>: VirtualBox default disk format</li> <li>● If not set, defaults to same type as source image.</li> </ul>
<b>snapshots_directory = \$instances_path/snapshots</b>	string value	Location where libvirt driver will store snapshots before uploading them to image service
<b>sparse_logical_volumes = False</b>	boolean value	Create sparse logical volumes (with virtalsize) if this flag is set to True.
<b>sysinfo_serial = auto</b>	string value	The data source used to the populate the host "serial" UUID exposed to guest in the virtual BIOS.
<b>tx_queue_size = None</b>	integer value	<p>Configure virtio tx queue size.</p> <p>This option is only usable for virtio-net device with vhost-user backend. Available only with QEMU/KVM. Requires libvirt v3.7 QEMU v2.10.</p>

Configuration option = Default value	Type	Description
<b>uid_maps = []</b>	list value	List of uid targets and ranges. Syntax is guest-uid:host-uid:count. Maximum of 5 allowed.
<b>use_usb_tablet = True</b>	boolean value	<p>Enable a mouse cursor within a graphical VNC or SPICE sessions.</p> <p>This will only be taken into account if the VM is fully virtualized and VNC and/or SPICE is enabled. If the node doesn't support a graphical framebuffer, then it is valid to set this to False.</p> <p>Related options: * <b>[vnc]enabled</b>: If VNC is enabled, <b>use_usb_tablet</b> will have an effect. * <b>[spice]enabled + [spice].agent_enabled</b>: If SPICE is enabled and the spice agent is disabled, the config value of <b>use_usb_tablet</b> will have an effect.</p>
<b>use_virtio_for_bridges = True</b>	boolean value	Use virtio for bridge interfaces with KVM/QEMU
<b>virt_type = kvm</b>	string value	<p>Describes the virtualization type (or so called domain type) libvirt should use.</p> <p>The choice of this type must match the underlying virtualization strategy you have chosen for this host.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● See the predefined set of case-sensitive values.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>connection_uri</b>: depends on this</li> <li>● <b>disk_prefix</b>: depends on this</li> <li>● <b>cpu_mode</b>: depends on this</li> <li>● <b>cpu_model</b>: depends on this</li> </ul>

Configuration option = Default value	Type	Description
<b>volume_clear = zero</b>	string value	<p>Method used to wipe ephemeral disks when they are deleted. Only takes effect if LVM is set as backing storage.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● none - do not wipe deleted volumes</li> <li>● zero - overwrite volumes with zeroes</li> <li>● shred - overwrite volume repeatedly</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● images_type - must be set to <b>lvm</b></li> <li>● volume_clear_size</li> </ul>
<b>volume_clear_size = 0</b>	integer value	<p>Size of area in MiB, counting from the beginning of the allocated volume, that will be cleared using method set in <b>volume_clear</b> option.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● 0 - clear whole volume</li> <li>● &gt;0 - clear specified amount of MiB</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● images_type - must be set to <b>lvm</b></li> <li>● volume_clear - must be set and the value must be different than <b>none</b> for this option to have any impact</li> </ul>
<b>volume_use_multipath = False</b>	boolean value	<p>Use multipath connection of the iSCSI or FC volume</p> <p>Volumes can be connected in the LibVirt as multipath devices. This will provide high availability and fault tolerance.</p>

Configuration option = Default value	Type	Description
<b>vzstorage_cache_path = None</b>	string value	<p>Path to the SSD cache file.</p> <p>You can attach an SSD drive to a client and configure the drive to store a local cache of frequently accessed data. By having a local cache on a client's SSD drive, you can increase the overall cluster performance by up to 10 and more times. <b>WARNING!</b> There is a lot of SSD models which are not server grade and may loose arbitrary set of data changes on power loss. Such SSDs should not be used in Vstorage and are dangerous as may lead to data corruptions and inconsistencies. Please consult with the manual on which SSD models are known to be safe or verify it using <code>vstorage-hwflush-check(1)</code> utility.</p> <p>This option defines the path which should include "%(cluster_name)s" template to separate caches from multiple shares.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>vzstorage_mount_opts</code> may include more detailed cache options.</li> </ul>
<b>vzstorage_log_path = /var/log/vstorage/% (cluster_name)s/nova.log. gz</b>	string value	<p>Path to vzstorage client log.</p> <p>This option defines the log of cluster operations, it should include "%(cluster_name)s" template to separate logs from multiple shares.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>vzstorage_mount_opts</code> may include more detailed logging options.</li> </ul>
<b>vzstorage_mount_group = qemu</b>	string value	<p>Mount owner group name.</p> <p>This option defines the owner group of Vstorage cluster mountpoint.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <code>vzstorage_mount_*</code> group of parameters</li> </ul>

Configuration option = Default value	Type	Description
<b>vzstorage_mount_opts = []</b>	list value	<p>Extra mount options for pstorage-mount</p> <p>For full description of them, see <a href="https://static.openvz.org/vz-man/man1/pstorage-mount.1.gz.html">https://static.openvz.org/vz-man/man1/pstorage-mount.1.gz.html</a> Format is a python string representation of arguments list, like: "[-v, -R, 500]" Shouldn't include -c, -l, -C, -u, -g and -m as those have explicit vzstorage_* options.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● All other vzstorage_* options</li> </ul>
<b>vzstorage_mount_perms = 0770</b>	string value	<p>Mount access mode.</p> <p>This option defines the access bits of Vzstorage cluster mountpoint, in the format similar to one of chmod(1) utility, like this: 0770. It consists of one to four digits ranging from 0 to 7, with missing lead digits assumed to be 0's.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● vzstorage_mount_* group of parameters</li> </ul>
<b>vzstorage_mount_point_base = \$state_path/mnt</b>	string value	<p>Directory where the Virtuozzo Storage clusters are mounted on the compute node.</p> <p>This option defines non-standard mountpoint for Vzstorage cluster.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● vzstorage_mount_* group of parameters</li> </ul>
<b>vzstorage_mount_user = stack</b>	string value	<p>Mount owner user name.</p> <p>This option defines the owner user of Vzstorage cluster mountpoint.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● vzstorage_mount_* group of parameters</li> </ul>
<b>wait_soft_reboot_seconds = 120</b>	integer value	<p>Number of seconds to wait for instance to shut down after soft reboot request is made. We fall back to hard reboot if instance does not shutdown within this window.</p>



Configuration option = Default value	Type	Description
<b>xen_hvmloader_path = /usr/lib/xen/boot/hvmloder</b>	string value	Location where the Xen hvmloader is kept

### 11.1.27. matchmaker\_redis

The following table outlines the options available under the **[matchmaker\_redis]** group in the **/etc/nova/nova.conf** file.

Table 11.27. matchmaker\_redis

Configuration option = Default value	Type	Description
<b>check_timeout = 20000</b>	integer value	Time in ms to wait before the transaction is killed.
<b>host = 127.0.0.1</b>	string value	Host to locate redis.
<code>password = `</code>	string value	Password for Redis server (optional).
<b>port = 6379</b>	port value	Use this port to connect to redis host.
<b>sentinel_group_name = oslo-messaging-zeromq</b>	string value	Redis replica set name.
<b>sentinel_hosts = []</b>	list value	List of Redis Sentinel hosts (fault tolerance mode), e.g., [host:port, host1:port ... ]
<b>socket_timeout = 10000</b>	integer value	Timeout in ms on blocking socket operations.
<b>wait_timeout = 2000</b>	integer value	Time in ms to wait between connection attempts.

### 11.1.28. metrics

The following table outlines the options available under the **[metrics]** group in the **/etc/nova/nova.conf** file.

Table 11.28. metrics

Configuration option = Default value	Type	Description
---	------	-------------

Configuration option = Default value	Type	Description
<b>required = True</b>	boolean value	<p>This setting determines how any unavailable metrics are treated. If this option is set to True, any hosts for which a metric is unavailable will raise an exception, so it is recommended to also use the MetricFilter to filter out those hosts before weighing.</p> <p>This option is only used by the FilterScheduler and its subclasses; if you use a different scheduler, this option has no effect.</p> <p>Possible values:</p> <ul style="list-style-type: none"><li>• True or False, where False ensures any metric being unavailable for a host will set the host weight to <i>weight_of_unavailable</i>.</li></ul> <p>Related options:</p> <ul style="list-style-type: none"><li>• <i>weight_of_unavailable</i></li></ul>

Configuration option = Default value	Type	Description
<b>weight_multiplier = 1.0</b>	floating point value	<p>When using metrics to weight the suitability of a host, you can use this option to change how the calculated weight influences the weight assigned to a host as follows:</p> <ul style="list-style-type: none"> <li>● &gt;1.0: increases the effect of the metric on overall weight</li> <li>● 1.0: no change to the calculated weight</li> <li>● &gt;0.0,&lt;1.0: reduces the effect of the metric on overall weight</li> <li>● 0.0: the metric value is ignored, and the value of the <i>weight_of_unavailable</i> option is returned instead</li> <li>● &gt;-1.0,&lt;0.0: the effect is reduced and reversed</li> <li>● -1.0: the effect is reversed</li> <li>● ←-1.0: the effect is increased proportionally and reversed</li> </ul> <p>This option is only used by the FilterScheduler and its subclasses; if you use a different scheduler, this option has no effect.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● An integer or float value, where the value corresponds to the multiplier ratio for this weigher.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <i>weight_of_unavailable</i></li> </ul>

Configuration option = Default value	Type	Description
<b>weight_of_unavailable = -10000.0</b>	floating point value	<p>When any of the following conditions are met, this value will be used in place of any actual metric value:</p> <ul style="list-style-type: none"> <li>● One of the metrics named in <i>weight_setting</i> is not available for a host, and the value of <i>required</i> is False</li> <li>● The ratio specified for a metric in <i>weight_setting</i> is 0</li> <li>● The <i>weight_multiplier</i> option is set to 0</li> </ul> <p>This option is only used by the FilterScheduler and its subclasses; if you use a different scheduler, this option has no effect.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● An integer or float value, where the value corresponds to the multiplier ratio for this weigher.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <i>weight_setting</i></li> <li>● <i>required</i></li> <li>● <i>weight_multiplier</i></li> </ul>

Configuration option = Default value	Type	Description
<b>weight_setting = []</b>	list value	<p>This setting specifies the metrics to be weighed and the relative ratios for each metric. This should be a single string value, consisting of a series of one or more <i>name=ratio</i> pairs, separated by commas, where <i>name</i> is the name of the metric to be weighed, and <i>ratio</i> is the relative weight for that metric.</p> <p>Note that if the ratio is set to 0, the metric value is ignored, and instead the weight will be set to the value of the <i>weight_of_unavailable</i> option.</p> <p>As an example, let's consider the case where this option is set to:</p> <pre>name1=1.0, name2=-1.3`</pre> <p>The final weight will be:</p> <pre>(name1.value * 1.0) + (name2.value * -1.3)`</pre> <p>This option is only used by the FilterScheduler and its subclasses; if you use a different scheduler, this option has no effect.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>A list of zero or more key/value pairs separated by commas, where the key is a string representing the name of a metric and the value is a numeric weight for that metric. If any value is set to 0, the value is ignored and the weight will be set to the value of the <i>weight_of_unavailable</i> option.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li><code>weight_of_unavailable</code></li> </ul>

### 11.1.29. mks

The following table outlines the options available under the **[mks]** group in the `/etc/nova/nova.conf` file.

Table 11.29. mks

Configuration option = Default value	Type	Description
<b>enabled = False</b>	boolean value	Enables graphical console access for virtual machines.

Configuration option = Default value	Type	Description
<b>mksproxy_base_url =</b> <b>http://127.0.0.1:6090/</b>	uri value	<p>Location of MKS web console proxy</p> <p>The URL in the response points to a WebMKS proxy which starts proxying between client and corresponding vCenter server where instance runs. In order to use the web based console access, WebMKS proxy should be installed and configured</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Must be a valid URL of the form:<b>http://host:port/</b> or <b>https://host:port/</b></li> </ul>

### 11.1.30. neutron

The following table outlines the options available under the **[neutron]** group in the `/etc/nova/nova.conf` file.

Table 11.30. neutron

Configuration option = Default value	Type	Description
<b>auth-url = None</b>	string value	Authentication URL
<b>auth_section = None</b>	string value	Config Section from which to load plugin specific options
<b>auth_type = None</b>	string value	Authentication type to load
<b>cafile = None</b>	string value	PEM encoded Certificate Authority to use when verifying HTTPs connections.
<b>certfile = None</b>	string value	PEM encoded client certificate cert file
<b>default-domain-id = None</b>	string value	Optional domain ID to use with v3 and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.
<b>default-domain-name = None</b>	string value	Optional domain name to use with v3 API and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.

Configuration option = Default value	Type	Description
<b>default_floating_pool = nova</b>	string value	Default name for the floating IP pool.  Specifies the name of floating IP pool used for allocating floating IPs. This option is only used if Neutron does not specify the floating IP pool name in port binding responses.
<b>domain-id = None</b>	string value	Domain ID to scope to
<b>domain-name = None</b>	string value	Domain name to scope to
<b>endpoint-override = None</b>	string value	Always use this endpoint URL for requests for this client. NOTE: The unversioned endpoint should be specified here; to request a particular API version, use the <b>version</b> , <b>min-version</b> , and/or <b>max-version</b> options.
<b>extension_sync_interval = 600</b>	integer value	Integer value representing the number of seconds to wait before querying Neutron for extensions. After this number of seconds the next time Nova needs to create a resource in Neutron it will query Neutron for the extensions that it has loaded. Setting value to 0 will refresh the extensions with no wait.
<b>insecure = False</b>	boolean value	Verify HTTPS connections.
<b>keyfile = None</b>	string value	PEM encoded client certificate key file
<code>metadata_proxy_shared_secret = `</code>	string value	This option holds the shared secret string used to validate proxy requests to Neutron metadata requests. In order to be used, the <i>X-Metadata-Provider-Signature</i> header must be supplied in the request.  Related options: <ul style="list-style-type: none"><li>● <code>service_metadata_proxy</code></li></ul>
<b>ovs_bridge = br-int</b>	string value	Default name for the Open vSwitch integration bridge.  Specifies the name of an integration bridge interface used by OpenvSwitch. This option is only used if Neutron does not specify the OVS bridge name in port binding responses.
<b>password = None</b>	string value	User's password

Configuration option = Default value	Type	Description
<b>physnets = []</b>	list value	<p>List of physnets present on this host.</p> <p>For each <b>physnet</b> listed, an additional section, <b>[neutron_physnet_\$PHYSNET]</b>, will be added to the configuration file. Each section must be configured with a single configuration option, <b>numa_nodes</b>, which should be a list of node IDs for all NUMA nodes this physnet is associated with. For example::</p> <pre>[neutron] physnets = foo, bar</pre> <pre>[neutron_physnet_foo] numa_nodes = 0</pre> <pre>[neutron_physnet_bar] numa_nodes = 0,1</pre> <p>Any <b>physnet</b> that is not listed using this option will be treated as having no particular NUMA node affinity.</p> <p>Tunnelled networks (VXLAN, GRE, ...) cannot be accounted for in this way and are instead configured using the <b>[neutron_tunnel]</b> group. For example::</p> <pre>[neutron_tunnel] numa_nodes = 1</pre> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <b>[neutron_tunnel] numa_nodes</b> can be used to configure NUMA affinity for all tunneled networks</li> <li>• <b>[neutron_physnet_\$PHYSNET] numa_nodes</b> must be configured for each value of <b>\$PHYSNET</b> specified by this option</li> </ul>
<b>project-domain-id = None</b>	string value	Domain ID containing project
<b>project-domain-name = None</b>	string value	Domain name containing project
<b>project-id = None</b>	string value	Project ID to scope to



Configuration option = Default value	Type	Description
<b>project-name = None</b>	string value	Project name to scope to
<b>region-name = None</b>	string value	The default region_name for endpoint URL discovery.
<b>service-name = None</b>	string value	The default service_name for endpoint URL discovery.
<b>service-type = network</b>	string value	The default service_type for endpoint URL discovery.
<b>service_metadata_proxy = False</b>	boolean value	<p>When set to True, this option indicates that Neutron will be used to proxy metadata requests and resolve instance ids. Otherwise, the instance ID must be passed to the metadata request in the <i>X-Instance-ID</i> header.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● metadata_proxy_shared_secret</li> </ul>
<b>system-scope = None</b>	string value	Scope for system operations
<b>tenant-id = None</b>	string value	Tenant ID
<b>tenant-name = None</b>	string value	Tenant Name
<b>timeout = None</b>	integer value	Timeout value for http requests
<b>trust-id = None</b>	string value	Trust ID
<b>url = None</b>	uri value	<p>This option specifies the URL for connecting to Neutron.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any valid URL that points to the Neutron API service is appropriate here. This typically matches the URL returned for the <i>network</i> service type from the Keystone service catalog.</li> </ul>
<b>user-domain-id = None</b>	string value	User's domain id
<b>user-domain-name = None</b>	string value	User's domain name

Configuration option = Default value	Type	Description
<b>user-id = None</b>	string value	User ID
<b>username = None</b>	string value	Username
<b>valid-interfaces = [u'internal', u'public']</b>	list value	List of interfaces, in order of preference, for endpoint URL.

### 11.1.31. notifications

The following table outlines the options available under the **[notifications]** group in the **/etc/nova/nova.conf** file.

Table 11.31. notifications

Configuration option = Default value	Type	Description
<b>bdms_in_notifications = False</b>	boolean value	If enabled, include block device information in the versioned notification payload. Sending block device information is disabled by default as providing that information can incur some overhead on the system since the information may need to be loaded from the database.
<b>default_level = INFO</b>	string value	Default notification level for outgoing notifications.
<b>default_publisher_id = \$host</b>	string value	<p>Default publisher_id for outgoing notifications. If you consider routing notifications using different publisher, change this value accordingly.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Defaults to the current hostname of this host, but it can be any valid oslo.messaging publisher_id</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• host - Hostname, FQDN or IP address of this host.</li> </ul>

Configuration option = Default value	Type	Description
<b>notification_format = both</b>	string value	<p>Specifies which notification format shall be used by nova.</p> <p>The default value is <code>fine</code> for most deployments and rarely needs to be changed. This value can be set to <code>versioned</code> once the infrastructure moves closer to consuming the newer format of notifications. After this occurs, this option will be removed.</p> <p>Note that notifications can be completely disabled by setting <b>driver=noop</b> in the <b>[oslo_messaging_notifications]</b> group.</p> <p>Possible values: * <code>unversioned</code>: Only the legacy unversioned notifications are emitted. * <code>versioned</code>: Only the new versioned notifications are emitted. * <code>both</code>: Both the legacy unversioned and the new versioned notifications are emitted. (Default)</p> <p>The list of versioned notifications is visible in <a href="https://docs.openstack.org/nova/latest/reference/notifications.html">https://docs.openstack.org/nova/latest/reference/notifications.html</a></p>
<b>notify_on_state_change = None</b>	string value	<p>If set, send <code>compute.instance.update</code> notifications on instance state changes.</p> <p>Please refer to <a href="https://docs.openstack.org/nova/latest/reference/notifications.html">https://docs.openstack.org/nova/latest/reference/notifications.html</a> for additional information on notifications.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● <code>None</code> - no notifications</li> <li>● <code>"vm_state"</code> - notifications are sent with VM state transition information in the <b>old_state</b> and <b>state</b> fields. The <b>old_task_state</b> and <b>new_task_state</b> fields will be set to the current <code>task_state</code> of the instance.</li> <li>● <code>"vm_and_task_state"</code> - notifications are sent with VM and task state transition information.</li> </ul>

Configuration option = Default value	Type	Description
<b>versioned_notifications_topics =</b> [u'versioned_notifications']	list value	<p>Specifies the topics for the versioned notifications issued by nova.</p> <p>The default value is fine for most deployments and rarely needs to be changed. However, if you have a third-party service that consumes versioned notifications, it might be worth getting a topic for that service. Nova will send a message containing a versioned notification payload to each topic queue in this list.</p> <p>The list of versioned notifications is visible in <a href="https://docs.openstack.org/nova/latest/reference/notifications.html">https://docs.openstack.org/nova/latest/reference/notifications.html</a></p>

### 11.1.32. osapi\_v21

The following table outlines the options available under the **[osapi\_v21]** group in the **/etc/nova/nova.conf** file.

Table 11.32. osapi\_v21

Configuration option = Default value	Type	Description
<b>project_id_regex = None</b>	string value	<p>This option is a string representing a regular expression (regex) that matches the project_id as contained in URLs. If not set, it will match normal UUIDs created by keystone.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A string representing any legal regular expression</li> </ul>

### 11.1.33. oslo\_concurrency

The following table outlines the options available under the **[oslo\_concurrency]** group in the **/etc/nova/nova.conf** file.

Table 11.33. oslo\_concurrency

Configuration option = Default value	Type	Description
<b>disable_process_locking = False</b>	boolean value	Enables or disables inter-process locks.

Configuration option = Default value	Type	Description
<b>lock_path = None</b>	string value	Directory to use for lock files. For security, the specified directory should only be writable by the user running the processes that need locking. Defaults to environment variable OSLO_LOCK_PATH. If external locks are used, a lock path must be set.

### 11.1.34. oslo\_messaging\_amqp

The following table outlines the options available under the **[oslo\_messaging\_amqp]** group in the **/etc/nova/nova.conf** file.

Table 11.34. oslo\_messaging\_amqp

Configuration option = Default value	Type	Description
<b>addressing_mode = dynamic</b>	string value	Indicates the addressing mode used by the driver. Permitted values: <i>legacy</i> - use legacy non-routable addressing <i>routable</i> - use routable addresses <i>dynamic</i> - use legacy addresses if the message bus does not support routing otherwise use routable addressing
<b>allow_insecure_clients = False</b>	boolean value	Accept clients using either SSL or plain TCP
<b>anycast_address = anycast</b>	string value	Appended to the address prefix when sending to a group of consumers. Used by the message bus to identify messages that should be delivered in a round-robin fashion across consumers.
<b>broadcast_prefix = broadcast</b>	string value	address prefix used when broadcasting to all servers
<b>connection_retry_backoff = 2</b>	integer value	Increase the connection_retry_interval by this many seconds after each unsuccessful failover attempt.
<b>connection_retry_interval = 1</b>	integer value	Seconds to pause before attempting to re-connect.
<b>connection_retry_interval_max = 30</b>	integer value	Maximum limit for connection_retry_interval + connection_retry_backoff
<b>container_name = None</b>	string value	Name for the AMQP container. must be globally unique. Defaults to a generated UUID

Configuration option = Default value	Type	Description
<b>default_notification_exchange = None</b>	string value	Exchange name used in notification addresses. Exchange name resolution precedence: Target.exchange if set else default_notification_exchange if set else control_exchange if set else <i>notify</i>
<b>default_notify_timeout = 30</b>	integer value	The deadline for a sent notification message delivery. Only used when caller does not provide a timeout expiry.
<b>default_reply_retry = 0</b>	integer value	The maximum number of attempts to re-send a reply message which failed due to a recoverable error.
<b>default_reply_timeout = 30</b>	integer value	The deadline for an rpc reply message delivery.
<b>default_rpc_exchange = None</b>	string value	Exchange name used in RPC addresses. Exchange name resolution precedence: Target.exchange if set else default_rpc_exchange if set else control_exchange if set else <i>rpc</i>
<b>default_send_timeout = 30</b>	integer value	The deadline for an rpc cast or call message delivery. Only used when caller does not provide a timeout expiry.
<b>default_sender_link_timeout = 600</b>	integer value	The duration to schedule a purge of idle sender links. Detach link after expiry.
<b>group_request_prefix = unicast</b>	string value	address prefix when sending to any server in group
<b>idle_timeout = 0</b>	integer value	Timeout for inactive connections (in seconds)
<b>link_retry_delay = 10</b>	integer value	Time to pause between re-connecting an AMQP 1.0 link that failed due to a recoverable error.
<b>multicast_address = multicast</b>	string value	Appended to the address prefix when sending a fanout message. Used by the message bus to identify fanout messages.
<b>notify_address_prefix = openstack.org/om/notify</b>	string value	Address prefix for all generated Notification addresses
<b>notify_server_credit = 100</b>	integer value	Window size for incoming Notification messages
<b>`password = `</b>	string value	Password for message broker authentication

Configuration option = Default value	Type	Description
<b>pre_settled = [u'rpc-cast', u'rpc-reply']</b>	multi valued	Send messages of this type pre-settled. Pre-settled messages will not receive acknowledgement from the peer. Note well: pre-settled messages may be silently discarded if the delivery fails. Permitted values: <i>rpc-call</i> - send RPC Calls pre-settled <i>rpc-reply</i> - send RPC Replies pre-settled <i>rpc-cast</i> - Send RPC Casts pre-settled <i>notify</i> - Send Notifications pre-settled
<b>pseudo_vhost = True</b>	boolean value	Enable virtual host support for those message buses that do not natively support virtual hosting (such as <i>qpidd</i> ). When set to true the virtual host name will be added to all message bus addresses, effectively creating a private <i>subnet</i> per virtual host. Set to False if the message bus supports virtual hosting using the <i>hostname</i> field in the AMQP 1.0 Open performative as the name of the virtual host.
<b>reply_link_credit = 200</b>	integer value	Window size for incoming RPC Reply messages.
<b>rpc_address_prefix = openstack.org/om/rpc</b>	string value	Address prefix for all generated RPC addresses
<b>rpc_server_credit = 100</b>	integer value	Window size for incoming RPC Request messages
<code>`sasl_config_dir = `</code>	string value	Path to directory that contains the SASL configuration
<code>`sasl_config_name = `</code>	string value	Name of configuration file (without <i>.conf</i> suffix)
<code>`sasl_default_realm = `</code>	string value	SASL realm to use if no realm present in username
<code>`sasl_mechanisms = `</code>	string value	Space separated list of acceptable SASL mechanisms
<b>server_request_prefix = exclusive</b>	string value	address prefix used when sending to a specific server
<b>ssl = False</b>	boolean value	Attempt to connect via SSL. If no other <i>ssl</i> -related parameters are given, it will use the system's CA-bundle to verify the server's certificate.
<code>`ssl_ca_file = `</code>	string value	CA certificate PEM file used to verify the server's certificate
<code>`ssl_cert_file = `</code>	string value	Self-identifying certificate PEM file for client authentication

Configuration option = Default value	Type	Description
<code>`ssl_key_file = `</code>	string value	Private key PEM file used to sign <code>ssl_cert_file</code> certificate (optional)
<b>ssl_key_password = None</b>	string value	Password for decrypting <code>ssl_key_file</code> (if encrypted)
<b>ssl_verify_vhost = False</b>	boolean value	By default SSL checks that the name in the server's certificate matches the hostname in the <code>transport_url</code> . In some configurations it may be preferable to use the virtual hostname instead, for example if the server uses the Server Name Indication TLS extension (rfc6066) to provide a certificate per virtual host. Set <code>ssl_verify_vhost</code> to True if the server's SSL certificate uses the virtual host name instead of the DNS name.
<b>trace = False</b>	boolean value	Debug: dump AMQP frames to stdout
<b>unicast_address = unicast</b>	string value	Appended to the address prefix when sending to a particular RPC/Notification server. Used by the message bus to identify messages sent to a single destination.
<code>`username = `</code>	string value	User name for message broker authentication

### 11.1.35. oslo\_messaging\_kafka

The following table outlines the options available under the **[oslo\_messaging\_kafka]** group in the `/etc/nova/nova.conf` file.

Table 11.35. oslo\_messaging\_kafka

Configuration option = Default value	Type	Description
<b>conn_pool_min_size = 2</b>	integer value	The pool size limit for connections expiration policy
<b>conn_pool_ttl = 1200</b>	integer value	The time-to-live in sec of idle connections in the pool
<b>consumer_group = oslo_messaging_consumer</b>	string value	Group id for Kafka consumer. Consumers in one group will coordinate message consumption
<b>kafka_consumer_timeout = 1.0</b>	floating point value	Default timeout(s) for Kafka consumers



Configuration option = Default value	Type	Description
<b>kafka_default_host = localhost</b>	string value	Default Kafka broker Host
<b>kafka_default_port = 9092</b>	port value	Default Kafka broker Port
<b>kafka_max_fetch_bytes = 1048576</b>	integer value	Max fetch bytes of Kafka consumer
<b>pool_size = 10</b>	integer value	Pool Size for Kafka Consumers
<b>producer_batch_size = 16384</b>	integer value	Size of batch for the producer async send
<b>producer_batch_timeout = 0.0</b>	floating point value	Upper bound on the delay for KafkaProducer batching in seconds

### 11.1.36. oslo\_messaging\_notifications

The following table outlines the options available under the **[oslo\_messaging\_notifications]** group in the **/etc/nova/nova.conf** file.

Table 11.36. oslo\_messaging\_notifications

Configuration option = Default value	Type	Description
<b>driver = []</b>	multi valued	The Drivers(s) to handle sending notifications. Possible values are messaging, messagingv2, routing, log, test, noop
<b>retry = -1</b>	integer value	The maximum number of attempts to re-send a notification message which failed to be delivered due to a recoverable error. 0 - No retry, -1 - indefinite
<b>topics = [u'notifications']</b>	list value	AMQP topic used for OpenStack notifications.
<b>transport_url = None</b>	string value	A URL representing the messaging driver to use for notifications. If not set, we fall back to the same configuration used for RPC.

### 11.1.37. oslo\_messaging\_rabbit

The following table outlines the options available under the **[oslo\_messaging\_rabbit]** group in the **/etc/nova/nova.conf** file.

Table 11.37. oslo\_messaging\_rabbit

Configuration option = Default value	Type	Description
<b>amqp_auto_delete = False</b>	boolean value	Auto-delete queues in AMQP.
<b>amqp_durable_queues = False</b>	boolean value	Use durable queues in AMQP.
<b>channel_max = None</b>	integer value	Maximum number of channels to allow
<b>connection_factory = single</b>	string value	Connection factory implementation
<b>default_notification_exchange = \${control_exchange}_notification</b>	string value	Exchange name for sending notifications
<b>default_notification_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending notification, -1 means infinite retry.
<b>default_rpc_exchange = \${control_exchange}_rpc</b>	string value	Exchange name for sending RPC messages
<b>default_rpc_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending RPC message, -1 means infinite retry. If actual retry attempts in not 0 the rpc request could be processed more than one time
<b>default_serializer_type = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>fake_rabbit = False</b>	boolean value	Deprecated, use rpc_backend=kombu+memory or rpc_backend=fake
<b>frame_max = None</b>	integer value	The maximum byte size for an AMQP frame
<b>heartbeat_interval = 3</b>	integer value	How often to send heartbeats for consumer's connections
<b>heartbeat_rate = 2</b>	integer value	How often times during the heartbeat_timeout_threshold we check the heartbeat.

Configuration option = Default value	Type	Description
<b>heartbeat_timeout_thresh old = 60</b>	integer value	Number of seconds after which the Rabbit broker is considered down if heartbeat's keep-alive fails (0 disable the heartbeat). EXPERIMENTAL
<b>host_connection_reconn ect_delay = 0.25</b>	floating point value	Set delay for reconnection to some host which has connection error
<b>kombu_compression = None</b>	string value	EXPERIMENTAL: Possible values are: gzip, bz2. If not set compression will not be used. This option may not be available in future versions.
<b>kombu_failover_strategy = round-robin</b>	string value	Determines how the next RabbitMQ node is chosen in case the one we are currently connected to becomes unavailable. Takes effect only if more than one RabbitMQ node is provided in config.
<b>kombu_missing_consum er_retry_timeout = 60</b>	integer value	How long to wait a missing client before abandoning to send it its replies. This value should not be longer than <code>rpc_response_timeout</code> .
<b>kombu_reconnect_delay = 1.0</b>	floating point value	How long to wait before reconnecting in response to an AMQP consumer cancel notification.
<b>notification_listener_pref etch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to notification listener.
<b>notification_persistence = False</b>	boolean value	Persist notification messages.
<b>notification_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending notification message
<b>pool_max_overflow = 0</b>	integer value	Maximum number of connections to create above <b>pool_max_size</b> .
<b>pool_max_size = 30</b>	integer value	Maximum number of connections to keep queued.
<b>pool_recycle = 600</b>	integer value	Lifetime of a connection (since creation) in seconds or None for no recycling. Expired connections are closed on acquire.
<b>pool_stale = 60</b>	integer value	Threshold at which inactive (since release) connections are considered stale in seconds or None for no staleness. Stale connections are closed on acquire.

Configuration option = Default value	Type	Description
<b>pool_timeout = 30</b>	integer value	Default number of seconds to wait for a connections to available
<b>rabbit_ha_queues = False</b>	boolean value	Try to use HA queues in RabbitMQ (x-ha-policy: all). If you change this option, you must wipe the RabbitMQ database. In RabbitMQ 3.0, queue mirroring is no longer controlled by the x-ha-policy argument when declaring a queue. If you just want to make sure that all queues (except those with auto-generated names) are mirrored across all nodes, run: <code>"rabbitmqctl set_policy HA ^(?!amq\.).*{"ha-mode": "all"}"</code>
<b>rabbit_host = localhost</b>	string value	The RabbitMQ broker address where a single node is used.
<b>rabbit_hosts = [u'\$rabbit_host:\$rabbit_port']</b>	list value	RabbitMQ HA cluster host:port pairs.
<b>rabbit_interval_max = 30</b>	integer value	Maximum interval of RabbitMQ connection retries. Default is 30 seconds.
<b>rabbit_login_method = AMQPLAIN</b>	string value	The RabbitMQ login method.
<b>rabbit_max_retries = 0</b>	integer value	Maximum number of RabbitMQ connection retries. Default is 0 (infinite retry count).
<b>rabbit_password = guest</b>	string value	The RabbitMQ password.
<b>rabbit_port = 5672</b>	port value	The RabbitMQ broker port where a single node is used.
<b>rabbit_qos_prefetch_count = 0</b>	integer value	Specifies the number of messages to prefetch. Setting to zero allows unlimited messages.
<b>rabbit_retry_backoff = 2</b>	integer value	How long to backoff for between retries when connecting to RabbitMQ.
<b>rabbit_retry_interval = 1</b>	integer value	How frequently to retry connecting with RabbitMQ.
<b>rabbit_transient_queues_ttl = 1800</b>	integer value	Positive integer representing duration in seconds for queue TTL (x-expires). Queues which are unused for the duration of the TTL are automatically deleted. The parameter affects only reply and fanout queues.

Configuration option = Default value	Type	Description
<b>rabbit_userid = guest</b>	string value	The RabbitMQ userid.
<b>rabbit_virtual_host = /</b>	string value	The RabbitMQ virtual host.
<b>rpc_listener_prefetch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to rpc listener.
<b>rpc_queue_expiration = 60</b>	integer value	Time to live for rpc queues without consumers in seconds.
<b>rpc_reply_exchange = \${control_exchange}_rpc_reply</b>	string value	Exchange name for receiving RPC replies
<b>rpc_reply_listener_prefetch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to rpc reply listener.
<b>rpc_reply_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending reply. -1 means infinite retry during rpc_timeout
<b>rpc_reply_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending reply.
<b>rpc_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending RPC message
<b>socket_timeout = 0.25</b>	floating point value	Set socket timeout in seconds for connection's socket
<b>ssl = None</b>	boolean value	Enable SSL
<code>`ssl_ca_file = `</code>	string value	SSL certification authority file (valid only if SSL enabled).
<code>`ssl_cert_file = `</code>	string value	SSL cert file (valid only if SSL enabled).
<code>`ssl_key_file = `</code>	string value	SSL key file (valid only if SSL enabled).
<b>ssl_options = None</b>	dict value	Arguments passed to ssl.wrap_socket
<code>`ssl_version = `</code>	string value	SSL version to use (valid only if SSL enabled). Valid values are TLSv1 and SSLv23. SSLv2, SSLv3, TLSv1_1, and TLSv1_2 may be available on some distributions.

Configuration option = Default value	Type	Description
<b>tcp_user_timeout = 0.25</b>	floating point value	Set TCP_USER_TIMEOUT in seconds for connection's socket

### 11.1.38. oslo\_messaging\_zmq

The following table outlines the options available under the **[oslo\_messaging\_zmq]** group in the **/etc/nova/nova.conf** file.

Table 11.38. oslo\_messaging\_zmq

Configuration option = Default value	Type	Description
<b>rpc_ack_timeout_base = 15</b>	integer value	Number of seconds to wait for an ack from a cast/call. After each retry attempt this timeout is multiplied by some specified multiplier.
<b>rpc_ack_timeout_multiplier = 2</b>	integer value	Number to multiply base ack timeout by after each retry attempt.
<b>rpc_message_ttl = 300</b>	integer value	Expiration timeout in seconds of a sent/received message after which it is not tracked anymore by a client/server.
<b>rpc_poll_timeout = 1</b>	integer value	The default number of seconds that poll should wait. Poll raises timeout exception when timeout expired.
<b>rpc_retry_attempts = 3</b>	integer value	Default number of message sending attempts in case of any problems occurred: positive value N means at most N retries, 0 means no retries, None or -1 (or any other negative values) mean to retry forever. This option is used only if acknowledgments are enabled.
<b>rpc_thread_pool_size = 100</b>	integer value	Maximum number of (green) threads to work concurrently.
<b>rpc_use_acks = False</b>	boolean value	Wait for message acknowledgements from receivers. This mechanism works only via proxy without PUB/SUB.
<b>rpc_zmq_bind_address = *</b>	string value	ZeroMQ bind address. Should be a wildcard (*), an ethernet interface, or IP. The "host" option should point or resolve to this address.
<b>rpc_zmq_bind_port_retries = 100</b>	integer value	Number of retries to find free port number before fail with ZMQBindError.

Configuration option = Default value	Type	Description
<b>rpc_zmq_contexts = 1</b>	integer value	Number of ZeroMQ contexts, defaults to 1.
<b>rpc_zmq_host = ccosp-undercloud.usersys.redhat.com</b>	string value	Name of this node. Must be a valid hostname, FQDN, or IP address. Must match "host" option, if running Nova.
<b>rpc_zmq_ipc_dir = /var/run/openstack</b>	string value	Directory for holding IPC sockets.
<b>rpc_zmq_matchmaker = redis</b>	string value	MatchMaker driver.
<b>rpc_zmq_max_port = 65536</b>	integer value	Maximal port number for random ports range.
<b>rpc_zmq_min_port = 49153</b>	port value	Minimal port number for random ports range.
<b>rpc_zmq_serialization = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>rpc_zmq_topic_backlog = None</b>	integer value	Maximum number of ingress messages to locally buffer per topic. Default is unlimited.
<b>subscribe_on = []</b>	list value	List of publisher hosts SubConsumer can subscribe on. This option has higher priority then the default publishers list taken from the matchmaker.
<b>use_dynamic_connections = False</b>	boolean value	This option makes direct connections dynamic or static. It makes sense only with use_router_proxy=False which means to use direct connections for direct message types (ignored otherwise).
<b>use_pub_sub = False</b>	boolean value	Use PUB/SUB pattern for fanout methods. PUB/SUB always uses proxy.
<b>use_router_proxy = False</b>	boolean value	Use ROUTER remote proxy.
<b>zmq_failover_connections = 2</b>	integer value	How many additional connections to a host will be made for failover reasons. This option is actual only in dynamic connections mode.

Configuration option = Default value	Type	Description
<b>zmq_immediate = True</b>	boolean value	This option configures round-robin mode in zmq socket. True means not keeping a queue when server side disconnects. False means to keep queue and messages even if server is disconnected, when the server appears we send all accumulated messages to it.
<b>zmq_linger = -1</b>	integer value	Number of seconds to wait before all pending messages will be sent after closing a socket. The default value of -1 specifies an infinite linger period. The value of 0 specifies no linger period. Pending messages shall be discarded immediately when the socket is closed. Positive values specify an upper bound for the linger period.
<b>zmq_target_expire = 300</b>	integer value	Expiration timeout in seconds of a name service record about existing target (< 0 means no timeout).
<b>zmq_target_update = 180</b>	integer value	Update period in seconds of a name service record about existing target.
<b>zmq_tcp_keepalive = -1</b>	integer value	Enable/disable TCP keepalive (KA) mechanism. The default value of -1 (or any other negative value) means to skip any overrides and leave it to OS default; 0 and 1 (or any other positive value) mean to disable and enable the option respectively.
<b>zmq_tcp_keepalive_cnt = -1</b>	integer value	The number of retransmissions to be carried out before declaring that remote end is not available. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_idle = -1</b>	integer value	The duration between two keepalive transmissions in idle condition. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_intvl = -1</b>	integer value	The duration between two successive keepalive retransmissions, if acknowledgement to the previous keepalive transmission is not received. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.



### 11.1.39. oslo\_middlewares

The following table outlines the options available under the **[oslo\_middlewares]** group in the `/etc/nova/nova.conf` file.

Table 11.39. oslo\_middlewares

Configuration option = Default value	Type	Description
<b>enable_proxy_headers_parsing = False</b>	boolean value	Whether the application is behind a proxy or not. This determines if the middleware should parse the headers or not.
<b>max_request_body_size = 114688</b>	integer value	The maximum body size for each request, in bytes.
<b>secure_proxy_ssl_header = X-Forwarded-Proto</b>	string value	The HTTP Header that will be used to determine what the original request protocol scheme was, even if it was hidden by a SSL termination proxy.

### 11.1.40. oslo\_policy

The following table outlines the options available under the **[oslo\_policy]** group in the `/etc/nova/nova.conf` file.

Table 11.40. oslo\_policy

Configuration option = Default value	Type	Description
<b>enforce_scope = False</b>	boolean value	This option controls whether or not to enforce scope when evaluating policies. If <b>True</b> , the scope of the token used in the request is compared to the <b>scope_types</b> of the policy being enforced. If the scopes do not match, an <b>InvalidScope</b> exception will be raised. If <b>False</b> , a message will be logged informing operators that policies are being invoked with mismatching scope.
<b>policy_default_rule = default</b>	string value	Default rule. Enforced when a requested rule is not found.
<b>policy_dirs = [u'policy.d']</b>	multi valued	Directories where policy configuration files are stored. They can be relative to any directory in the search path defined by the <code>config_dir</code> option, or absolute paths. The file defined by <code>policy_file</code> must exist for these directories to be searched. Missing or empty directories are ignored.
<b>policy_file = policy.json</b>	string value	The file that defines policies.

Configuration option = Default value	Type	Description
<b>remote_content_type = application/x-www-form- urlencoded</b>	string value	Content Type to send and receive data for REST based policy check
<b>remote_ssl_ca_cert_file = None</b>	string value	Absolute path to ca cert file for REST based policy check
<b>remote_ssl_client_cert_file = None</b>	string value	Absolute path to client cert for REST based policy check
<b>remote_ssl_client_key_file = None</b>	string value	Absolute path client key file REST based policy check
<b>remote_ssl_verify_server _cert = False</b>	boolean value	server identity verification for REST based policy check

#### 11.1.41. pci

The following table outlines the options available under the **[pci]** group in the **/etc/nova/nova.conf** file.

Table 11.41. pci

Configuration option = Default value	Type	Description
---	------	-------------

Configuration option = Default value	Type	Description
<b>alias = []</b>	multi valued	<p>An alias for a PCI passthrough device requirement.</p> <p>This allows users to specify the alias in the extra specs for a flavor, without needing to repeat all the PCI property requirements.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>• A list of JSON values which describe the aliases. For example::</li> </ul> <pre>alias = {   "name": "QuickAssist",   "product_id": "0443",   "vendor_id": "8086",   "device_type": "type-PCI",   "numa_policy": "required" }</pre> <p>This defines an alias for the Intel QuickAssist card. (multi valued). Valid key values are :</p> <ul style="list-style-type: none"> <li><code>`name`</code> Name of the PCI alias.</li> <li><code>`product_id`</code> Product ID of the device in hexadecimal.</li> <li><code>`vendor_id`</code> Vendor ID of the device in hexadecimal.</li> <li><code>`device_type`</code> Type of PCI device. Valid values are: <code>`type-PCI`</code>, <code>`type-PF`</code> and <code>`type-VF`</code>.</li> <li><code>`numa_policy`</code> Required NUMA affinity of device. Valid values are: <code>`legacy`</code>, <code>`preferred`</code> and <code>`required`</code>.</li> </ul>
<b>passthrough_whitelist = []</b>	multi valued	<p>White list of PCI devices available to VMs.</p> <p>Possible values:</p>

Configuration option = Default value	Type	<ul style="list-style-type: none"> <li>A JSON dictionary which describe a whitelisted PCI device. It should take the following format:</li> </ul>
		<pre data-bbox="959 241 1422 456">["vendor_id": "&lt;id&gt;"], ["product_id": "&lt;id&gt;"], ["address": "[[[[&lt;domain&gt;]:]&lt;bus&gt;:]&lt;slot&gt;].[&lt;function&gt;]]" &amp;verbar; "devname": "&lt;name&gt;"], {"&lt;tag&gt;": "&lt;tag_value&gt;"},</pre> <p data-bbox="959 510 1422 757">Where '[' indicates zero or one occurrences, '{' indicates zero or multiple occurrences, '&amp;verbar;' mutually exclusive options. Note that any missing fields are automatically wildcarded.</p> <p data-bbox="959 810 1230 846">Valid key values are :</p> <ul data-bbox="879 891 1422 1352" style="list-style-type: none"> <li>"vendor_id": Vendor ID of the device in hexadecimal.</li> <li>"product_id": Product ID of the device in hexadecimal.</li> <li>"address": PCI address of the device.</li> <li>"devname": Device name of the device (for e.g. interface name). Not all PCI devices have a name.</li> <li>"&lt;tag&gt;": Additional &lt;tag&gt; and &lt;tag_value&gt; used for matching PCI devices. Supported &lt;tag&gt;: "physical_network", "trusted".</li> </ul> <p data-bbox="959 1393 1422 1496">The address key supports traditional glob style and regular expression syntax. Valid examples are:</p> <pre data-bbox="959 1554 1422 2123">passthrough_whitelist = {"devname":"eth0", "physical_network":"physnet"} passthrough_whitelist = {"address":"*:0a:00.*"} passthrough_whitelist = {"address":":0a:00.", "physical_network":"physnet1"} passthrough_whitelist = {"vendor_id":"1137", "product_id":"0071"} passthrough_whitelist = {"vendor_id":"1137", "product_id":"0071",</pre>

Configuration option = Default value	Type	Description
		<pre> "address": "0000:0a:00.1",  "physical_network":"physnet1"} passthrough_whitelist = {"address": {"domain": ". *", "bus": "02", "slot": "01", "function": "[2- 7]"},  "physical_network":"physnet1"} passthrough_whitelist = {"address": {"domain": ". *", "bus": "02", "slot": "0[1-2]", "function": ". *"}},  "physical_network":"physnet1"} passthrough_whitelist = {"devname": "eth0", "physical_network":"physnet1", "trusted": "true"} </pre> <p>The following are invalid, as they specify mutually exclusive options:</p> <pre> passthrough_whitelist = {"devname":"eth0",  "physical_network":"physnet", "address":":*:0a:00.*"} </pre> <ul style="list-style-type: none"> <li>• A JSON list of JSON dictionaries corresponding to the above format. For example: <pre> passthrough_whitelist = [{"product_id":"0001", "vendor_id":"8086"}, {"product_id":"0002", "vendor_id":"8086"}] </pre> </li> </ul>

### 11.1.42. placement

The following table outlines the options available under the **[placement]** group in the **/etc/nova/nova.conf** file.

Table 11.42. placement

Configuration option = Default value	Type	Description
<b>auth-url = None</b>	string value	Authentication URL
<b>auth_section = None</b>	string value	Config Section from which to load plugin specific options
<b>auth_type = None</b>	string value	Authentication type to load
<b>cafile = None</b>	string value	PEM encoded Certificate Authority to use when verifying HTTPs connections.
<b>certfile = None</b>	string value	PEM encoded client certificate cert file
<b>default-domain-id = None</b>	string value	Optional domain ID to use with v3 and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.
<b>default-domain-name = None</b>	string value	Optional domain name to use with v3 API and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.
<b>domain-id = None</b>	string value	Domain ID to scope to
<b>domain-name = None</b>	string value	Domain name to scope to
<b>endpoint-override = None</b>	string value	Always use this endpoint URL for requests for this client. NOTE: The unversioned endpoint should be specified here; to request a particular API version, use the <b>version</b> , <b>min-version</b> , and/or <b>max-version</b> options.
<b>insecure = False</b>	boolean value	Verify HTTPS connections.
<b>keyfile = None</b>	string value	PEM encoded client certificate key file
<b>os_interface = None</b>	string value	Endpoint interface for this node. This is used when picking the URL in the service catalog.
<b>os_region_name = None</b>	string value	Region name of this node. This is used when picking the URL in the service catalog.  Possible values: <ul style="list-style-type: none"> <li>Any string representing region name</li> </ul>

Configuration option = Default value	Type	Description
<b>password = None</b>	string value	User's password
<b>project-domain-id = None</b>	string value	Domain ID containing project
<b>project-domain-name = None</b>	string value	Domain name containing project
<b>project-id = None</b>	string value	Project ID to scope to
<b>project-name = None</b>	string value	Project name to scope to
<b>randomize_allocation_candidates = False</b>	boolean value	If True, when limiting allocation candidate results, the results will be a random sampling of the full result set. If False, allocation candidates are returned in a deterministic but undefined order. That is, all things being equal, two requests for allocation candidates will return the same results in the same order; but no guarantees are made as to how that order is determined.
<b>region-name = None</b>	string value	The default region_name for endpoint URL discovery.
<b>service-name = None</b>	string value	The default service_name for endpoint URL discovery.
<b>service-type = placement</b>	string value	The default service_type for endpoint URL discovery.
<b>system-scope = None</b>	string value	Scope for system operations
<b>tenant-id = None</b>	string value	Tenant ID
<b>tenant-name = None</b>	string value	Tenant Name
<b>timeout = None</b>	integer value	Timeout value for http requests
<b>trust-id = None</b>	string value	Trust ID
<b>user-domain-id = None</b>	string value	User's domain id
<b>user-domain-name = None</b>	string value	User's domain name
<b>user-id = None</b>	string value	User ID
<b>username = None</b>	string value	Username

Configuration option = Default value	Type	Description
<b>valid-interfaces =</b> <b>[u'internal', u'public']</b>	list value	List of interfaces, in order of preference, for endpoint URL.

### 11.1.43. quota

The following table outlines the options available under the **[quota]** group in the `/etc/nova/nova.conf` file.

Table 11.43. quota

Configuration option = Default value	Type	Description
<b>cores = 20</b>	integer value	<p>The number of instance cores or vCPUs allowed per project.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A positive integer or 0.</li> <li>• -1 to disable the quota.</li> </ul>
<b>driver =</b> <b>nova.quota.DbQuotaDriver</b>	string value	<p>The quota enforcer driver.</p> <p>Provides abstraction for quota checks. Users can configure a specific driver to use for quota checks.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• <code>nova.quota.DbQuotaDriver</code> (default) or any string representing fully qualified class name.</li> </ul>
<b>fixed_ips = -1</b>	integer value	<p>The number of fixed IPs allowed per project.</p> <p>Unlike floating IPs, fixed IPs are allocated dynamically by the network component when instances boot up. This quota value should be at least the number of instances allowed</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A positive integer or 0.</li> <li>• -1 to disable the quota.</li> </ul>



Configuration option = Default value	Type	Description
<b>floating_ips = 10</b>	integer value	<p>The number of floating IPs allowed per project.</p> <p>Floating IPs are not allocated to instances by default. Users need to select them from the pool configured by the OpenStack administrator to attach to their instances.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A positive integer or 0.</li> <li>● -1 to disable the quota.</li> </ul>
<b>injected_file_content_bytes = 10240</b>	integer value	<p>The number of bytes allowed per injected file.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A positive integer or 0.</li> <li>● -1 to disable the quota.</li> </ul>
<b>injected_file_path_length = 255</b>	integer value	<p>The maximum allowed injected file path length.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A positive integer or 0.</li> <li>● -1 to disable the quota.</li> </ul>
<b>injected_files = 5</b>	integer value	<p>The number of injected files allowed.</p> <p>File injection allows users to customize the personality of an instance by injecting data into it upon boot. Only text file injection is permitted: binary or ZIP files are not accepted. During file injection, any existing files that match specified files are renamed to include <b>.bak</b> extension appended with a timestamp.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A positive integer or 0.</li> <li>● -1 to disable the quota.</li> </ul>

Configuration option = Default value	Type	Description
<b>instances = 10</b>	integer value	<p>The number of instances allowed per project.</p> <p>Possible Values</p> <ul style="list-style-type: none"> <li>• A positive integer or 0.</li> <li>• -1 to disable the quota.</li> </ul>
<b>key_pairs = 100</b>	integer value	<p>The maximum number of key pairs allowed per user.</p> <p>Users can create at least one key pair for each project and use the key pair for multiple instances that belong to that project.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A positive integer or 0.</li> <li>• -1 to disable the quota.</li> </ul>
<b>max_age = 0</b>	integer value	<p>The number of seconds between subsequent usage refreshes.</p> <p>This defaults to 0 (off) to avoid additional load but it is useful to turn on to help keep quota usage up-to-date and reduce the impact of out of sync usage issues. Note that quotas are not updated on a periodic task, they will update on a new reservation if max_age has passed since the last reservation.</p>
<b>metadata_items = 128</b>	integer value	<p>The number of metadata items allowed per instance.</p> <p>Users can associate metadata with an instance during instance creation. This metadata takes the form of key-value pairs.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A positive integer or 0.</li> <li>• -1 to disable the quota.</li> </ul>
<b>ram = 51200</b>	integer value	<p>The number of megabytes of instance RAM allowed per project.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A positive integer or 0.</li> <li>• -1 to disable the quota.</li> </ul>

Configuration option = Default value	Type	Description
<b>recheck_quota = True</b>	boolean value	<p>Recheck quota after resource creation to prevent allowing quota to be exceeded.</p> <p>This defaults to True (recheck quota after resource creation) but can be set to False to avoid additional load if allowing quota to be exceeded because of racing requests is considered acceptable. For example, when set to False, if a user makes highly parallel REST API requests to create servers, it will be possible for them to create more servers than their allowed quota during the race. If their quota is 10 servers, they might be able to create 50 during the burst. After the burst, they will not be able to create any more servers but they will be able to keep their 50 servers until they delete them.</p> <p>The initial quota check is done before resources are created, so if multiple parallel requests arrive at the same time, all could pass the quota check and create resources, potentially exceeding quota. When recheck_quota is True, quota will be checked a second time after resources have been created and if the resource is over quota, it will be deleted and OverQuota will be raised, usually resulting in a 403 response to the REST API user. This makes it impossible for a user to exceed their quota with the caveat that it will, however, be possible for a REST API user to be rejected with a 403 response in the event of a collision close to reaching their quota limit, even if the user has enough quota available when they made the request.</p>
<b>reservation_expire = 86400</b>	integer value	<p>The number of seconds until a reservation expires.</p> <p>This quota represents the time period for invalidating quota reservations.</p>
<b>security_group_rules = 20</b>	integer value	<p>The number of security rules per security group.</p> <p>The associated rules in each security group control the traffic to instances in the group.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A positive integer or 0.</li> <li>● -1 to disable the quota.</li> </ul>

Configuration option = Default value	Type	Description
<b>security_groups = 10</b>	integer value	The number of security groups per project.  Possible values: <ul style="list-style-type: none"> <li>• A positive integer or 0.</li> <li>• -1 to disable the quota.</li> </ul>
<b>server_group_members = 10</b>	integer value	The maximum number of servers per server group.  Possible values: <ul style="list-style-type: none"> <li>• A positive integer or 0.</li> <li>• -1 to disable the quota.</li> </ul>
<b>server_groups = 10</b>	integer value	The maximum number of server groups per project.  Server groups are used to control the affinity and anti-affinity scheduling policy for a group of servers or instances. Reducing the quota will not affect any existing group, but new servers will not be allowed into groups that have become over quota.  Possible values: <ul style="list-style-type: none"> <li>• A positive integer or 0.</li> <li>• -1 to disable the quota.</li> </ul>
<b>until_refresh = 0</b>	integer value	The count of reservations until usage is refreshed.  This defaults to 0 (off) to avoid additional load but it is useful to turn on to help keep quota usage up-to-date and reduce the impact of out of sync usage issues.

#### 11.1.44. rdp

The following table outlines the options available under the **[rdp]** group in the `/etc/nova/nova.conf` file.

Table 11.44. rdp

Configuration option = Default value	Type	Description
---	------	-------------

Configuration option = Default value	Type	Description
<b>enabled = False</b>	boolean value	<p>Enable Remote Desktop Protocol (RDP) related features.</p> <p>Hyper-V, unlike the majority of the hypervisors employed on Nova compute nodes, uses RDP instead of VNC and SPICE as a desktop sharing protocol to provide instance console access. This option enables RDP for graphical console access for virtual machines created by Hyper-V.</p> <p><b>Note:</b> RDP should only be enabled on compute nodes that support the Hyper-V virtualization platform.</p> <p>Related options:</p> <ul style="list-style-type: none"><li>• <b>compute_driver:</b> Must be hyperv.</li></ul>

Configuration option = Default value	Type	Description
<b>html5_proxy_base_url =</b> <b>http://127.0.0.1:6083/</b>	uri value	<p>The URL an end user would use to connect to the RDP HTML5 console proxy. The console proxy service is called with this token-embedded URL and establishes the connection to the proper instance.</p> <p>An RDP HTML5 console proxy service will need to be configured to listen on the address configured here. Typically the console proxy service would be run on a controller node. The localhost address used as default would only work in a single node environment i.e. devstack.</p> <p>An RDP HTML5 proxy allows a user to access via the web the text or graphical console of any Windows server or workstation using RDP. RDP HTML5 console proxy services include FreeRDP, wsgate. See <a href="https://github.com/FreeRDP/FreeRDP-WebConnect">https://github.com/FreeRDP/FreeRDP-WebConnect</a></p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• &lt;scheme&gt;://&lt;ip-address&gt;:&lt;port-number&gt;/ <ul style="list-style-type: none"> <li>The scheme must be identical to the scheme configured for the RDP HTML5 console proxy service. It is `http` or `https`.</li> <li>The IP address must be identical to the address on which the RDP HTML5 console proxy service is listening.</li> <li>The port must be identical to the port on which the RDP HTML5 console proxy service is listening.</li> </ul> </li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• <b>rdp.enabled</b>: Must be set to <b>True</b> for <b>html5_proxy_base_url</b> to be effective.</li> </ul>

### 11.1.45. remote\_debug

The following table outlines the options available under the **[remote\_debug]** group in the **/etc/nova/nova.conf** file.

Table 11.45. remote\_debug

Configuration option = Default value	Type	Description
<b>host = None</b>	unknown value	<p>Debug host (IP or name) to connect to. This command line parameter is used when you want to connect to a nova service via a debugger running on a different host.</p> <p>Note that using the remote debug option changes how Nova uses the eventlet library to support async IO. This could result in failures that do not occur under normal operation. Use at your own risk.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>IP address of a remote host as a command line parameter to a nova service. For Example:</li> </ul> <pre>/usr/local/bin/nova-compute --config-file /etc/nova/nova.conf --remote_debug-host &lt;IP address where the debugger is running&gt;</pre>
<b>port = None</b>	port value	<p>Debug port to connect to. This command line parameter allows you to specify the port you want to use to connect to a nova service via a debugger running on different host.</p> <p>Note that using the remote debug option changes how Nova uses the eventlet library to support async IO. This could result in failures that do not occur under normal operation. Use at your own risk.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> <li>Port number you want to use as a command line parameter to a nova service. For Example:</li> </ul> <pre>/usr/local/bin/nova-compute --config-file /etc/nova/nova.conf --remote_debug-host &lt;IP address where the debugger is running&gt; --remote_debug-port &lt;port&gt; it's listening on&gt;.</pre>


### 11.1.46. scheduler

The following table outlines the options available under the **[scheduler]** group in the **/etc/nova/nova.conf** file.

Table 11.46. scheduler

Configuration option = Default value	Type	Description
<b>discover_hosts_in_cells_interval = -1</b>	integer value	<p>Periodic task interval.</p> <p>This value controls how often (in seconds) the scheduler should attempt to discover new hosts that have been added to cells. If negative (the default), no automatic discovery will occur.</p> <p>Deployments where compute nodes come and go frequently may want this enabled, where others may prefer to manually discover hosts when one is added to avoid any overhead from constantly checking. If enabled, every time this runs, we will select any unmapped hosts out of each cell database on every run.</p>
<b>driver = filter_scheduler</b>	string value	<p>The class of the driver used by the scheduler. This should be chosen from one of the entrypoints under the namespace <i>nova.scheduler.driver</i> of file <i>setup.cfg</i>. If nothing is specified in this option, the <i>filter_scheduler</i> is used.</p> <p>Other options are:</p> <ul style="list-style-type: none"> <li>● <i> caching_scheduler </i> which aggressively caches the system state for better individual scheduler performance at the risk of more retries when running multiple schedulers. [DEPRECATED]</li> <li>● <i> chance_scheduler </i> which simply picks a host at random. [DEPRECATED]</li> <li>● <i> fake_scheduler </i> which is used for testing.</li> </ul> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any of the drivers included in Nova: <ul style="list-style-type: none"> <li>○ <i> filter_scheduler </i></li> <li>○ <i> caching_scheduler </i></li> <li>○ <i> chance_scheduler </i></li> <li>○ <i> fake_scheduler </i></li> </ul> </li> <li>● You may also set this to the entry point name of a custom scheduler driver, but you will be responsible for creating and maintaining it in your <i>setup.cfg</i> file.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <i> workers </i></li> </ul>



Configuration option = Default value	Type	Description
<b>host_manager = host_manager</b>	string value	<p>The scheduler host manager to use.</p> <p>The host manager manages the in-memory picture of the hosts that the scheduler uses. The options values are chosen from the entry points under the namespace <i>nova.scheduler.host_manager</i> in <i>setup.cfg</i>.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p>The "ironic_host_manager" option is deprecated as of the 17.0.0 Queens release.</p> </div> </div>
<b>max_attempts = 3</b>	integer value	<p>This is the maximum number of attempts that will be made for a given instance build/move operation. It limits the number of alternate hosts returned by the scheduler. When that list of hosts is exhausted, a <code>MaxRetriesExceeded</code> exception is raised and the instance is set to an error state.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A positive integer, where the integer corresponds to the max number of attempts that can be made when building or moving an instance.</li> </ul>
<b>max_placement_results = 1000</b>	integer value	<p>This setting determines the maximum limit on results received from the placement service during a scheduling operation. It effectively limits the number of hosts that may be considered for scheduling requests that match a large number of candidates.</p> <p>A value of 1 (the minimum) will effectively defer scheduling to the placement service strictly on "will it fit" grounds. A higher value will put an upper cap on the number of results the scheduler will consider during the filtering and weighing process. Large deployments may need to set this lower than the total number of hosts available to limit memory consumption, network traffic, etc. of the scheduler.</p> <p>This option is only used by the <code>FilterScheduler</code>; if you use a different scheduler, this option has no effect.</p>

Configuration option = Default value	Type	Description
<b>periodic_task_interval = 60</b>	integer value	<p>Periodic task interval.</p> <p>This value controls how often (in seconds) to run periodic tasks in the scheduler. The specific tasks that are run for each period are determined by the particular scheduler being used.</p> <p>If this is larger than the nova-service <i>service_down_time</i> setting, Nova may report the scheduler service as down. This is because the scheduler driver is responsible for sending a heartbeat and it will only do that as often as this option allows. As each scheduler can work a little differently than the others, be sure to test this with your selected scheduler.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● An integer, where the integer corresponds to periodic task interval in seconds. 0 uses the default interval (60 seconds). A negative value disables periodic tasks.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>nova-service service_down_time</b></li> </ul>
<b>workers = None</b>	integer value	<p>Number of workers for the nova-scheduler service. The default will be the number of CPUs available if using the "filter_scheduler" scheduler driver, otherwise the default will be 1.</p>

### 11.1.47. serial\_console

The following table outlines the options available under the **[serial\_console]** group in the **/etc/nova/nova.conf** file.

Table 11.47. serial\_console

Configuration option = Default value	Type	Description
---	------	-------------

Configuration option = Default value	Type	Description
<b>base_url = ws://127.0.0.1:6083/</b>	uri value	<p>The URL an end user would use to connect to the <b>nova-serialproxy</b> service.</p> <p>The <b>nova-serialproxy</b> service is called with this token enriched URL and establishes the connection to the proper instance.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>• The IP address must be identical to the address to which the <b>nova-serialproxy</b> service is listening (see option <b>serialproxy_host</b> in this section).</li> <li>• The port must be the same as in the option <b>serialproxy_port</b> of this section.</li> <li>• If you choose to use a secured websocket connection, then start this option with <b>wss://</b> instead of the unsecured <b>ws://</b>. The options <b>cert</b> and <b>key</b> in the <b>[DEFAULT]</b> section have to be set for that.</li> </ul>
<b>enabled = False</b>	boolean value	<p>Enable the serial console feature.</p> <p>In order to use this feature, the service <b>nova-serialproxy</b> needs to run. This service is typically executed on the controller node.</p>
<b>port_range = 10000:20000</b>	string value	<p>A range of TCP ports a guest can use for its backend.</p> <p>Each instance which gets created will use one port out of this range. If the range is not big enough to provide another port for an new instance, this instance won't get launched.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Each string which passes the regex <b>\d+:\d+</b>. For example <b>10000:20000</b>. Be sure that the first port number is lower than the second port number and that both are in range from 0 to 65535.</li> </ul>
<b>proxycient_address = 127.0.0.1</b>	string value	<p>The IP address to which proxy clients (like <b>nova-serialproxy</b>) should connect to get the serial console of an instance.</p> <p>This is typically the IP address of the host of a <b>nova-compute</b> service.</p>

Configuration option = Default value	Type	Description
<b>serialproxy_host = 0.0.0.0</b>	string value	<p>The IP address which is used by the <b>nova-serialproxy</b> service to listen for incoming requests.</p> <p>The <b>nova-serialproxy</b> service listens on this IP address for incoming connection requests to instances which expose serial console.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>Ensure that this is the same IP address which is defined in the option <b>base_url</b> of this section or use <b>0.0.0.0</b> to listen on all addresses.</li> </ul>
<b>serialproxy_port = 6083</b>	port value	<p>The port number which is used by the <b>nova-serialproxy</b> service to listen for incoming requests.</p> <p>The <b>nova-serialproxy</b> service listens on this port number for incoming connection requests to instances which expose serial console.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>Ensure that this is the same port number which is defined in the option <b>base_url</b> of this section.</li> </ul>

### 11.1.48. service\_user

The following table outlines the options available under the **[service\_user]** group in the **/etc/nova/nova.conf** file.

Table 11.48. service\_user

Configuration option = Default value	Type	Description
<b>auth-url = None</b>	string value	Authentication URL
<b>auth_section = None</b>	string value	Config Section from which to load plugin specific options
<b>auth_type = None</b>	string value	Authentication type to load
<b>cafile = None</b>	string value	PEM encoded Certificate Authority to use when verifying HTTPs connections.
<b>certfile = None</b>	string value	PEM encoded client certificate cert file

Configuration option = Default value	Type	Description
<b>default-domain-id = None</b>	string value	Optional domain ID to use with v3 and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.
<b>default-domain-name = None</b>	string value	Optional domain name to use with v3 API and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.
<b>domain-id = None</b>	string value	Domain ID to scope to
<b>domain-name = None</b>	string value	Domain name to scope to
<b>insecure = False</b>	boolean value	Verify HTTPS connections.
<b>keyfile = None</b>	string value	PEM encoded client certificate key file
<b>password = None</b>	string value	User's password
<b>project-domain-id = None</b>	string value	Domain ID containing project
<b>project-domain-name = None</b>	string value	Domain name containing project
<b>project-id = None</b>	string value	Project ID to scope to
<b>project-name = None</b>	string value	Project name to scope to
<b>send_service_user_token = False</b>	boolean value	<p>When True, if sending a user token to a REST API, also send a service token.</p> <p>Nova often reuses the user token provided to the nova-api to talk to other REST APIs, such as Cinder, Glance and Neutron. It is possible that while the user token was valid when the request was made to Nova, the token may expire before it reaches the other service. To avoid any failures, and to make it clear it is Nova calling the service on the user's behalf, we include a service token along with the user token. Should the user's token have expired, a valid service token ensures the REST API request will still be accepted by the keystone middleware.</p>
<b>system-scope = None</b>	string value	Scope for system operations
<b>tenant-id = None</b>	string value	Tenant ID

Configuration option = Default value	Type	Description
<b>tenant-name = None</b>	string value	Tenant Name
<b>timeout = None</b>	integer value	Timeout value for http requests
<b>trust-id = None</b>	string value	Trust ID
<b>user-domain-id = None</b>	string value	User's domain id
<b>user-domain-name = None</b>	string value	User's domain name
<b>user-id = None</b>	string value	User ID
<b>username = None</b>	string value	Username

### 11.1.49. spice

The following table outlines the options available under the **[spice]** group in the `/etc/nova/nova.conf` file.

Table 11.49. spice

Configuration option = Default value	Type	Description
<b>agent_enabled = True</b>	boolean value	<p>Enable the SPICE guest agent support on the instances.</p> <p>The Spice agent works with the Spice protocol to offer a better guest console experience. However, the Spice console can still be used without the Spice Agent. With the Spice agent installed the following features are enabled:</p> <ul style="list-style-type: none"> <li>● Copy &amp; Paste of text and images between the guest and client machine</li> <li>● Automatic adjustment of resolution when the client screen changes - e.g. if you make the Spice console full screen the guest resolution will adjust to match it rather than letterboxing.</li> <li>● Better mouse integration - The mouse can be captured and released without needing to click inside the console or press keys to release it. The performance of mouse movement is also improved.</li> </ul>

Configuration option = Default value	Type	Description
<b>enabled = False</b>	boolean value	<p>Enable SPICE related features.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>• VNC must be explicitly disabled to get access to the SPICE console. Set the <code>enabled</code> option to <code>False</code> in the <code>[vnc]</code> section to disable the VNC console.</li> </ul>
<b>html5proxy_base_url = http://127.0.0.1:6082/spice _auto.html</b>	uri value	<p>Location of the SPICE HTML5 console proxy.</p> <p>End user would use this URL to connect to the <b>nova-spicehtml5proxy</b> service. This service will forward request to the console of an instance.</p> <p>In order to use SPICE console, the service <b>nova-spicehtml5proxy</b> should be running. This service is typically launched on the controller node.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• Must be a valid URL of the form: <b>http://host:port/spice_auto.html</b> where <code>host</code> is the node running <b>nova-spicehtml5proxy</b> and the port is typically 6082. Consider not using default value as it is not well defined for any real deployment.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• This option depends on <b>html5proxy_host</b> and <b>html5proxy_port</b> options. The access URL returned by the compute node must have the host and port where the <b>nova-spicehtml5proxy</b> service is listening.</li> </ul>
<b>html5proxy_host = 0.0.0.0</b>	unknown value	<p>IP address or a hostname on which the <b>nova-spicehtml5proxy</b> service listens for incoming requests.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>• This option depends on the <b>html5proxy_base_url</b> option. The <b>nova-spicehtml5proxy</b> service must be listening on a host that is accessible from the HTML5 client.</li> </ul>

Configuration option = Default value	Type	Description
<b>html5proxy_port = 6082</b>	port value	<p>Port on which the <b>nova-spicehtml5proxy</b> service listens for incoming requests.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>This option depends on the <b>html5proxy_base_url</b> option. The <b>nova-spicehtml5proxy</b> service must be listening on a port that is accessible from the HTML5 client.</li> </ul>
<b>keymap = en-us</b>	string value	<p>A keyboard layout which is supported by the underlying hypervisor on this node.</p> <p>Possible values: * This is usually an <i>IETF language tag</i> (default is <i>en-us</i>). If you use QEMU as hypervisor, you should find the list of supported keyboard layouts at <code>/usr/share/qemu/keymaps</code>.</p>
<b>server_listen = 127.0.0.1</b>	string value	<p>The address where the SPICE server running on the instances should listen.</p> <p>Typically, the <b>nova-spicehtml5proxy</b> proxy client runs on the controller node and connects over the private network to this address on the compute node(s).</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>IP address to listen on.</li> </ul>
<b>server_proxyclient_addresses = 127.0.0.1</b>	string value	<p>The address used by <b>nova-spicehtml5proxy</b> client to connect to instance console.</p> <p>Typically, the <b>nova-spicehtml5proxy</b> proxy client runs on the controller node and connects over the private network to this address on the compute node(s).</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any valid IP address on the compute node.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>This option depends on the <b>server_listen</b> option. The proxy client must be able to access the address specified in <b>server_listen</b> using the value of this option.</li> </ul>



### 11.1.50. upgrade\_levels

The following table outlines the options available under the **[upgrade\_levels]** group in the `/etc/nova/nova.conf` file.

Table 11.50. upgrade\_levels

Configuration option = Default value	Type	Description
<b>baseapi = None</b>	string value	Base API RPC API version cap
<b>cells = None</b>	string value	Cells RPC API version cap
<b>cert = None</b>	string value	Cert RPC API version cap
<b>compute = None</b>	string value	<p>Compute RPC API version cap.</p> <p>By default, we always send messages using the most recent version the client knows about.</p> <p>Where you have old and new compute services running, you should set this to the lowest deployed version. This is to guarantee that all services never send messages that one of the compute nodes can't understand. Note that we only support upgrading from release N to release N+1.</p> <p>Set this option to "auto" if you want to let the compute RPC module automatically determine what version to use based on the service versions in the deployment.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● By default send the latest version the client knows about</li> <li>● <i>auto</i>: Automatically determines what version to use based on the service versions in the deployment.</li> <li>● A string representing a version number in the format <i>N.N</i>; for example, possible values might be <i>1.12</i> or <i>2.0</i>.</li> <li>● An OpenStack release name, in lower case, such as <i>mitaka</i> or <i>liberty</i>.</li> </ul>
<b>conductor = None</b>	string value	Conductor RPC API version cap
<b>console = None</b>	string value	Console RPC API version cap
<b>consoleauth = None</b>	string value	Consoleauth RPC API version cap

Configuration option = Default value	Type	Description
<b>intercell = None</b>	string value	Inter-cell RPC API version cap
<b>network = None</b>	string value	Network RPC API version cap
<b>scheduler = None</b>	string value	Scheduler RPC API version cap

### 11.1.51. vault

The following table outlines the options available under the **[vault]** group in the **/etc/nova/nova.conf** file.

Table 11.51. vault

Configuration option = Default value	Type	Description
<b>root_token_id = None</b>	string value	root token for vault
<b>ssl_ca_cert_file = None</b>	string value	Absolute path to ca cert file
<b>use_ssl = False</b>	boolean value	SSL Enabled/Disabled
<b>vault_url = http://127.0.0.1:8200</b>	string value	Use this endpoint to connect to Vault, for example: "http://127.0.0.1:8200"

### 11.1.52. vendordata\_dynamic\_auth

The following table outlines the options available under the **[vendordata\_dynamic\_auth]** group in the **/etc/nova/nova.conf** file.

Table 11.52. vendordata\_dynamic\_auth

Configuration option = Default value	Type	Description
<b>auth-url = None</b>	string value	Authentication URL
<b>auth_section = None</b>	string value	Config Section from which to load plugin specific options
<b>auth_type = None</b>	string value	Authentication type to load
<b>cafile = None</b>	string value	PEM encoded Certificate Authority to use when verifying HTTPs connections.

Configuration option = Default value	Type	Description
<b>certfile = None</b>	string value	PEM encoded client certificate cert file
<b>default-domain-id = None</b>	string value	Optional domain ID to use with v3 and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.
<b>default-domain-name = None</b>	string value	Optional domain name to use with v3 API and v2 parameters. It will be used for both the user and project domain in v3 and ignored in v2 authentication.
<b>domain-id = None</b>	string value	Domain ID to scope to
<b>domain-name = None</b>	string value	Domain name to scope to
<b>insecure = False</b>	boolean value	Verify HTTPS connections.
<b>keyfile = None</b>	string value	PEM encoded client certificate key file
<b>password = None</b>	string value	User's password
<b>project-domain-id = None</b>	string value	Domain ID containing project
<b>project-domain-name = None</b>	string value	Domain name containing project
<b>project-id = None</b>	string value	Project ID to scope to
<b>project-name = None</b>	string value	Project name to scope to
<b>system-scope = None</b>	string value	Scope for system operations
<b>tenant-id = None</b>	string value	Tenant ID
<b>tenant-name = None</b>	string value	Tenant Name
<b>timeout = None</b>	integer value	Timeout value for http requests
<b>trust-id = None</b>	string value	Trust ID
<b>user-domain-id = None</b>	string value	User's domain id
<b>user-domain-name = None</b>	string value	User's domain name

Configuration option = Default value	Type	Description
<b>user-id = None</b>	string value	User ID
<b>username = None</b>	string value	Username

### 11.1.53. vmware

The following table outlines the options available under the **[vmware]** group in the `/etc/nova/nova.conf` file.

Table 11.53. vmware

Configuration option = Default value	Type	Description
<b>api_retry_count = 10</b>	integer value	Number of times VMware vCenter server API must be retried on connection failures, e.g. socket error, etc.
<b>ca_file = None</b>	string value	Specifies the CA bundle file to be used in verifying the vCenter server certificate.
<b>cache_prefix = None</b>	string value	<p>This option adds a prefix to the folder where cached images are stored</p> <p>This is not the full path - just a folder prefix. This should only be used when a datastore cache is shared between compute nodes.</p> <p>Note: This should only be used when the compute nodes are running on same host or they have a shared file system.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any string representing the cache prefix to the folder</li> </ul>
<b>cluster_name = None</b>	string value	Name of a VMware Cluster ComputeResource.
<b>connection_pool_size = 10</b>	integer value	<p>This option sets the http connection pool size</p> <p>The connection pool size is the maximum number of connections from nova to vSphere. It should only be increased if there are warnings indicating that the connection pool is full, otherwise, the default should suffice.</p>

Configuration option = Default value	Type	Description
<b>console_delay_seconds = None</b>	integer value	Set this value if affected by an increased network latency causing repeated characters when typing in a remote console.
<b>datastore_regex = None</b>	string value	<p>Regular expression pattern to match the name of datastore.</p> <p>The <code>datastore_regex</code> setting specifies the datastores to use with Compute. For example, <code>datastore_regex="nas.*"</code> selects all the data stores that have a name starting with "nas".</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p><b>NOTE</b></p> <p>If no regex is given, it just picks the datastore with the most freespace.</p> </div> </div> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any matching regular expression to a datastore must be given</li> </ul>
<b>host_ip = None</b>	unknown value	Hostname or IP address for connection to VMware vCenter host.
<b>host_password = None</b>	string value	Password for connection to VMware vCenter host.
<b>host_port = 443</b>	port value	Port for connection to VMware vCenter host.
<b>host_username = None</b>	string value	Username for connection to VMware vCenter host.
<b>insecure = False</b>	boolean value	<p>If true, the vCenter server certificate is not verified. If false, then the default CA truststore is used for verification.</p> <p>Related options: * <code>ca_file</code>: This option is ignored if "<code>ca_file</code>" is set.</p>

Configuration option = Default value	Type	Description
<b>integration_bridge = None</b>	string value	<p>This option should be configured only when using the NSX-MH Neutron plugin. This is the name of the integration bridge on the ESXi server or host. This should not be set for any other Neutron plugin. Hence the default value is not set.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any valid string representing the name of the integration bridge</li> </ul>
<b>maximum_objects = 100</b>	integer value	<p>This option specifies the limit on the maximum number of objects to return in a single result.</p> <p>A positive value will cause the operation to suspend the retrieval when the count of objects reaches the specified limit. The server may still limit the count to something less than the configured value. Any remaining objects may be retrieved with additional requests.</p>
<b>pbm_default_policy = None</b>	string value	<p>This option specifies the default policy to be used.</p> <p>If <code>pbm_enabled</code> is set and there is no defined storage policy for the specific request, then this policy will be used.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any valid storage policy such as VSAN default storage policy</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li><code>pbm_enabled</code></li> </ul>
<b>pbm_enabled = False</b>	boolean value	<p>This option enables or disables storage policy based placement of instances.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li><code>pbm_default_policy</code></li> </ul>

Configuration option = Default value	Type	Description
<b>pbm_wsdl_location = None</b>	string value	<p>This option specifies the PBM service WSDL file location URL.</p> <p>Setting this will disable storage policy based placement of instances.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any valid file path e.g <a href="file:///opt/SDK/spbm/wsdl/pbmService.wsdl">file:///opt/SDK/spbm/wsdl/pbmService.wsdl</a></li> </ul>
<b>serial_log_dir = /opt/vmware/vspc</b>	string value	Specifies the directory where the Virtual Serial Port Concentrator is storing console log files. It should match the <i>serial_log_dir</i> config value of VSPC.
<b>serial_port_proxy_uri = None</b>	uri value	<p>Identifies a proxy service that provides network access to the <i>serial_port_service_uri</i>.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any valid URI (The scheme is <i>telnet</i> or <i>telnets</i>.)</li> </ul> <p>Related options: This option is ignored if <i>serial_port_service_uri</i> is not specified. * <i>serial_port_service_uri</i></p>
<b>serial_port_service_uri = None</b>	string value	<p>Identifies the remote system where the serial port traffic will be sent.</p> <p>This option adds a virtual serial port which sends console output to a configurable service URI. At the service URI address there will be virtual serial port concentrator that will collect console logs. If this is not set, no serial ports will be added to the created VMs.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any valid URI</li> </ul>
<b>task_poll_interval = 0.5</b>	floating point value	Time interval in seconds to poll remote tasks invoked on VMware VC server.

Configuration option = Default value	Type	Description
<b>use_linked_clone = True</b>	boolean value	<p>This option enables/disables the use of linked clone.</p> <p>The ESX hypervisor requires a copy of the VMDK file in order to boot up a virtual machine. The compute driver must download the VMDK via HTTP from the OpenStack Image service to a datastore that is visible to the hypervisor and cache it. Subsequent virtual machines that need the VMDK use the cached version and don't have to copy the file again from the OpenStack Image service.</p> <p>If set to false, even with a cached VMDK, there is still a copy operation from the cache location to the hypervisor file directory in the shared datastore. If set to true, the above copy operation is avoided as it creates copy of the virtual machine that shares virtual disks with its parent VM.</p>
<b>vlan_interface = vmnic0</b>	string value	<p>This option specifies the physical ethernet adapter name for VLAN networking.</p> <p>Set the <code>vlan_interface</code> configuration option to match the ESX host interface that handles VLAN-tagged VM traffic.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any valid string representing VLAN interface name</li> </ul>
<b>vnc_port = 5900</b>	port value	<p>This option specifies VNC starting port.</p> <p>Every VM created by ESX host has an option of enabling VNC client for remote connection. Above option <code>vnc_port</code> helps you to set default starting port for the VNC client.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any valid port number within 5900 -(5900 + <code>vnc_port_total</code>)</li> </ul> <p>Related options: Below options should be set to enable VNC client. * <code>vnc.enabled = True</code> * <code>vnc_port_total</code></p>
<b>vnc_port_total = 10000</b>	integer value	Total number of VNC ports.

### 11.1.54. vnc



The following table outlines the options available under the **[vnc]** group in the `/etc/nova/nova.conf` file.

**Table 11.54. vnc**

Configuration option = Default value	Type	Description
<b>auth_schemes = [u'none']</b>	list value	<p>The authentication schemes to use with the compute node.</p> <p>Control what RFB authentication schemes are permitted for connections between the proxy and the compute host. If multiple schemes are enabled, the first matching scheme will be used, thus the strongest schemes should be listed first.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● <b>none</b>: allow connection without authentication</li> <li>● <b>vencrypt</b>: use VeNCrypt authentication scheme</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>[vnc]vencrypt_client_key</b>, <b>[vnc]vencrypt_client_cert</b>: must also be set</li> </ul>
<b>enabled = True</b>	boolean value	<p>Enable VNC related features.</p> <p>Guests will get created with graphical devices to support this. Clients (for example Horizon) can then establish a VNC connection to the guest.</p>
<b>keymap = en-us</b>	string value	<p>Keymap for VNC.</p> <p>The keyboard mapping (keymap) determines which keyboard layout a VNC session should use by default.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● A keyboard layout which is supported by the underlying hypervisor on this node. This is usually an <i>IETF language tag</i> (for example <i>en-us</i>). If you use QEMU as hypervisor, you should find the list of supported keyboard layouts at <code>/usr/share/qemu/keymaps</code>.</li> </ul>

Configuration option = Default value	Type	Description
<b>novncproxy_base_url = http://127.0.0.1:6080/vnc_ auto.html</b>	uri value	<p>Public address of noVNC VNC console proxy.</p> <p>The VNC proxy is an OpenStack component that enables compute service users to access their instances through VNC clients. noVNC provides VNC support through a websocket-based client.</p> <p>This option sets the public base URL to which client systems will connect. noVNC clients can use this address to connect to the noVNC instance and, by extension, the VNC sessions.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● novncproxy_host</li> <li>● novncproxy_port</li> </ul>
<b>novncproxy_host = 0.0.0.0</b>	string value	<p>IP address that the noVNC console proxy should bind to.</p> <p>The VNC proxy is an OpenStack component that enables compute service users to access their instances through VNC clients. noVNC provides VNC support through a websocket-based client.</p> <p>This option sets the private address to which the noVNC console proxy service should bind to.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● novncproxy_port</li> <li>● novncproxy_base_url</li> </ul>
<b>novncproxy_port = 6080</b>	port value	<p>Port that the noVNC console proxy should bind to.</p> <p>The VNC proxy is an OpenStack component that enables compute service users to access their instances through VNC clients. noVNC provides VNC support through a websocket-based client.</p> <p>This option sets the private port to which the noVNC console proxy service should bind to.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● novncproxy_host</li> <li>● novncproxy_base_url</li> </ul>

Configuration option = Default value	Type	Description
<b>server_listen = 127.0.0.1</b>	unknown value	The IP address or hostname on which an instance should listen to for incoming VNC connection requests on this node.
<b>server_proxyclient_address = 127.0.0.1</b>	unknown value	<p>Private, internal IP address or hostname of VNC console proxy.</p> <p>The VNC proxy is an OpenStack component that enables compute service users to access their instances through VNC clients.</p> <p>This option sets the private address to which proxy clients, such as <b>nova-xvncproxy</b>, should connect to.</p>
<b>vencrypt_ca_certs = None</b>	string value	<p>The path to the CA certificate PEM file</p> <p>The fully qualified path to a PEM file containing one or more x509 certificates for the certificate authorities used by the compute node VNC server.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>vnc.auth_schemes</b>: must include <b>vencrypt</b></li> </ul>
<b>vencrypt_client_cert = None</b>	string value	<p>The path to the client key file (for x509)</p> <p>The fully qualified path to a PEM file containing the x509 certificate which the VNC proxy server presents to the compute node during VNC authentication.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>vnc.auth_schemes</b>: must include <b>vencrypt</b></li> <li>● <b>vnc.vencrypt_client_key</b>: must also be set</li> </ul>

Configuration option = Default value	Type	Description
<b>vencrypt_client_key = None</b>	string value	<p>The path to the client certificate PEM file (for x509)</p> <p>The fully qualified path to a PEM file containing the private key which the VNC proxy server presents to the compute node during VNC authentication.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>vnc.auth_schemes</b>: must include <b>vencrypt</b></li> <li>● <b>vnc.vencrypt_client_cert</b>: must also be set</li> </ul>
<b>xvpvncproxy_base_url = http://127.0.0.1:6081/console</b>	uri value	<p>Public URL address of XVP VNC console proxy.</p> <p>The VNC proxy is an OpenStack component that enables compute service users to access their instances through VNC clients. Xen provides the Xenserver VNC Proxy, or XVP, as an alternative to the websocket-based noVNC proxy used by Libvirt. In contrast to noVNC, XVP clients are Java-based.</p> <p>This option sets the public base URL to which client systems will connect. XVP clients can use this address to connect to the XVP instance and, by extension, the VNC sessions.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● xvpvncproxy_host</li> <li>● xvpvncproxy_port</li> </ul>

Configuration option = Default value	Type	Description
<b>xvpvncproxy_host = 0.0.0.0</b>	unknown value	<p>IP address or hostname that the XVP VNC console proxy should bind to.</p> <p>The VNC proxy is an OpenStack component that enables compute service users to access their instances through VNC clients. Xen provides the Xenserver VNC Proxy, or XVP, as an alternative to the websocket-based noVNC proxy used by Libvirt. In contrast to noVNC, XVP clients are Java-based.</p> <p>This option sets the private address to which the XVP VNC console proxy service should bind to.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● xvpvncproxy_port</li> <li>● xvpvncproxy_base_url</li> </ul>
<b>xvpvncproxy_port = 6081</b>	port value	<p>Port that the XVP VNC console proxy should bind to.</p> <p>The VNC proxy is an OpenStack component that enables compute service users to access their instances through VNC clients. Xen provides the Xenserver VNC Proxy, or XVP, as an alternative to the websocket-based noVNC proxy used by Libvirt. In contrast to noVNC, XVP clients are Java-based.</p> <p>This option sets the private port to which the XVP VNC console proxy service should bind to.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● xvpvncproxy_host</li> <li>● xvpvncproxy_base_url</li> </ul>

### 11.1.55. workarounds

The following table outlines the options available under the **[workarounds]** group in the **/etc/nova/nova.conf** file.

Table 11.55. workarounds

Configuration option = Default value	Type	Description
---	------	-------------

Configuration option = Default value	Type	Description
<b>disable_group_policy_check_upcall = False</b>	boolean value	<p>Disable the server group policy check upcall in compute.</p> <p>In order to detect races with server group affinity policy, the compute service attempts to validate that the policy was not violated by the scheduler. It does this by making an upcall to the API database to list the instances in the server group for one that it is booting, which violates our api/cell isolation goals. Eventually this will be solved by proper affinity guarantees in the scheduler and placement service, but until then, this late check is needed to ensure proper affinity policy.</p> <p>Operators that desire api/cell isolation over this check should enable this flag, which will avoid making that upcall from compute.</p> <p>Related options:</p> <ul style="list-style-type: none"><li>● [filter_scheduler]/track_instance_changes also relies on upcalls from the compute service to the scheduler service.</li></ul>

Configuration option = Default value	Type	Description
<b>disable_libvirt_livesnaps hot = False</b>	boolean value	<p>Disable live snapshots when using the libvirt driver.</p> <p>Live snapshots allow the snapshot of the disk to happen without an interruption to the guest, using coordination with a guest agent to quiesce the filesystem.</p> <p>When using libvirt 1.2.2 live snapshots fail intermittently under load (likely related to concurrent libvirt/qemu operations). This config option provides a mechanism to disable live snapshot, in favor of cold snapshot, while this is resolved. Cold snapshot causes an instance outage while the guest is going through the snapshotting process.</p> <p>For more information, refer to the bug report:</p> <p><a href="https://bugs.launchpad.net/nova/+bug/1334398">https://bugs.launchpad.net/nova/+bug/1334398</a></p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True: Live snapshot is disabled when using libvirt</li> <li>● False: Live snapshots are always used when snapshotting (as long as there is a new enough libvirt and the backend storage supports it)</li> </ul>
<b>disable_rootwrap = False</b>	boolean value	<p>Use sudo instead of rootwrap.</p> <p>Allow fallback to sudo for performance reasons.</p> <p>For more information, refer to the bug report:</p> <p><a href="https://bugs.launchpad.net/nova/+bug/1415106">https://bugs.launchpad.net/nova/+bug/1415106</a></p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● True: Use sudo instead of rootwrap</li> <li>● False: Use rootwrap as usual</li> </ul> <p>Interdependencies to other options:</p> <ul style="list-style-type: none"> <li>● Any options that affect <i>rootwrap</i> will be ignored.</li> </ul>

Configuration option = Default value	Type	Description
<b>ensure_libvirt_rbd_instance_dir_cleanup = False</b>	boolean value	<p>Ensure the instance directory is removed during clean up when using rbd.</p> <p>When enabled this workaround will ensure that the instance directory is always removed during cleanup on hosts using <b>[libvirt]/images_type=rbd</b>. This avoids the following bugs with evacuation and revert resize clean up that lead to the instance directory remaining on the host:</p> <p><a href="https://bugs.launchpad.net/nova/+bug/1414895">https://bugs.launchpad.net/nova/+bug/1414895</a></p> <p><a href="https://bugs.launchpad.net/nova/+bug/1761062">https://bugs.launchpad.net/nova/+bug/1761062</a></p> <p>Both of these bugs can then result in <b>DestinationDiskExists</b> errors being raised if the instances ever attempt to return to the host.</p> <ol style="list-style-type: none"> <li>warning:: Operators will need to ensure that the instance directory itself, specified by <b>[DEFAULT]/instances_path</b>, is not shared between computes before enabling this workaround otherwise the console.log, kernels, ramdisks and any additional files being used by the running instance will be lost.</li> </ol> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>compute_driver</b> (libvirt)</li> <li>● <b>[libvirt]/images_type</b> (rbd)</li> <li>● <b>instances_path</b></li> </ul>



Configuration option = Default value	Type	Description
<b>handle_virt_lifecycle_events = True</b>	boolean value	<p>Enable handling of events emitted from compute drivers.</p> <p>Many compute drivers emit lifecycle events, which are events that occur when, for example, an instance is starting or stopping. If the instance is going through task state changes due to an API operation, like <code>resize</code>, the events are ignored.</p> <p>This is an advanced feature which allows the hypervisor to signal to the compute service that an unexpected state change has occurred in an instance and that the instance can be shutdown automatically. Unfortunately, this can race in some conditions, for example in reboot operations or when the compute service or when host is rebooted (planned or due to an outage). If such races are common, then it is advisable to disable this feature.</p> <p>Care should be taken when this feature is disabled and <code>sync_power_state_interval</code> is set to a negative value. In this case, any instances that get out of sync between the hypervisor and the Nova database will have to be synchronized manually.</p> <p>For more information, refer to the bug report:</p> <p><a href="https://bugs.launchpad.net/bugs/1444630">https://bugs.launchpad.net/bugs/1444630</a></p> <p>Interdependencies to other options:</p> <ul style="list-style-type: none"> <li>• If <b><code>sync_power_state_interval</code></b> is negative and this feature is disabled, then instances that get out of sync between the hypervisor and the Nova database will have to be synchronized manually.</li> </ul>


### 11.1.56. wsgi

The following table outlines the options available under the **[wsgi]** group in the `/etc/nova/nova.conf` file.

Table 11.56. wsgi

Configuration option = Default value	Type	Description
---	------	-------------

Configuration option = Default value	Type	Description
<b>api_paste_config = api-paste.ini</b>	string value	<p>This option represents a file name for the paste.deploy config for nova-api.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• A string representing file name for the paste.deploy config.</li> </ul>
<b>client_socket_timeout = 900</b>	integer value	<p>This option specifies the timeout for client connections' socket operations. If an incoming connection is idle for this number of seconds it will be closed. It indicates timeout on individual read/writes on the socket connection. To wait forever set to 0.</p>
<b>default_pool_size = 1000</b>	integer value	<p>This option specifies the size of the pool of greenthreads used by wsgi. It is possible to limit the number of concurrent connections using this option.</p>
<b>keep_alive = True</b>	boolean value	<p>This option allows using the same TCP connection to send and receive multiple HTTP requests/responses, as opposed to opening a new one for every single request/response pair. HTTP keep-alive indicates HTTP connection reuse.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• True : reuse HTTP connection.</li> <li>• False : closes the client socket connection explicitly.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>• tcp_keepidle</li> </ul>
<b>max_header_line = 16384</b>	integer value	<p>This option specifies the maximum line size of message headers to be accepted. max_header_line may need to be increased when using large tokens (typically those generated by the Keystone v3 API with big service catalogs).</p> <p>Since TCP is a stream based protocol, in order to reuse a connection, the HTTP has to have a way to indicate the end of the previous response and beginning of the next. Hence, in a keep_alive case, all messages must have a self-defined message length.</p>

Configuration option = Default value	Type	Description
<b>secure_proxy_ssl_header = None</b>	string value	<p>This option specifies the HTTP header used to determine the protocol scheme for the original request, even if it was removed by a SSL terminating proxy.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● None (default) - the request scheme is not influenced by any HTTP headers</li> <li>● Valid HTTP header, like HTTP_X_FORWARDED_PROTO</li> </ul> <div data-bbox="817 696 1428 983" style="background-color: #fff9c4; padding: 10px; border: 1px solid #ccc;">  <p><b>WARNING</b></p> <p>Do not set this unless you know what you are doing.</p> </div> <p>Make sure ALL of the following are true before setting this (assuming the values from the example above): * Your API is behind a proxy. * Your proxy strips the X-Forwarded-Proto header from all incoming requests. In other words, if end users include that header in their requests, the proxy will discard it. * Your proxy sets the X-Forwarded-Proto header and sends it to API, but only for requests that originally come in via HTTPS.</p> <p>If any of those are not true, you should keep this setting set to None.</p>
<b>ssl_ca_file = None</b>	string value	<p>This option allows setting path to the CA certificate file that should be used to verify connecting clients.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● String representing path to the CA certificate file.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● enabled_ssl_apis</li> </ul>

Configuration option = Default value	Type	Description
<b>ssl_cert_file = None</b>	string value	<p>This option allows setting path to the SSL certificate of API server.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● String representing path to the SSL certificate.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● enabled_ssl_apis</li> </ul>
<b>ssl_key_file = None</b>	string value	<p>This option specifies the path to the file where SSL private key of API server is stored when SSL is in effect.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● String representing path to the SSL private key.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● enabled_ssl_apis</li> </ul>
<b>tcp_keepidle = 600</b>	integer value	<p>This option sets the value of TCP_KEEPIDLE in seconds for each server socket. It specifies the duration of time to keep connection active. TCP generates a KEEPALIVE transmission for an application that requests to keep connection active. Not supported on OS X.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● keep_alive</li> </ul>

Configuration option = Default value	Type	Description
<b>wsgi_log_format = % (client_ip)s "% (request_line)s" status: % (status_code)s len: % (body_length)s time: % (wall_seconds).7f</b>	string value	<p>It represents a python format string that is used as the template to generate log lines. The following values can be formatted into it: <code>client_ip</code>, <code>date_time</code>, <code>request_line</code>, <code>status_code</code>, <code>body_length</code>, <code>wall_seconds</code>.</p> <p>This option is used for building custom request loglines when running nova-api under eventlet. If used under uwsgi or apache, this option has no effect.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>• <code>%(client_ip)s "%(request_line)s" status: %(status_code)s len: %(body_length)s time: %(wall_seconds).7f</code> (default)</li> <li>• Any formatted string formed by specific values.</li> </ul>

### 11.1.57. xenserver

The following table outlines the options available under the **[xenserver]** group in the `/etc/nova/nova.conf` file.

Table 11.57. xenserver

Configuration option = Default value	Type	Description
<b>agent_path = usr/sbin/xen- update-networking</b>	string value	<p>Path to locate guest agent on the server.</p> <p>Specifies the path in which the XenAPI guest agent should be located. If the agent is present, network configuration is not injected into the image.</p> <p>Related options:</p> <p>For this option to have an effect: * <b>flat_injected</b> should be set to <b>True</b> * <b>compute_driver</b> should be set to <b>xenapi.XenAPIDriver</b></p>
<b>agent_resetnetwork_time out = 60</b>	integer value	<p>Number of seconds to wait for agent's reply to resetnetwork request.</p> <p>This indicates the amount of time xapi agent plugin waits for the agent to respond to the <code>resetnetwork</code> request specifically. The generic timeout for agent communication <b>agent_timeout</b> is ignored in this case.</p>

Configuration option = Default value	Type	Description
<b>agent_timeout = 30</b>	integer value	<p>Number of seconds to wait for agent's reply to a request.</p> <p>Nova configures/performs certain administrative actions on a server with the help of an agent that's installed on the server. The communication between Nova and the agent is achieved via sharing messages, called records, over xenstore, a shared storage across all the domains on a Xenserver host. Operations performed by the agent on behalf of nova are: <i>version</i>, <i>key_init</i>, <i>password</i>, <i>resetnetwork</i>, <i>inject_file</i>, and <i>agentupdate</i>.</p> <p>To perform one of the above operations, the <i>xapi agent</i> plugin writes the command and its associated parameters to a certain location known to the domain and awaits response. On being notified of the message, the agent performs appropriate actions on the server and writes the result back to xenstore. This result is then read by the <i>xapi agent</i> plugin to determine the success/failure of the operation.</p> <p>This config option determines how long the <i>xapi agent</i> plugin shall wait to read the response off of xenstore for a given request/command. If the agent on the instance fails to write the result in this time period, the operation is considered to have timed out.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>agent_version_timeout</b></li> <li>● <b>agent_resetnetwork_timeout</b></li> </ul>
<b>agent_version_timeout = 300</b>	integer value	<p>Number of seconds to wait for agent's reply to version request.</p> <p>This indicates the amount of time <i>xapi agent</i> plugin waits for the agent to respond to the <i>version</i> request specifically. The generic timeout for agent communication <b>agent_timeout</b> is ignored in this case.</p> <p>During the build process the <i>version</i> request is used to determine if the agent is available/operational to perform other requests such as <i>resetnetwork</i>, <i>password</i>, <i>key_init</i> and <i>inject_file</i>. If the <i>version</i> call fails, the other configuration is skipped. So, this configuration option can also be interpreted as time in which agent is expected to be fully operational.</p>

Configuration option = Default value	Type	Description
<b>block_device_creation_timeout = 10</b>	integer value	Time in secs to wait for a block device to be created
<b>cache_images = all</b>	string value	<p>Cache glance images locally.</p> <p>The value for this option must be chosen from the choices listed here. Configuring a value other than these will default to <i>all</i>.</p> <p>Note: There is nothing that deletes these images.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● <b>all</b>: will cache all images.</li> <li>● <b>some</b>: will only cache images that have the image_property <b>cache_in_nova=True</b>.</li> <li>● <b>none</b>: turns off caching entirely.</li> </ul>
<b>check_host = True</b>	boolean value	<p>Ensure compute service is running on host XenAPI connects to. This option must be set to false if the <i>independent_compute</i> option is set to true.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Setting this option to true will make sure that compute service is running on the same host that is specified by <i>connection_url</i>.</li> <li>● Setting this option to false, doesn't perform the check.</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>independent_compute</b></li> </ul>
<b>connection_concurrent = 5</b>	integer value	<p>Maximum number of concurrent XenAPI connections.</p> <p>In nova, multiple XenAPI requests can happen at a time. Configuring this option will parallelize access to the XenAPI session, which allows you to make concurrent XenAPI connections.</p>
<b>connection_password = None</b>	string value	Password for connection to XenServer/Xen Cloud Platform

Configuration option = Default value	Type	Description
<b>connection_url = None</b>	string value	<p>URL for connection to XenServer/Xen Cloud Platform. A special value of <code>unix://local</code> can be used to connect to the local unix socket.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any string that represents a URL. The <code>connection_url</code> is generally the management network IP address of the XenServer.</li> <li>This option must be set if you chose the XenServer driver.</li> </ul>
<b>connection_username = root</b>	string value	Username for connection to XenServer/Xen Cloud Platform
<b>console_public_hostname = ccsosp-undercloud.usersys.redhat.com</b>	string value	<p>Publicly visible name for this console host.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Current hostname (default) or any string representing hostname.</li> </ul>
<b>default_os_type = linux</b>	string value	Default OS type used when uploading an image to glance
<b>disable_agent = False</b>	boolean value	<p>Disables the use of XenAPI agent.</p> <p>This configuration option suggests whether the use of agent should be enabled or not regardless of what image properties are present. Image properties have an effect only when this is set to <b>True</b>. Read description of config option <b>use_agent_default</b> for more information.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li><b>use_agent_default</b></li> </ul>



Configuration option = Default value	Type	Description
<b>image_compression_level = None</b>	integer value	<p>Compression level for images.</p> <p>By setting this option we can configure the gzip compression level. This option sets GZIP environment variable before spawning tar -cz to force the compression level. It defaults to none, which means the GZIP environment variable is not set and the default (usually -6) is used.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Range is 1-9, e.g., 9 for gzip -9, 9 being most compressed but most CPU intensive on dom0.</li> <li>● Any values out of this range will default to None.</li> </ul>
<b>image_upload_handler = nova.virt.xenapi.image.glance.GlanceStore</b>	string value	<p>Dom0 plugin driver used to handle image uploads.</p> <p>Provide a string value representing a plugin driver required to handle the image uploading to GlanceStore.</p> <p>Images, and snapshots from XenServer need to be uploaded to the data store for use. image_upload_handler takes in a value for the Dom0 plugin driver. This driver is then called to upload images to the GlanceStore.</p>
<b>independent_compute = False</b>	boolean value	<p>Used to prevent attempts to attach VBDs locally, so Nova can be run in a VM on a different host.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>CONF.flat_injected</b> (Must be False)</li> <li>● <b>CONF.xenserver.check_host</b> (Must be False)</li> <li>● <b>CONF.default_ephemeral_format</b> (Must be unset or ext3)</li> <li>● Joining host aggregates (will error if attempted)</li> <li>● Swap disks for Windows VMs (will error if attempted)</li> <li>● Nova-based auto_configure_disk (will error if attempted)</li> </ul>

Configuration option = Default value	Type	Description
<b>introduce_vdi_retry_wait = 20</b>	integer value	<p>Number of seconds to wait for SR to settle if the VDI does not exist when first introduced.</p> <p>Some SRs, particularly iSCSI connections are slow to see the VDIs right after they got introduced. Setting this option to a time interval will make the SR to wait for that time period before raising VDI not found exception.</p>
<b>ipxe_boot_menu_url = None</b>	string value	<p>URL to the iPXE boot menu.</p> <p>An iPXE ISO is a specially crafted ISO which supports iPXE booting. This feature gives a means to roll your own image.</p> <p>By default this option is not set. Enable this option to boot an iPXE ISO.</p> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● <b>ipxe_network_name</b></li> <li>● <b>ipxe_mkisofs_cmd</b></li> </ul>
<b>ipxe_mkisofs_cmd = mkisofs</b>	string value	<p>Name and optionally path of the tool used for ISO image creation.</p> <p>An iPXE ISO is a specially crafted ISO which supports iPXE booting. This feature gives a means to roll your own image.</p> <p>Note: By default <b>mkisofs</b> is not present in the Dom0, so the package can either be manually added to Dom0 or include the <b>mkisofs</b> binary in the image itself.</p> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● <b>ipxe_network_name</b></li> <li>● <b>ipxe_boot_menu_url</b></li> </ul>

Configuration option = Default value	Type	Description
<b>ipxe_network_name = None</b>	string value	<p>Name of network to use for booting iPXE ISOs.</p> <p>An iPXE ISO is a specially crafted ISO which supports iPXE booting. This feature gives a means to roll your own image.</p> <p>By default this option is not set. Enable this option to boot an iPXE ISO.</p> <p>Related Options:</p> <ul style="list-style-type: none"> <li>● <b>ipxe_boot_menu_url</b></li> <li>● <b>ipxe_mkisofs_cmd</b></li> </ul>
<b>login_timeout = 10</b>	integer value	Timeout in seconds for XenAPI login.
<b>max_kernel_ramdisk_size = 16777216</b>	integer value	<p>Maximum size in bytes of kernel or ramdisk images.</p> <p>Specifying the maximum size of kernel or ramdisk will avoid copying large files to dom0 and fill up /boot/guest.</p>
<b>num_vbd_unplug_retries = 10</b>	integer value	Maximum number of retries to unplug VBD. If set to 0, should try once, no retries.
<b>ovs_integration_bridge = None</b>	string value	<p>The name of the integration Bridge that is used with xenapi when connecting with Open vSwitch.</p> <p>Note: The value of this config option is dependent on the environment, therefore this configuration value must be set accordingly if you are using XenAPI.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● Any string that represents a bridge name.</li> </ul>

Configuration option = Default value	Type	Description
<b>running_timeout = 60</b>	integer value	<p>Wait time for instances to go to running state.</p> <p>Provide an integer value representing time in seconds to set the wait time for an instance to go to running state.</p> <p>When a request to create an instance is received by nova-api and communicated to nova-compute, the creation of the instance occurs through interaction with Xen via XenAPI in the compute node. Once the node on which the instance(s) are to be launched is decided by nova-schedule and the launch is triggered, a certain amount of wait time is involved until the instance(s) can become available and <i>running</i>. This wait time is defined by <code>running_timeout</code>. If the instances do not go to running state within this specified wait time, the launch expires and the instance(s) are set to <i>error</i> state.</p>
<b>sparse_copy = True</b>	boolean value	Whether to use <code>sparse_copy</code> for copying data on a resize down. (False will use standard <code>dd</code> ). This speeds up resizes down considerably since large runs of zeros won't have to be rsynced.
<b>sr_base_path = /var/run/sr-mount</b>	string value	Base path to the storage repository on the XenServer host.
<b>sr_matching_filter = default-sr:true</b>	string value	<p>Filter for finding the SR to be used to install guest instances on.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● To use the Local Storage in default XenServer/XCP installations set this flag to <code>other-config:i18n-key=local-storage</code>.</li> <li>● To select an SR with a different matching criteria, you could set it to <code>other-config:my_favorite_sr=true</code>.</li> <li>● To fall back on the Default SR, as displayed by XenCenter, set this flag to: <code>default-sr:true</code>.</li> </ul>

Configuration option = Default value	Type	Description
<b>target_host = None</b>	unknown value	<p>The iSCSI Target Host.</p> <p>This option represents the hostname or ip of the iSCSI Target. If the target host is not present in the connection information from the volume provider then the value from this option is taken.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>Any string that represents hostname/ip of Target.</li> </ul>
<b>target_port = 3260</b>	port value	<p>The iSCSI Target Port.</p> <p>This option represents the port of the iSCSI Target. If the target port is not present in the connection information from the volume provider then the value from this option is taken.</p>
<b>use_agent_default = False</b>	boolean value	<p>Whether or not to use the agent by default when its usage is enabled but not indicated by the image.</p> <p>The use of XenAPI agent can be disabled altogether using the configuration option <b>disable_agent</b>. However, if it is not disabled, the use of an agent can still be controlled by the image in use through one of its properties, <b>xenapi_use_agent</b>. If this property is either not present or specified incorrectly on the image, the use of agent is determined by this configuration option.</p> <p>Note that if this configuration is set to <b>True</b> when the agent is not present, the boot times will increase significantly.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li><b>disable_agent</b></li> </ul>

Configuration option = Default value	Type	Description
<b>use_join_force = True</b>	boolean value	<p>When adding new host to a pool, this will append a --force flag to the command, forcing hosts to join a pool, even if they have different CPUs.</p> <p>Since XenServer version 5.6 it is possible to create a pool of hosts that have different CPU capabilities. To accommodate CPU differences, XenServer limited features it uses to determine CPU compatibility to only the ones that are exposed by CPU and support for CPU masking was added. Despite this effort to level differences between CPUs, it is still possible that adding new host will fail, thus option to force join was introduced.</p>
<b>vhd_coalesce_max_attempts = 20</b>	integer value	<p>Max number of times to poll for VHD to coalesce.</p> <p>This option determines the maximum number of attempts that can be made for coalescing the VHD before giving up.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>vhd_coalesce_poll_interval</b></li> </ul>
<b>vhd_coalesce_poll_interval = 5.0</b>	floating point value	<p>The interval used for polling of coalescing vhds.</p> <p>This is the interval after which the task of coalesce VHD is performed, until it reaches the max attempts that is set by vhd_coalesce_max_attempts.</p> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>vhd_coalesce_max_attempts</b></li> </ul>

Configuration option = Default value	Type	Description
<b>vif_driver = nova.virt.xenapi.vif.XenAPIOpenVswitchDriver</b>	string value	<p>The XenAPI VIF driver using XenServer Network APIs.</p> <p>Provide a string value representing the VIF XenAPI vif driver to use for plugging virtual network interfaces.</p> <p>Xen configuration uses bridging within the backend domain to allow all VMs to appear on the network as individual hosts. Bridge interfaces are used to create a XenServer VLAN network in which the VIFs for the VM instances are plugged. If no VIF bridge driver is plugged, the bridge is not made available. This configuration option takes in a value for the VIF driver.</p> <p>Possible values:</p> <ul style="list-style-type: none"> <li>● <code>nova.virt.xenapi.vif.XenAPIOpenVswitchDriver</code> (default)</li> <li>● <code>nova.virt.xenapi.vif.XenAPIBridgeDriver</code> (deprecated)</li> </ul> <p>Related options:</p> <ul style="list-style-type: none"> <li>● <b>vlan_interface</b></li> <li>● <b>ovs_integration_bridge</b></li> </ul>

### 11.1.58. xvp

The following table outlines the options available under the **[xvp]** group in the `/etc/nova/nova.conf` file.

Table 11.58. xvp

Configuration option = Default value	Type	Description
<b>console_xvp_conf = /etc/xvp.conf</b>	string value	Generated XVP conf file
<b>console_xvp_conf_template = \$pybasedir/nova/console/xvp.conf.template</b>	string value	XVP conf template
<b>console_xvp_log = /var/log/xvp.log</b>	string value	XVP log file

Configuration option = Default value	Type	Description
<b>console_xvp_multiplex_port = 5900</b>	port value	Port for XVP to multiplex VNC connections on
<b>console_xvp_pid = /var/run/xvp.pid</b>	string value	XVP master process pid file



## CHAPTER 12. SAHARA

The following chapter contains information about the configuration options in the **sahara** service.

### 12.1. SAHARA.CONF

This section contains options for the `/etc/sahara/sahara.conf` file.

#### 12.1.1. DEFAULT

The following table outlines the options available under the **[DEFAULT]** group in the `/etc/sahara/sahara.conf` file.

Table 12.1. DEFAULT

Configuration option = Default value	Type	Description
<b>api_paste_config = api-paste.ini</b>	string value	File name for the paste.deploy config for api service
<b>api_workers = 1</b>	integer value	Number of workers for Sahara API service (0 means all-in-one-thread configuration).
<b>cleanup_time_for_incomplete_clusters = 0</b>	integer value	Maximal time (in hours) for clusters allowed to be in states other than "Active", "Deleting" or "Error". If a cluster is not in "Active", "Deleting" or "Error" state and last update of it was longer than "cleanup_time_for_incomplete_clusters" hours ago then it will be deleted automatically. (0 value means that automatic clean up is disabled).
<b>client_socket_timeout = 900</b>	integer value	Timeout for client connections' socket operations. If an incoming connection is idle for this number of seconds it will be closed. A value of 0 means wait forever.
<b>cluster_remote_threshold = 70</b>	integer value	The same as global_remote_threshold, but for a single cluster.
<b>compute_topology_file = etc/sahara/compute.topology</b>	string value	File with nova compute topology. It should contain mapping between nova computes and racks.
<b>conn_pool_min_size = 2</b>	integer value	The pool size limit for connections expiration policy
<b>conn_pool_ttl = 1200</b>	integer value	The time-to-live in sec of idle connections in the pool
<b>control_exchange = openstack</b>	string value	The default exchange under which topics are scoped. May be overridden by an exchange name specified in the transport_url option.

Configuration option = Default value	Type	Description
<b>coordinator_heartbeat_interval = 1</b>	integer value	Interval size between heartbeat execution in seconds. Heartbeats are executed to make sure that connection to the coordination server is active.
<b>data_source_types = [u'swift', u'hdfs', u'maprfs', u'manila']</b>	list value	List of data sources types to be loaded. Sahara preserves the order of the list when returning it.
<b>db_driver = sahara.db</b>	string value	Driver to use for database access.
<b>debug = False</b>	boolean value	If set to true, the logging level will be set to DEBUG instead of the default INFO level.
<b>default_log_levels = [u'amqp=WARN', u'amqplib=WARN', u'boto=WARN', u'qpid=WARN', u'sqlalchemy=WARN', u'suds=INFO', u'oslo.messaging=INFO', u'oslo_messaging=INFO', u'iso8601=WARN', u'requests.packages.urllib3.connectionpool=WARN', u'urllib3.connectionpool=WARN', u'websocket=WARN', u'requests.packages.urllib3.util.retry=WARN', u'urllib3.util.retry=WARN', u'keystonemiddleware=WARN', u'routes.middleware=WARN', u'stevedore=WARN', u'taskflow=WARN', u'keystoneauth=WARN', u'oslo.cache=INFO', u'dogpile.core.dogpile=INFO', u'stevedore=INFO', u'eventlet.wsgi.server=WARN', u'paramiko=WARN', u'requests=WARN', u'neutronclient=INFO']</b>	list value	List of package logging levels in logger=LEVEL pairs. This option is ignored if log_config_append is set.
<b>disable_event_log = False</b>	boolean value	Disables event log feature.

Configuration option = Default value	Type	Description
<b>edp_internal_db_enabled = True</b>	boolean value	Use Sahara internal db to store job binaries.
<b>enable_data_locality = False</b>	boolean value	Enables data locality for hadoop cluster. Also enables data locality for Swift used by hadoop. If enabled, <i>compute_topology</i> and <i>swift_topology</i> configuration parameters should point to OpenStack and Swift topology correspondingly.
<b>enable_hypervisor_awareness = True</b>	boolean value	Enables four-level topology for data locality. Works only if corresponding plugin supports such mode.
<b>executor_thread_pool_size = 64</b>	integer value	Size of executor thread pool when executor is threading or eventlet.
<b>fatal_deprecations = False</b>	boolean value	Enables or disables fatal status of deprecations.
<b>global_remote_threshold = 100</b>	integer value	Maximum number of remote operations that will be running at the same time. Note that each remote operation requires its own process to run.
<b>hash_ring_replicas_count = 40</b>	integer value	Number of points that belongs to each member on a hash ring. The larger number leads to a better distribution.
<b>heat_enable_wait_condition = True</b>	boolean value	Enable wait condition feature to reduce polling during cluster creation
<b>heat_stack_tags = [u'data-processing-cluster']</b>	list value	List of tags to be used during operating with stack.
<b>host = 0.0.0.0</b>	unknown value	Hostname or IP address that will be used to listen on.
<code>`instance_format = [instance: %(uuid)s] `</code>	string value	The format for an instance that is passed with the log message.
<code>`instance_uuid_format = [instance: %(uuid)s] `</code>	string value	The format for an instance UUID that is passed with the log message.
<b>job_binary_max_KB = 5120</b>	integer value	Maximum length of job binary data in kilobytes that may be stored or retrieved in a single operation.

Configuration option = Default value	Type	Description
<b>job_binary_types =</b> [u'swift', u'manila', u'internal-db', u's3']	list value	List of job binary types to be loaded. Sahara preserves the order of the list when returning it.
<b>job_canceling_timeout =</b> 300	integer value	Timeout for canceling job execution (in seconds). Sahara will try to cancel job execution during this time.
<code>`job_workflow_postfix = `</code>	string value	Postfix for storing jobs in hdfs. Will be added to <code>/user/&lt;hdfs user&gt;/</code> path.
<b>log-config-append = None</b>	string value	The name of a logging configuration file. This file is appended to any existing logging configuration files. For details about logging configuration files, see the Python logging module documentation. Note that when logging configuration files are used then all logging configuration is set in the configuration file and other logging configuration options are ignored (for example, <code>logging_context_format_string</code> ).
<b>log-date-format = %Y-%m-%d %H:%M:%S</b>	string value	Defines the format string for <code>%(asctime)s</code> in log records. Default: <code>%(default)s</code> . This option is ignored if <code>log_config_append</code> is set.
<b>log-dir = None</b>	string value	(Optional) The base directory used for relative <code>log_file</code> paths. This option is ignored if <code>log_config_append</code> is set.
<b>log-exchange = False</b>	boolean value	Log request/response exchange details: environ, headers and bodies.
<b>log-file = None</b>	string value	(Optional) Name of log file to send logging output to. If no default is set, logging will go to <code>stderr</code> as defined by <code>use_stderr</code> . This option is ignored if <code>log_config_append</code> is set.
<b>logging_context_format_string = %(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [%(request_id)s] %(user_identity)s %(instance)s%(message)s</b>	string value	Format string to use for log messages with context.
<b>logging_debug_format_s uffix = %(funcName)s % (pathname)s:%(lineno)d</b>	string value	Additional data to append to log message when logging level for the message is <code>DEBUG</code> .

Configuration option = Default value	Type	Description
<b>logging_default_format_string =</b> <code>%(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [-] %(instance)s%(message)s</code>	string value	Format string to use for log messages when context is undefined.
<b>logging_exception_prefix =</b> <code>%(asctime)s.%(msecs)03d %(process)d ERROR %(name)s %(instance)s</code>	string value	Prefix each line of exception output with this format.
<b>logging_user_identity_format =</b> <code>%(user)s %(tenant)s %(domain)s %(user_domain)s %(project_domain)s</code>	string value	Defines the format string for <code>%(user_identity)s</code> that is used in <code>logging_context_format_string</code> .
<b>max_header_line = 16384</b>	integer value	Maximum line size of message headers to be accepted. <code>max_header_line</code> may need to be increased when using large tokens (typically those generated when keystone is configured to use PKI tokens with big service catalogs).
<b>min_transient_cluster_active_time = 30</b>	integer value	Minimal "lifetime" in seconds for a transient cluster. Cluster is guaranteed to be "alive" within this time period.
<b>nameservers = []</b>	list value	IP addresses of Designate nameservers. This is required if <code>use_designate</code> is True
<b>node_domain = novalocal</b>	string value	The suffix of the node's FQDN.
<b>os_region_name = None</b>	string value	Region name used to get services endpoints.
<b>periodic_coordinator_backend_url = None</b>	string value	The backend URL to use for distributed periodic tasks coordination.
<b>periodic_enable = True</b>	boolean value	Enable periodic tasks.
<b>periodic_fuzzy_delay = 60</b>	integer value	Range in seconds to randomly delay when starting the periodic task scheduler to reduce stampeding. (Disable by setting to 0).
<b>periodic_interval_max = 60</b>	integer value	Max interval size between periodic tasks execution in seconds.

Configuration option = Default value	Type	Description
<b>periodic_workers_number = 1</b>	integer value	Number of threads to run periodic tasks.
<b>plugins = [u'vanilla', u'spark', u'cdh', u'ambari', u'storm', u'mapr']</b>	list value	List of plugins to be loaded. Sahara preserves the order of the list when returning it.
<b>port = 8386</b>	port value	Port that will be used to listen on.
<code>`proxy_command = `</code>	string value	Proxy command used to connect to instances. If set, this command should open a netcat socket, that Sahara will use for SSH and HTTP connections. Use {host} and {port} to describe the destination. Other available keywords: {tenant_id}, {network_id}, {router_id}.
<b>proxy_command_use_internal_ip = False</b>	boolean value	Force proxy_command usage to be consuming internal IP always, instead of management IP. Ignored if proxy_command is not set.
<b>proxy_user_domain_name = None</b>	string value	The domain Sahara will use to create new proxy users for Swift object access.
<b>proxy_user_role_names = [u'Member']</b>	list value	A list of the role names that the proxy user should assume through trust for Swift object access.
<b>publish_errors = False</b>	boolean value	Enables or disables publication of error events.
<b>rate_limit_burst = 0</b>	integer value	Maximum number of logged messages per rate_limit_interval.
<b>rate_limit_except_level = CRITICAL</b>	string value	Log level name used by rate limiting: CRITICAL, ERROR, INFO, WARNING, DEBUG or empty string. Logs with level greater or equal to rate_limit_except_level are not filtered. An empty string means that all levels are filtered.
<b>rate_limit_interval = 0</b>	integer value	Interval, number of seconds, of log rate limiting.
<b>remote = ssh</b>	string value	A method for Sahara to execute commands on VMs.
<b>retries_number = 5</b>	integer value	Number of times to retry the request to client before failing
<b>retry_after = 10</b>	integer value	Time between the retries to client (in seconds).

Configuration option = Default value	Type	Description
<b>rootwrap_command = sudo sahara-rootwrap /etc/sahara/rootwrap.conf</b>	string value	Rootwrap command to leverage. Use in conjunction with use_rootwrap=True
<b>rpc_ack_timeout_base = 15</b>	integer value	Number of seconds to wait for an ack from a cast/call. After each retry attempt this timeout is multiplied by some specified multiplier.
<b>rpc_ack_timeout_multipli er = 2</b>	integer value	Number to multiply base ack timeout by after each retry attempt.
<b>rpc_backend = rabbit</b>	string value	The messaging driver to use, defaults to rabbit. Other drivers include amqp and zmq.
<b>rpc_conn_pool_size = 30</b>	integer value	Size of RPC connection pool.
<b>rpc_message_ttl = 300</b>	integer value	Expiration timeout in seconds of a sent/received message after which it is not tracked anymore by a client/server.
<b>rpc_poll_timeout = 1</b>	integer value	The default number of seconds that poll should wait. Poll raises timeout exception when timeout expired.
<b>rpc_response_timeout = 60</b>	integer value	Seconds to wait for a response from a call.
<b>rpc_retry_attempts = 3</b>	integer value	Default number of message sending attempts in case of any problems occurred: positive value N means at most N retries, 0 means no retries, None or -1 (or any other negative values) mean to retry forever. This option is used only if acknowledgments are enabled.
<b>rpc_thread_pool_size = 100</b>	integer value	Maximum number of (green) threads to work concurrently.
<b>rpc_use_acks = False</b>	boolean value	Wait for message acknowledgements from receivers. This mechanism works only via proxy without PUB/SUB.
<b>rpc_zmq_bind_address = *</b>	string value	ZeroMQ bind address. Should be a wildcard (*), an ethernet interface, or IP. The "host" option should point or resolve to this address.
<b>rpc_zmq_bind_port_retrie s = 100</b>	integer value	Number of retries to find free port number before fail with ZMQBindError.

Configuration option = Default value	Type	Description
<b>rpc_zmq_contexts = 1</b>	integer value	Number of ZeroMQ contexts, defaults to 1.
<b>rpc_zmq_host = ccsosp-undercloud.usersys.redhat.com</b>	string value	Name of this node. Must be a valid hostname, FQDN, or IP address. Must match "host" option, if running Nova.
<b>rpc_zmq_ipc_dir = /var/run/openstack</b>	string value	Directory for holding IPC sockets.
<b>rpc_zmq_matchmaker = redis</b>	string value	MatchMaker driver.
<b>rpc_zmq_max_port = 65536</b>	integer value	Maximal port number for random ports range.
<b>rpc_zmq_min_port = 49153</b>	port value	Minimal port number for random ports range.
<b>rpc_zmq_serialization = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>rpc_zmq_topic_backlog = None</b>	integer value	Maximum number of ingress messages to locally buffer per topic. Default is unlimited.
<b>run_external_periodic_tasks = True</b>	boolean value	Some periodic tasks can be run in a separate process. Should we run them here?
<b>ssh_timeout_common = 300</b>	integer value	Overrides timeout for common ssh operations, in seconds
<b>ssh_timeout_files = 120</b>	integer value	Overrides timeout for ssh operations with files, in seconds
<b>ssh_timeout_interactive = 1800</b>	integer value	Overrides timeout for interactive ssh operations, in seconds
<b>subscribe_on = []</b>	list value	List of publisher hosts SubConsumer can subscribe on. This option has higher priority then the default publishers list taken from the matchmaker.
<b>swift_topology_file = etc/sahara/swift.topology</b>	string value	File with Swift topology.It should contain mapping between Swift nodes and racks.
<b>syslog-log-facility = LOG_USER</b>	string value	Syslog facility to receive log lines. This option is ignored if log_config_append is set.



Configuration option = Default value	Type	Description
<b>tcp_keepidle = 600</b>	integer value	Sets the value of TCP_KEEPIDLE in seconds for each server socket. Not supported on OS X.
<b>transport_url = None</b>	string value	<p>The network address and optional user credentials for connecting to the messaging backend, in URL format. The expected format is:</p> <pre>driver://[user:pass@]host:port[, [userN:passN@]hostN:portN]/virtual_host?query</pre> <p>Example: rabbit://rabbitmq:password@127.0.0.1:5672//</p> <p>For full details on the fields in the URL see the documentation of oslo_messaging.TransportURL at <a href="https://docs.openstack.org/oslo.messaging/latest/reference/transport.html">https://docs.openstack.org/oslo.messaging/latest/reference/transport.html</a></p>
<b>use-journal = False</b>	boolean value	Enable journald for logging. If running in a systemd environment you may wish to enable journal support. Doing so will use the journal native protocol which includes structured metadata in addition to log messages. This option is ignored if log_config_append is set.
<b>use-json = False</b>	boolean value	Use JSON formatting for logging. This option is ignored if log_config_append is set.
<b>use-syslog = False</b>	boolean value	Use syslog for logging. Existing syslog format is DEPRECATED and will be changed later to honor RFC5424. This option is ignored if log_config_append is set.
<b>use_barbican_key_manager = False</b>	boolean value	Enable the usage of the OpenStack Key Management service provided by barbican.
<b>use_designate = False</b>	boolean value	Use Designate for internal and external hostnames resolution
<b>use_domain_for_proxy_users = False</b>	boolean value	Enables Sahara to use a domain for creating temporary proxy users to access Swift. If this is enabled a domain must be created for Sahara to use.
<b>use_dynamic_connections = False</b>	boolean value	This option makes direct connections dynamic or static. It makes sense only with use_router_proxy=False which means to use direct connections for direct message types (ignored otherwise).

Configuration option = Default value	Type	Description
<b>use_floating_ips = True</b>	boolean value	If set to True, Sahara will use floating IPs to communicate with instances. To make sure that all instances have floating IPs assigned, make sure that all Node Groups have "floating_ip_pool" parameter defined.
<b>use_identity_api_v3 = True</b>	boolean value	Enables Sahara to use Keystone API v3. If that flag is disabled, per-job clusters will not be terminated automatically.
<b>use_namespaces = False</b>	boolean value	Use network namespaces for communication.
<b>use_pub_sub = False</b>	boolean value	Use PUB/SUB pattern for fanout methods. PUB/SUB always uses proxy.
<b>use_rootwrap = False</b>	boolean value	Use rootwrap facility to allow non-root users to run the sahara services and access private network IPs (only valid to use in conjunction with use_namespaces=True)
<b>use_router_proxy = False</b>	boolean value	Use ROUTER remote proxy.
<b>use_stderr = False</b>	boolean value	Log output to standard error. This option is ignored if log_config_append is set.
<b>watch-log-file = False</b>	boolean value	Uses logging handler designed to watch file system. When log file is moved or removed this handler will open a new log file with specified path instantaneously. It makes sense only if log_file option is specified and Linux platform is used. This option is ignored if log_config_append is set.
<b>wsgi_default_pool_size = 100</b>	integer value	Size of the pool of greenthreads used by wsgi
<b>wsgi_keep_alive = True</b>	boolean value	If False, closes the client socket connection explicitly.
<b>wsgi_log_format = %(client_ip)s "%(request_line)s" status: %(status_code)s len: %(body_length)s time: %(wall_seconds).7f</b>	string value	A python format string that is used as the template to generate log lines. The following values can be formatted into it: client_ip, date_time, request_line, status_code, body_length, wall_seconds.

Configuration option = Default value	Type	Description
<b>zmq_failover_connections = 2</b>	integer value	How many additional connections to a host will be made for failover reasons. This option is actual only in dynamic connections mode.
<b>zmq_immediate = True</b>	boolean value	This option configures round-robin mode in zmq socket. True means not keeping a queue when server side disconnects. False means to keep queue and messages even if server is disconnected, when the server appears we send all accumulated messages to it.
<b>zmq_linger = -1</b>	integer value	Number of seconds to wait before all pending messages will be sent after closing a socket. The default value of -1 specifies an infinite linger period. The value of 0 specifies no linger period. Pending messages shall be discarded immediately when the socket is closed. Positive values specify an upper bound for the linger period.
<b>zmq_target_expire = 300</b>	integer value	Expiration timeout in seconds of a name service record about existing target (< 0 means no timeout).
<b>zmq_target_update = 180</b>	integer value	Update period in seconds of a name service record about existing target.
<b>zmq_tcp_keepalive = -1</b>	integer value	Enable/disable TCP keepalive (KA) mechanism. The default value of -1 (or any other negative value) means to skip any overrides and leave it to OS default; 0 and 1 (or any other positive value) mean to disable and enable the option respectively.
<b>zmq_tcp_keepalive_cnt = -1</b>	integer value	The number of retransmissions to be carried out before declaring that remote end is not available. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_idle = -1</b>	integer value	The duration between two keepalive transmissions in idle condition. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.

Configuration option = Default value	Type	Description
<b>zmq_tcp_keepalive_intvl = -1</b>	integer value	The duration between two successive keepalive retransmissions, if acknowledgement to the previous keepalive transmission is not received. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.

### 12.1.2. castellan

The following table outlines the options available under the **[castellan]** group in the **/etc/sahara/sahara.conf** file.

Table 12.2. castellan

Configuration option = Default value	Type	Description
<b>barbican_api_endpoint = None</b>	string value	The endpoint to use for connecting to the barbican api controller. By default, castellan will use the URL from the service catalog.
<b>barbican_api_version = v1</b>	string value	Version of the barbican API, for example: "v1"

### 12.1.3. cinder

The following table outlines the options available under the **[cinder]** group in the **/etc/sahara/sahara.conf** file.

Table 12.3. cinder

Configuration option = Default value	Type	Description
<b>api_insecure = False</b>	boolean value	Allow to perform insecure SSL requests to cinder.
<b>api_version = 3</b>	integer value	Version of the Cinder API to use.
<b>ca_file = None</b>	string value	Location of ca certificates file to use for cinder client requests.
<b>endpoint_type = internalURL</b>	string value	Endpoint type for cinder client requests

### 12.1.4. conductor

The following table outlines the options available under the **[conductor]** group in the `/etc/sahara/sahara.conf` file.

Table 12.4. conductor

Configuration option = Default value	Type	Description
<b>use_local = True</b>	boolean value	Perform sahara-conductor operations locally.

### 12.1.5. cors

The following table outlines the options available under the **[cors]** group in the `/etc/sahara/sahara.conf` file.

Table 12.5. cors

Configuration option = Default value	Type	Description
<b>allow_credentials = True</b>	boolean value	Indicate that the actual request can include user credentials
<b>allow_headers = [u'X-Auth-Token', u'X-Identity-Status', u'X-Roles', u'X-Service-Catalog', u'X-User-Id', u'X-Tenant-Id', u'X-OpenStack-Request-ID']</b>	list value	Indicate which header field names may be used during the actual request.
<b>allow_methods = [u'GET', u'PUT', u'POST', u'DELETE', u'PATCH']</b>	list value	Indicate which methods can be used during the actual request.
<b>allowed_origin = None</b>	list value	Indicate whether this resource may be shared with the domain received in the requests "origin" header. Format: "<protocol>://<host>[:<port>]", no trailing slash. Example: <a href="https://horizon.example.com">https://horizon.example.com</a>
<b>expose_headers = [u'X-Auth-Token', u'X-Subject-Token', u'X-Service-Token', u'X-OpenStack-Request-ID']</b>	list value	Indicate which headers are safe to expose to the API. Defaults to HTTP Simple Headers.
<b>max_age = 3600</b>	integer value	Maximum cache age of CORS preflight requests.

### 12.1.6. database

The following table outlines the options available under the **[database]** group in the `/etc/sahara/sahara.conf` file.

Table 12.6. database

Configuration option = Default value	Type	Description
<b>backend = sqlalchemy</b>	string value	The back end to use for the database.
<b>connection = None</b>	string value	The SQLAlchemy connection string to use to connect to the database.
<b>connection_debug = 0</b>	integer value	Verbosity of SQL debugging information: 0=None, 100=Everything.
<b>connection_recycle_time = 3600</b>	integer value	Connections which have been present in the connection pool longer than this number of seconds will be replaced with a new one the next time they are checked out from the pool.
<b>connection_trace = False</b>	boolean value	Add Python stack traces to SQL as comment strings.
<b>db_inc_retry_interval = True</b>	boolean value	If True, increases the interval between retries of a database operation up to <code>db_max_retry_interval</code> .
<b>db_max_retries = 20</b>	integer value	Maximum retries in case of connection error or deadlock error before error is raised. Set to -1 to specify an infinite retry count.
<b>db_max_retry_interval = 10</b>	integer value	If <code>db_inc_retry_interval</code> is set, the maximum seconds between retries of a database operation.
<b>db_retry_interval = 1</b>	integer value	Seconds between retries of a database transaction.
<b>max_overflow = 50</b>	integer value	If set, use this value for <code>max_overflow</code> with SQLAlchemy.
<b>max_pool_size = 5</b>	integer value	Maximum number of SQL connections to keep open in a pool. Setting a value of 0 indicates no limit.
<b>max_retries = 10</b>	integer value	Maximum number of database connection retries during startup. Set to -1 to specify an infinite retry count.
<b>min_pool_size = 1</b>	integer value	Minimum number of SQL connections to keep open in a pool.
<b>mysql_enable_ndb = False</b>	boolean value	If True, transparently enables support for handling MySQL Cluster (NDB).

Configuration option = Default value	Type	Description
<b>mysql_sql_mode = TRADITIONAL</b>	string value	The SQL mode to be used for MySQL sessions. This option, including the default, overrides any server-set SQL mode. To use whatever SQL mode is set by the server configuration, set this to no value. Example: <code>mysql_sql_mode=</code>
<b>pool_timeout = None</b>	integer value	If set, use this value for <code>pool_timeout</code> with SQLAlchemy.
<b>retry_interval = 10</b>	integer value	Interval between retries of opening a SQL connection.
<b>slave_connection = None</b>	string value	The SQLAlchemy connection string to use to connect to the slave database.
<b>sqlite_synchronous = True</b>	boolean value	If True, SQLite uses synchronous mode.
<b>use_db_reconnect = False</b>	boolean value	Enable the experimental use of database reconnect on connection lost.

### 12.1.7. glance

The following table outlines the options available under the **[glance]** group in the `/etc/sahara/sahara.conf` file.

Table 12.7. glance

Configuration option = Default value	Type	Description
<b>api_insecure = False</b>	boolean value	Allow to perform insecure SSL requests to glance.
<b>ca_file = None</b>	string value	Location of ca certificates file to use for glance client requests.
<b>endpoint_type = internalURL</b>	string value	Endpoint type for glance client requests

### 12.1.8. heat

The following table outlines the options available under the **[heat]** group in the `/etc/sahara/sahara.conf` file.

Table 12.8. heat

Configuration option = Default value	Type	Description
<b>api_insecure = False</b>	boolean value	Allow to perform insecure SSL requests to heat.
<b>ca_file = None</b>	string value	Location of ca certificates file to use for heat client requests.
<b>endpoint_type = internalURL</b>	string value	Endpoint type for heat client requests

### 12.1.9. keystone

The following table outlines the options available under the **[keystone]** group in the `/etc/sahara/sahara.conf` file.

Table 12.9. keystone

Configuration option = Default value	Type	Description
<b>api_insecure = False</b>	boolean value	Allow to perform insecure SSL requests to keystone.
<b>ca_file = None</b>	string value	Location of ca certificates file to use for keystone client requests.
<b>endpoint_type = internalURL</b>	string value	Endpoint type for keystone client requests

### 12.1.10. keystone\_authtoken

The following table outlines the options available under the **[keystone\_authtoken]** group in the `/etc/sahara/sahara.conf` file.

Table 12.10. keystone\_authtoken

Configuration option = Default value	Type	Description
<b>auth_section = None</b>	string value	Config Section from which to load plugin specific options
<b>auth_type = None</b>	string value	Authentication type to load



Configuration option = Default value	Type	Description
<b>auth_uri = None</b>	string value	Complete "public" Identity API endpoint. This endpoint should not be an "admin" endpoint, as it should be accessible by all end users. Unauthenticated clients are redirected to this endpoint to authenticate. Although this endpoint should ideally be unversioned, client support in the wild varies. If you're using a versioned v2 endpoint here, then this should <b>not</b> be the same endpoint the service user utilizes for validating tokens, because normal end users may not be able to reach that endpoint. This option is deprecated in favor of <code>www_authenticate_uri</code> and will be removed in the S release.
<b>auth_version = None</b>	string value	API version of the admin Identity API endpoint.
<b>cache = None</b>	string value	Request environment key where the Swift cache object is stored. When <code>auth_token</code> middleware is deployed with a Swift cache, use this option to have the middleware share a caching backend with swift. Otherwise, use the <b>memcached_servers</b> option instead.
<b>cafile = None</b>	string value	A PEM encoded Certificate Authority to use when verifying HTTPs connections. Defaults to system CAs.
<b>certfile = None</b>	string value	Required if identity server requires client certificate
<b>check_revocations_for_cached = False</b>	boolean value	If true, the revocation list will be checked for cached tokens. This requires that PKI tokens are configured on the identity server.
<b>delay_auth_decision = False</b>	boolean value	Do not handle authorization requests within the middleware, but delegate the authorization decision to downstream WSGI components.
<b>enforce_token_bind = permissive</b>	string value	Used to control the use and type of token binding. Can be set to: "disabled" to not check token binding. "permissive" (default) to validate binding information if the bind type is of a form known to the server and ignore it if not. "strict" like "permissive" but if the bind type is unknown the token will be rejected. "required" any form of token binding is needed to be allowed. Finally the name of a binding method that must be present in tokens.

Configuration option = Default value	Type	Description
<b>hash_algorithms = [u'md5']</b>	list value	Hash algorithms to use for hashing PKI tokens. This may be a single algorithm or multiple. The algorithms are those supported by Python standard <code>hashlib.new()</code> . The hashes will be tried in the order given, so put the preferred one first for performance. The result of the first hash will be stored in the cache. This will typically be set to multiple values only while migrating from a less secure algorithm to a more secure one. Once all the old tokens are expired this option should be set to a single value for better performance.
<b>http_connect_timeout = None</b>	integer value	Request timeout value for communicating with Identity API server.
<b>http_request_max_retries = 3</b>	integer value	How many times are we trying to reconnect when communicating with Identity API Server.
<b>include_service_catalog = True</b>	boolean value	(Optional) Indicate whether to set the X-Service-Catalog header. If False, middleware will not ask for service catalog on token validation and will not set the X-Service-Catalog header.
<b>insecure = False</b>	boolean value	Verify HTTPS connections.
<b>keyfile = None</b>	string value	Required if identity server requires client certificate
<b>memcache_pool_connection_timeout = 10</b>	integer value	(Optional) Number of seconds that an operation will wait to get a memcached client connection from the pool.
<b>memcache_pool_dead_retry = 300</b>	integer value	(Optional) Number of seconds memcached server is considered dead before it is tried again.
<b>memcache_pool_maxsize = 10</b>	integer value	(Optional) Maximum total number of open connections to every memcached server.
<b>memcache_pool_socket_timeout = 3</b>	integer value	(Optional) Socket timeout in seconds for communicating with a memcached server.
<b>memcache_pool_unused_timeout = 60</b>	integer value	(Optional) Number of seconds a connection to memcached is held unused in the pool before it is closed.
<b>memcache_secret_key = None</b>	string value	(Optional, mandatory if <code>memcache_security_strategy</code> is defined) This string is used for key derivation.

Configuration option = Default value	Type	Description
<b>memcache_security_strategy = None</b>	string value	(Optional) If defined, indicate whether token data should be authenticated or authenticated and encrypted. If MAC, token data is authenticated (with HMAC) in the cache. If ENCRYPT, token data is encrypted and authenticated in the cache. If the value is not one of these options or empty, auth_token will raise an exception on initialization.
<b>memcache_use_advanced_pool = False</b>	boolean value	(Optional) Use the advanced (eventlet safe) memcached client pool. The advanced pool will only work under python 2.x.
<b>memcached_servers = None</b>	list value	Optionally specify a list of memcached server(s) to use for caching. If left undefined, tokens will instead be cached in-process.
<b>region_name = None</b>	string value	The region in which the identity server can be found.
<b>revocation_cache_time = 10</b>	integer value	Determines the frequency at which the list of revoked tokens is retrieved from the Identity service (in seconds). A high number of revocation events combined with a low cache duration may significantly reduce performance. Only valid for PKI tokens. This option has been deprecated in the Ocata release and will be removed in the P release.
<b>service_token_roles = [u'service']</b>	list value	A choice of roles that must be present in a service token. Service tokens are allowed to request that an expired token can be used and so this check should tightly control that only actual services should be sending this token. Roles here are applied as an ANY check so any role in this list must be present. For backwards compatibility reasons this currently only affects the allow_expired check.
<b>service_token_roles_required = False</b>	boolean value	For backwards compatibility reasons we must let valid service tokens pass that don't pass the service_token_roles check as valid. Setting this true will become the default in a future release and should be enabled if possible.
<b>signing_dir = None</b>	string value	Directory used to cache files related to PKI tokens. This option has been deprecated in the Ocata release and will be removed in the P release.
<b>token_cache_time = 300</b>	integer value	In order to prevent excessive effort spent validating tokens, the middleware caches previously-seen tokens for a configurable duration (in seconds). Set to -1 to disable caching completely.

Configuration option = Default value	Type	Description
<b>www_authenticate_uri = None</b>	string value	Complete "public" Identity API endpoint. This endpoint should not be an "admin" endpoint, as it should be accessible by all end users. Unauthenticated clients are redirected to this endpoint to authenticate. Although this endpoint should ideally be unversioned, client support in the wild varies. If you're using a versioned v2 endpoint here, then this should <b>not</b> be the same endpoint the service user utilizes for validating tokens, because normal end users may not be able to reach that endpoint.

### 12.1.11. manila

The following table outlines the options available under the **[manila]** group in the **/etc/sahara/sahara.conf** file.

Table 12.11. manila

Configuration option = Default value	Type	Description
<b>api_insecure = True</b>	boolean value	Allow to perform insecure SSL requests to manila.
<b>api_version = 1</b>	string value	Version of the manila API to use.
<b>ca_file = None</b>	string value	Location of ca certificates file to use for manila client requests.

### 12.1.12. matchmaker\_redis

The following table outlines the options available under the **[matchmaker\_redis]** group in the **/etc/sahara/sahara.conf** file.

Table 12.12. matchmaker\_redis

Configuration option = Default value	Type	Description
<b>check_timeout = 20000</b>	integer value	Time in ms to wait before the transaction is killed.
<b>host = 127.0.0.1</b>	string value	Host to locate redis.
<b>password = ``</b>	string value	Password for Redis server (optional).
<b>port = 6379</b>	port value	Use this port to connect to redis host.

Configuration option = Default value	Type	Description
<b>sentinel_group_name = oslo-messaging-zeromq</b>	string value	Redis replica set name.
<b>sentinel_hosts = []</b>	list value	List of Redis Sentinel hosts (fault tolerance mode), e.g., [host:port, host1:port ... ]
<b>socket_timeout = 10000</b>	integer value	Timeout in ms on blocking socket operations.
<b>wait_timeout = 2000</b>	integer value	Time in ms to wait between connection attempts.

### 12.1.13. neutron

The following table outlines the options available under the **[neutron]** group in the `/etc/sahara/sahara.conf` file.

Table 12.13. neutron

Configuration option = Default value	Type	Description
<b>api_insecure = False</b>	boolean value	Allow to perform insecure SSL requests to neutron.
<b>ca_file = None</b>	string value	Location of ca certificates file to use for neutron client requests.
<b>endpoint_type = internalURL</b>	string value	Endpoint type for neutron client requests

### 12.1.14. nova

The following table outlines the options available under the **[nova]** group in the `/etc/sahara/sahara.conf` file.

Table 12.14. nova

Configuration option = Default value	Type	Description
<b>api_insecure = False</b>	boolean value	Allow to perform insecure SSL requests to nova.
<b>ca_file = None</b>	string value	Location of ca certificates file to use for nova client requests.
<b>endpoint_type = internalURL</b>	string value	Endpoint type for nova client requests

### 12.1.15. object\_store\_access

The following table outlines the options available under the **[object\_store\_access]** group in the `/etc/sahara/sahara.conf` file.

Table 12.15. object\_store\_access

Configuration option = Default value	Type	Description
<b>public_identity_ca_file = None</b>	string value	Location of ca certificate file to use for identity client requests via public endpoint
<b>public_object_store_ca_file = None</b>	string value	Location of ca certificate file to use for object-store client requests via public endpoint

### 12.1.16. oslo\_concurrency

The following table outlines the options available under the **[oslo\_concurrency]** group in the `/etc/sahara/sahara.conf` file.

Table 12.16. oslo\_concurrency

Configuration option = Default value	Type	Description
<b>disable_process_locking = False</b>	boolean value	Enables or disables inter-process locks.
<b>lock_path = None</b>	string value	Directory to use for lock files. For security, the specified directory should only be writable by the user running the processes that need locking. Defaults to environment variable OSLO_LOCK_PATH. If external locks are used, a lock path must be set.

### 12.1.17. oslo\_messaging\_amqp

The following table outlines the options available under the **[oslo\_messaging\_amqp]** group in the `/etc/sahara/sahara.conf` file.

Table 12.17. oslo\_messaging\_amqp

Configuration option = Default value	Type	Description
---	------	-------------

Configuration option = Default value	Type	Description
<b>addressing_mode = dynamic</b>	string value	Indicates the addressing mode used by the driver. Permitted values: <i>legacy</i> - use legacy non-routable addressing <i>routable</i> - use routable addresses <i>dynamic</i> - use legacy addresses if the message bus does not support routing otherwise use routable addressing
<b>allow_insecure_clients = False</b>	boolean value	Accept clients using either SSL or plain TCP
<b>anycast_address = anycast</b>	string value	Appended to the address prefix when sending to a group of consumers. Used by the message bus to identify messages that should be delivered in a round-robin fashion across consumers.
<b>broadcast_prefix = broadcast</b>	string value	address prefix used when broadcasting to all servers
<b>connection_retry_backoff = 2</b>	integer value	Increase the <code>connection_retry_interval</code> by this many seconds after each unsuccessful failover attempt.
<b>connection_retry_interval = 1</b>	integer value	Seconds to pause before attempting to re-connect.
<b>connection_retry_interval _max = 30</b>	integer value	Maximum limit for <code>connection_retry_interval</code> + <code>connection_retry_backoff</code>
<b>container_name = None</b>	string value	Name for the AMQP container. must be globally unique. Defaults to a generated UUID
<b>default_notification_exch ange = None</b>	string value	Exchange name used in notification addresses. Exchange name resolution precedence: Target.exchange if set else default_notification_exchange if set else control_exchange if set else <i>notify</i>
<b>default_notify_timeout = 30</b>	integer value	The deadline for a sent notification message delivery. Only used when caller does not provide a timeout expiry.
<b>default_reply_retry = 0</b>	integer value	The maximum number of attempts to re-send a reply message which failed due to a recoverable error.
<b>default_reply_timeout = 30</b>	integer value	The deadline for an rpc reply message delivery.

Configuration option = Default value	Type	Description
<b>default_rpc_exchange = None</b>	string value	Exchange name used in RPC addresses. Exchange name resolution precedence: Target.exchange if set else default_rpc_exchange if set else control_exchange if set else <i>rpc</i>
<b>default_send_timeout = 30</b>	integer value	The deadline for an rpc cast or call message delivery. Only used when caller does not provide a timeout expiry.
<b>default_sender_link_timeout = 600</b>	integer value	The duration to schedule a purge of idle sender links. Detach link after expiry.
<b>group_request_prefix = unicast</b>	string value	address prefix when sending to any server in group
<b>idle_timeout = 0</b>	integer value	Timeout for inactive connections (in seconds)
<b>link_retry_delay = 10</b>	integer value	Time to pause between re-connecting an AMQP 1.0 link that failed due to a recoverable error.
<b>multicast_address = multicast</b>	string value	Appended to the address prefix when sending a fanout message. Used by the message bus to identify fanout messages.
<b>notify_address_prefix = openstack.org/om/notify</b>	string value	Address prefix for all generated Notification addresses
<b>notify_server_credit = 100</b>	integer value	Window size for incoming Notification messages
<b>password = `</b>	string value	Password for message broker authentication
<b>pre_settled = [u'rpc-cast', u'rpc-reply']</b>	multi valued	Send messages of this type pre-settled. Pre-settled messages will not receive acknowledgement from the peer. Note well: pre-settled messages may be silently discarded if the delivery fails. Permitted values: <i>rpc-call</i> - send RPC Calls pre-settled <i>rpc-reply</i> - send RPC Replies pre-settled <i>rpc-cast</i> - Send RPC Casts pre-settled <i>notify</i> - Send Notifications pre-settled
<b>pseudo_vhost = True</b>	boolean value	Enable virtual host support for those message buses that do not natively support virtual hosting (such as <i>qpidd</i> ). When set to true the virtual host name will be added to all message bus addresses, effectively creating a private <i>subnet</i> per virtual host. Set to False if the message bus supports virtual hosting using the <i>hostname</i> field in the AMQP 1.0 Open performative as the name of the virtual host.



Configuration option = Default value	Type	Description
<b>reply_link_credit = 200</b>	integer value	Window size for incoming RPC Reply messages.
<b>rpc_address_prefix = openstack.org/om/rpc</b>	string value	Address prefix for all generated RPC addresses
<b>rpc_server_credit = 100</b>	integer value	Window size for incoming RPC Request messages
<code>`sasl_config_dir = `</code>	string value	Path to directory that contains the SASL configuration
<code>`sasl_config_name = `</code>	string value	Name of configuration file (without .conf suffix)
<code>`sasl_default_realm = `</code>	string value	SASL realm to use if no realm present in username
<code>`sasl_mechanisms = `</code>	string value	Space separated list of acceptable SASL mechanisms
<b>server_request_prefix = exclusive</b>	string value	address prefix used when sending to a specific server
<b>ssl = False</b>	boolean value	Attempt to connect via SSL. If no other ssl-related parameters are given, it will use the system's CA-bundle to verify the server's certificate.
<code>`ssl_ca_file = `</code>	string value	CA certificate PEM file used to verify the server's certificate
<code>`ssl_cert_file = `</code>	string value	Self-identifying certificate PEM file for client authentication
<code>`ssl_key_file = `</code>	string value	Private key PEM file used to sign ssl_cert_file certificate (optional)
<b>ssl_key_password = None</b>	string value	Password for decrypting ssl_key_file (if encrypted)
<b>ssl_verify_vhost = False</b>	boolean value	By default SSL checks that the name in the server's certificate matches the hostname in the transport_url. In some configurations it may be preferable to use the virtual hostname instead, for example if the server uses the Server Name Indication TLS extension (rfc6066) to provide a certificate per virtual host. Set ssl_verify_vhost to True if the server's SSL certificate uses the virtual host name instead of the DNS name.

Configuration option = Default value	Type	Description
<b>trace = False</b>	boolean value	Debug: dump AMQP frames to stdout
<b>unicast_address = unicast</b>	string value	Appended to the address prefix when sending to a particular RPC/Notification server. Used by the message bus to identify messages sent to a single destination.
<code>`username = `</code>	string value	User name for message broker authentication

### 12.1.18. oslo\_messaging\_kafka

The following table outlines the options available under the **[oslo\_messaging\_kafka]** group in the `/etc/sahara/sahara.conf` file.

Table 12.18. oslo\_messaging\_kafka

Configuration option = Default value	Type	Description
<b>conn_pool_min_size = 2</b>	integer value	The pool size limit for connections expiration policy
<b>conn_pool_ttl = 1200</b>	integer value	The time-to-live in sec of idle connections in the pool
<b>consumer_group = oslo_messaging_consumer</b>	string value	Group id for Kafka consumer. Consumers in one group will coordinate message consumption
<b>kafka_consumer_timeout = 1.0</b>	floating point value	Default timeout(s) for Kafka consumers
<b>kafka_default_host = localhost</b>	string value	Default Kafka broker Host
<b>kafka_default_port = 9092</b>	port value	Default Kafka broker Port
<b>kafka_max_fetch_bytes = 1048576</b>	integer value	Max fetch bytes of Kafka consumer
<b>pool_size = 10</b>	integer value	Pool Size for Kafka Consumers
<b>producer_batch_size = 16384</b>	integer value	Size of batch for the producer async send
<b>producer_batch_timeout = 0.0</b>	floating point value	Upper bound on the delay for KafkaProducer batching in seconds

### 12.1.19. oslo\_messaging\_notifications

The following table outlines the options available under the **[oslo\_messaging\_notifications]** group in the **/etc/sahara/sahara.conf** file.

Table 12.19. oslo\_messaging\_notifications

Configuration option = Default value	Type	Description
<b>driver = []</b>	multi valued	The Drivers(s) to handle sending notifications. Possible values are messaging, messagingv2, routing, log, test, noop
<b>level = INFO</b>	string value	Notification level for outgoing notifications
<b>publisher_id = None</b>	string value	Identifier of the publisher
<b>retry = -1</b>	integer value	The maximum number of attempts to re-send a notification message which failed to be delivered due to a recoverable error. 0 - No retry, -1 - indefinite
<b>topics = [u'notifications']</b>	list value	AMQP topic used for OpenStack notifications.
<b>transport_url = None</b>	string value	A URL representing the messaging driver to use for notifications. If not set, we fall back to the same configuration used for RPC.

### 12.1.20. oslo\_messaging\_rabbit

The following table outlines the options available under the **[oslo\_messaging\_rabbit]** group in the **/etc/sahara/sahara.conf** file.

Table 12.20. oslo\_messaging\_rabbit

Configuration option = Default value	Type	Description
<b>amqp_auto_delete = False</b>	boolean value	Auto-delete queues in AMQP.
<b>amqp_durable_queues = False</b>	boolean value	Use durable queues in AMQP.
<b>channel_max = None</b>	integer value	Maximum number of channels to allow
<b>connection_factory = single</b>	string value	Connection factory implementation

Configuration option = Default value	Type	Description
<b>default_notification_exchange =</b> <b>`\${control_exchange}_notification`</b>	string value	Exchange name for sending notifications
<b>default_notification_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending notification, -1 means infinite retry.
<b>default_rpc_exchange =</b> <b>`\${control_exchange}_rpc`</b>	string value	Exchange name for sending RPC messages
<b>default_rpc_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending RPC message, -1 means infinite retry. If actual retry attempts in not 0 the rpc request could be processed more than one time
<b>default_serializer_type =</b> <b>json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>fake_rabbit = False</b>	boolean value	Deprecated, use <code>rpc_backend=kombu+memory</code> or <code>rpc_backend=fake</code>
<b>frame_max = None</b>	integer value	The maximum byte size for an AMQP frame
<b>heartbeat_interval = 3</b>	integer value	How often to send heartbeats for consumer's connections
<b>heartbeat_rate = 2</b>	integer value	How often times during the <code>heartbeat_timeout_threshold</code> we check the heartbeat.
<b>heartbeat_timeout_threshold = 60</b>	integer value	Number of seconds after which the Rabbit broker is considered down if heartbeat's keep-alive fails (0 disable the heartbeat). EXPERIMENTAL
<b>host_connection_reconnect_delay = 0.25</b>	floating point value	Set delay for reconnection to some host which has connection error
<b>kombu_compression =</b> <b>None</b>	string value	EXPERIMENTAL: Possible values are: <code>gzip</code> , <code>bz2</code> . If not set compression will not be used. This option may not be available in future versions.

Configuration option = Default value	Type	Description
<b>kombu_failover_strategy = round-robin</b>	string value	Determines how the next RabbitMQ node is chosen in case the one we are currently connected to becomes unavailable. Takes effect only if more than one RabbitMQ node is provided in config.
<b>kombu_missing_consumer_retry_timeout = 60</b>	integer value	How long to wait a missing client before abandoning to send it its replies. This value should not be longer than <code>rpc_response_timeout</code> .
<b>kombu_reconnect_delay = 1.0</b>	floating point value	How long to wait before reconnecting in response to an AMQP consumer cancel notification.
<b>notification_listener_prefetch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to notification listener.
<b>notification_persistence = False</b>	boolean value	Persist notification messages.
<b>notification_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending notification message
<b>pool_max_overflow = 0</b>	integer value	Maximum number of connections to create above <b>pool_max_size</b> .
<b>pool_max_size = 30</b>	integer value	Maximum number of connections to keep queued.
<b>pool_recycle = 600</b>	integer value	Lifetime of a connection (since creation) in seconds or None for no recycling. Expired connections are closed on acquire.
<b>pool_stale = 60</b>	integer value	Threshold at which inactive (since release) connections are considered stale in seconds or None for no staleness. Stale connections are closed on acquire.
<b>pool_timeout = 30</b>	integer value	Default number of seconds to wait for a connections to available
<b>rabbit_ha_queues = False</b>	boolean value	Try to use HA queues in RabbitMQ ( <code>x-ha-policy: all</code> ). If you change this option, you must wipe the RabbitMQ database. In RabbitMQ 3.0, queue mirroring is no longer controlled by the <code>x-ha-policy</code> argument when declaring a queue. If you just want to make sure that all queues (except those with auto-generated names) are mirrored across all nodes, run: <code>"rabbitmqctl set_policy HA ^{?!amq\\.}.* {\"ha-mode\": \"all\"}"</code>

Configuration option = Default value	Type	Description
<b>rabbit_host = localhost</b>	string value	The RabbitMQ broker address where a single node is used.
<b>rabbit_hosts =</b> <b>[u'\$rabbit_host:\$rabbit_p</b> <b>ort']</b>	list value	RabbitMQ HA cluster host:port pairs.
<b>rabbit_interval_max = 30</b>	integer value	Maximum interval of RabbitMQ connection retries. Default is 30 seconds.
<b>rabbit_login_method =</b> <b>AMQPLAIN</b>	string value	The RabbitMQ login method.
<b>rabbit_max_retries = 0</b>	integer value	Maximum number of RabbitMQ connection retries. Default is 0 (infinite retry count).
<b>rabbit_password = guest</b>	string value	The RabbitMQ password.
<b>rabbit_port = 5672</b>	port value	The RabbitMQ broker port where a single node is used.
<b>rabbit_qos_prefetch_cou</b> <b>nt = 0</b>	integer value	Specifies the number of messages to prefetch. Setting to zero allows unlimited messages.
<b>rabbit_retry_backoff = 2</b>	integer value	How long to backoff for between retries when connecting to RabbitMQ.
<b>rabbit_retry_interval = 1</b>	integer value	How frequently to retry connecting with RabbitMQ.
<b>rabbit_transient_queues_</b> <b>ttn = 1800</b>	integer value	Positive integer representing duration in seconds for queue TTL (x-expires). Queues which are unused for the duration of the TTL are automatically deleted. The parameter affects only reply and fanout queues.
<b>rabbit_userid = guest</b>	string value	The RabbitMQ userid.
<b>rabbit_virtual_host = /</b>	string value	The RabbitMQ virtual host.
<b>rpc_listener_prefetch_co</b> <b>unt = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to rpc listener.
<b>rpc_queue_expiration =</b> <b>60</b>	integer value	Time to live for rpc queues without consumers in seconds.

Configuration option = Default value	Type	Description
<b>rpc_reply_exchange = \${control_exchange}_rpc_reply</b>	string value	Exchange name for receiving RPC replies
<b>rpc_reply_listener_prefetch_count = 100</b>	integer value	Max number of not acknowledged message which RabbitMQ can send to rpc reply listener.
<b>rpc_reply_retry_attempts = -1</b>	integer value	Reconnecting retry count in case of connectivity problem during sending reply. -1 means infinite retry during rpc_timeout
<b>rpc_reply_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending reply.
<b>rpc_retry_delay = 0.25</b>	floating point value	Reconnecting retry delay in case of connectivity problem during sending RPC message
<b>socket_timeout = 0.25</b>	floating point value	Set socket timeout in seconds for connection's socket
<b>ssl = None</b>	boolean value	Enable SSL
<code>`ssl_ca_file = `</code>	string value	SSL certification authority file (valid only if SSL enabled).
<code>`ssl_cert_file = `</code>	string value	SSL cert file (valid only if SSL enabled).
<code>`ssl_key_file = `</code>	string value	SSL key file (valid only if SSL enabled).
<b>ssl_options = None</b>	dict value	Arguments passed to ssl.wrap_socket
<code>`ssl_version = `</code>	string value	SSL version to use (valid only if SSL enabled). Valid values are TLSv1 and SSLv23. SSLv2, SSLv3, TLSv1_1, and TLSv1_2 may be available on some distributions.
<b>tcp_user_timeout = 0.25</b>	floating point value	Set TCP_USER_TIMEOUT in seconds for connection's socket

### 12.1.21. oslo\_messaging\_zmq

The following table outlines the options available under the **[oslo\_messaging\_zmq]** group in the `/etc/sahara/sahara.conf` file.

Table 12.21. oslo\_messaging\_zmq

Configuration option = Default value	Type	Description
<b>rpc_ack_timeout_base = 15</b>	integer value	Number of seconds to wait for an ack from a cast/call. After each retry attempt this timeout is multiplied by some specified multiplier.
<b>rpc_ack_timeout_multiplier = 2</b>	integer value	Number to multiply base ack timeout by after each retry attempt.
<b>rpc_message_ttl = 300</b>	integer value	Expiration timeout in seconds of a sent/received message after which it is not tracked anymore by a client/server.
<b>rpc_poll_timeout = 1</b>	integer value	The default number of seconds that poll should wait. Poll raises timeout exception when timeout expired.
<b>rpc_retry_attempts = 3</b>	integer value	Default number of message sending attempts in case of any problems occurred: positive value N means at most N retries, 0 means no retries, None or -1 (or any other negative values) mean to retry forever. This option is used only if acknowledgments are enabled.
<b>rpc_thread_pool_size = 100</b>	integer value	Maximum number of (green) threads to work concurrently.
<b>rpc_use_acks = False</b>	boolean value	Wait for message acknowledgements from receivers. This mechanism works only via proxy without PUB/SUB.
<b>rpc_zmq_bind_address = *</b>	string value	ZeroMQ bind address. Should be a wildcard (*), an ethernet interface, or IP. The "host" option should point or resolve to this address.
<b>rpc_zmq_bind_port_retries = 100</b>	integer value	Number of retries to find free port number before fail with ZMQBindError.
<b>rpc_zmq_contexts = 1</b>	integer value	Number of ZeroMQ contexts, defaults to 1.
<b>rpc_zmq_host = ccosp-undercloud.usersys.redhat.com</b>	string value	Name of this node. Must be a valid hostname, FQDN, or IP address. Must match "host" option, if running Nova.
<b>rpc_zmq_ipc_dir = /var/run/openstack</b>	string value	Directory for holding IPC sockets.
<b>rpc_zmq_matchmaker = redis</b>	string value	MatchMaker driver.



Configuration option = Default value	Type	Description
<b>rpc_zmq_max_port = 65536</b>	integer value	Maximal port number for random ports range.
<b>rpc_zmq_min_port = 49153</b>	port value	Minimal port number for random ports range.
<b>rpc_zmq_serialization = json</b>	string value	Default serialization mechanism for serializing/deserializing outgoing/incoming messages
<b>rpc_zmq_topic_backlog = None</b>	integer value	Maximum number of ingress messages to locally buffer per topic. Default is unlimited.
<b>subscribe_on = []</b>	list value	List of publisher hosts SubConsumer can subscribe on. This option has higher priority then the default publishers list taken from the matchmaker.
<b>use_dynamic_connections = False</b>	boolean value	This option makes direct connections dynamic or static. It makes sense only with <code>use_router_proxy=False</code> which means to use direct connections for direct message types (ignored otherwise).
<b>use_pub_sub = False</b>	boolean value	Use PUB/SUB pattern for fanout methods. PUB/SUB always uses proxy.
<b>use_router_proxy = False</b>	boolean value	Use ROUTER remote proxy.
<b>zmq_failover_connections = 2</b>	integer value	How many additional connections to a host will be made for failover reasons. This option is actual only in dynamic connections mode.
<b>zmq_immediate = True</b>	boolean value	This option configures round-robin mode in zmq socket. True means not keeping a queue when server side disconnects. False means to keep queue and messages even if server is disconnected, when the server appears we send all accumulated messages to it.
<b>zmq_linger = -1</b>	integer value	Number of seconds to wait before all pending messages will be sent after closing a socket. The default value of -1 specifies an infinite linger period. The value of 0 specifies no linger period. Pending messages shall be discarded immediately when the socket is closed. Positive values specify an upper bound for the linger period.

Configuration option = Default value	Type	Description
<b>zmq_target_expire = 300</b>	integer value	Expiration timeout in seconds of a name service record about existing target ( < 0 means no timeout).
<b>zmq_target_update = 180</b>	integer value	Update period in seconds of a name service record about existing target.
<b>zmq_tcp_keepalive = -1</b>	integer value	Enable/disable TCP keepalive (KA) mechanism. The default value of -1 (or any other negative value) means to skip any overrides and leave it to OS default; 0 and 1 (or any other positive value) mean to disable and enable the option respectively.
<b>zmq_tcp_keepalive_cnt = -1</b>	integer value	The number of retransmissions to be carried out before declaring that remote end is not available. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_idle = -1</b>	integer value	The duration between two keepalive transmissions in idle condition. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.
<b>zmq_tcp_keepalive_intvl = -1</b>	integer value	The duration between two successive keepalive retransmissions, if acknowledgement to the previous keepalive transmission is not received. The unit is platform dependent, for example, seconds in Linux, milliseconds in Windows etc. The default value of -1 (or any other negative value and 0) means to skip any overrides and leave it to OS default.

### 12.1.22. oslo\_middlewares

The following table outlines the options available under the **[oslo\_middlewares]** group in the **/etc/sahara/sahara.conf** file.

Table 12.22. oslo\_middlewares

Configuration option = Default value	Type	Description
<b>enable_proxy_headers_parsing = False</b>	boolean value	Whether the application is behind a proxy or not. This determines if the middleware should parse the headers or not.

### 12.1.23. oslo\_policy

The following table outlines the options available under the **[oslo\_policy]** group in the `/etc/sahara/sahara.conf` file.

Table 12.23. oslo\_policy

Configuration option = Default value	Type	Description
<b>enforce_scope = False</b>	boolean value	This option controls whether or not to enforce scope when evaluating policies. If <b>True</b> , the scope of the token used in the request is compared to the <b>scope_types</b> of the policy being enforced. If the scopes do not match, an <b>InvalidScope</b> exception will be raised. If <b>False</b> , a message will be logged informing operators that policies are being invoked with mismatching scope.
<b>policy_default_rule = default</b>	string value	Default rule. Enforced when a requested rule is not found.
<b>policy_dirs = [u'policy.d']</b>	multi valued	Directories where policy configuration files are stored. They can be relative to any directory in the search path defined by the <code>config_dir</code> option, or absolute paths. The file defined by <code>policy_file</code> must exist for these directories to be searched. Missing or empty directories are ignored.
<b>policy_file = policy.json</b>	string value	The file that defines policies.
<b>remote_content_type = application/x-www-form-urlencoded</b>	string value	Content Type to send and receive data for REST based policy check
<b>remote_ssl_ca_cert_file = None</b>	string value	Absolute path to ca cert file for REST based policy check
<b>remote_ssl_client_cert_file = None</b>	string value	Absolute path to client cert for REST based policy check
<b>remote_ssl_client_key_file = None</b>	string value	Absolute path client key file REST based policy check
<b>remote_ssl_verify_server_cert = False</b>	boolean value	server identity verification for REST based policy check

### 12.1.24. retries

The following table outlines the options available under the **[retries]** group in the `/etc/sahara/sahara.conf` file.

Table 12.24. retries

Configuration option = Default value	Type	Description
<b>retries_number = 5</b>	integer value	Number of times to retry the request to client before failing
<b>retry_after = 10</b>	integer value	Time between the retries to client (in seconds).

### 12.1.25. ssl

The following table outlines the options available under the **[ssl]** group in the `/etc/sahara/sahara.conf` file.

Table 12.25. ssl

Configuration option = Default value	Type	Description
<b>ca_file = None</b>	string value	CA certificate file to use to verify connecting clients.
<b>cert_file = None</b>	string value	Certificate file to use when starting the server securely.
<b>ciphers = None</b>	string value	Sets the list of available ciphers. value should be a string in the OpenSSL cipher list format.
<b>key_file = None</b>	string value	Private key file to use when starting the server securely.
<b>version = None</b>	string value	SSL version to use (valid only if SSL enabled). Valid values are TLSv1 and SSLv23. SSLv2, SSLv3, TLSv1_1, and TLSv1_2 may be available on some distributions.

### 12.1.26. swift

The following table outlines the options available under the **[swift]** group in the `/etc/sahara/sahara.conf` file.

Table 12.26. swift

Configuration option = Default value	Type	Description
<b>api_insecure = False</b>	boolean value	Allow to perform insecure SSL requests to swift.
<b>ca_file = None</b>	string value	Location of ca certificates file to use for swift client requests.

Configuration option = Default value	Type	Description
<b>endpoint_type = internalURL</b>	string value	Endpoint type for swift client requests

### 12.1.27. timeouts

The following table outlines the options available under the **[timeouts]** group in the `/etc/sahara/sahara.conf` file.

Table 12.27. timeouts

Configuration option = Default value	Type	Description
<b>delete_instances_timeout = 10800</b>	integer value	Wait for instances to be deleted, in seconds
<b>detach_volume_timeout = 300</b>	integer value	Timeout for detaching volumes from instance, in seconds
<b>ips_assign_timeout = 10800</b>	integer value	Assign IPs timeout, in seconds
<b>wait_until_accessible = 10800</b>	integer value	Wait for instance accessibility, in seconds

### 12.1.28. trustee

The following table outlines the options available under the **[trustee]** group in the `/etc/sahara/sahara.conf` file.

Table 12.28. trustee

Configuration option = Default value	Type	Description
<b>auth_url = None</b>	string value	Auth url for trusts creation
<b>password = None</b>	string value	Password for trusts creation
<b>project_domain_name = Default</b>	string value	Project domain name for trusts creation
<b>project_name = None</b>	string value	Project name for trusts creation
<b>user_domain_name = Default</b>	string value	User domain name for trusts creation

Configuration option = Default value	Type	Description
<b>username = None</b>	string value	Username for trusts creation