

OpenStack

# Configuration Reference

liberty (June 30, 2015)

**DRAFT**  
Liberty



## OpenStack Configuration Reference

liberty (2015-06-30)

Copyright © 2013-2015 OpenStack Foundation All rights reserved.

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. It includes sample configuration files.

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

# Table of Contents

OpenStack configuration overview .....	xx
Conventions .....	xx
Document change history .....	xxi
Configuration file format .....	xxi
1. Bare metal .....	1
2. Block Storage .....	17
Introduction to the Block Storage service .....	17
Volume drivers .....	18
Backup drivers .....	119
Block Storage sample configuration files .....	123
Log files used by Block Storage .....	176
Fibre Channel Zone Manager .....	177
Volume encryption with static key .....	180
Additional options .....	183
New, updated and deprecated options in Kilo for OpenStack Block Storage .....	203
3. Compute .....	211
Overview of nova.conf .....	211
Configure logging .....	213
Configure authentication and authorization .....	213
Configure resize .....	213
Database configuration .....	214
Configure the Oslo RPC messaging system .....	214
Configure the Compute API .....	218
Configure the EC2 API .....	220
Fibre Channel support in Compute .....	221
iSCSI interface and offload support in Compute .....	221
Hypervisors .....	223
Scheduling .....	256
Cells .....	274
Conductor .....	279
Example <code>nova.conf</code> configuration files .....	279
Compute log files .....	284
Compute sample configuration files .....	284
New, updated and deprecated options in Kilo for OpenStack Compute .....	328
4. Dashboard .....	335
Configure the dashboard .....	335
Customize the dashboard .....	339
Additional sample configuration files .....	341
Dashboard log files .....	352
5. Database service .....	353
Configure the database .....	363
Configure the RPC messaging system .....	369
New, updated and deprecated options in Kilo for Database service .....	373
6. Data processing service .....	378
New, updated and deprecated options in Kilo for Data Processing service .....	389
7. Identity service .....	394
Caching layer .....	394
Identity service configuration file .....	396



---

## List of Figures

2.1. Ceph architecture .....	19
3.1. VMware driver architecture .....	239
3.2. Filtering .....	258
3.3. Weighting hosts .....	269
3.4. KVM, Flat, MySQL, and Glance, OpenStack or EC2 API .....	282
3.5. KVM, Flat, MySQL, and Glance, OpenStack or EC2 API .....	283

## List of Tables

- 1.1. Description of agent configuration options ..... 1
- 1.2. Description of AMQP configuration options ..... 1
- 1.3. Description of AMT configuration options ..... 1
- 1.4. Description of API configuration options ..... 2
- 1.5. Description of authorization token configuration options ..... 2
- 1.6. Description of authorization configuration options ..... 4
- 1.7. Description of Cisco UCS configuration options ..... 4
- 1.8. Description of common configuration options ..... 4
- 1.9. Description of conductor configuration options ..... 5
- 1.10. Description of console configuration options ..... 6
- 1.11. Description of database configuration options ..... 6
- 1.12. Description of logging configuration options ..... 7
- 1.13. Description of deploy configuration options ..... 8
- 1.14. Description of DHCP configuration options ..... 8
- 1.15. Description of disk partitioner configuration options ..... 8
- 1.16. Description of DRAC configuration options ..... 8
- 1.17. Description of glance configuration options ..... 8
- 1.18. Description of iLO configuration options ..... 9
- 1.19. Description of inspector configuration options ..... 10
- 1.20. Description of IPMI configuration options ..... 10
- 1.21. Description of iRMC configuration options ..... 10
- 1.22. Description of keystone configuration options ..... 10
- 1.23. Description of logging configuration options ..... 11
- 1.24. Description of neutron configuration options ..... 12
- 1.25. Description of policy configuration options ..... 12
- 1.26. Description of PXE configuration options ..... 12
- 1.27. Description of Redis configuration options ..... 13
- 1.28. Description of RPC configuration options ..... 13
- 1.29. Description of RabbitMQ configuration options ..... 14
- 1.30. Description of Qpid configuration options ..... 15
- 1.31. Description of SeaMicro configuration options ..... 15
- 1.32. Description of SNMP configuration options ..... 16
- 1.33. Description of SSH configuration options ..... 16
- 1.34. Description of swift configuration options ..... 16
- 1.35. Description of VirtualBox configuration options ..... 16
- 1.36. Description of ZeroMQ configuration options ..... 16
- 2.1. Description of Ceph storage configuration options ..... 20
- 2.2. Description of Dell EqualLogic volume driver configuration options ..... 21
- 2.3. Description of Dell Storage Center volume driver configuration options ..... 26
- 2.4. Description of GlusterFS storage configuration options ..... 45
- 2.5. Configuration options for service labels ..... 48
- 2.6. Configuration options ..... 48
- 2.7. Description of Hitachi storage volume driver configuration options ..... 53
- 2.8. Huawei storage driver configuration options ..... 66
- 2.9. Description of GPFS storage configuration options ..... 68
- 2.10. Volume Create Options for GPFS Volume Drive ..... 69
- 2.11. List of configuration flags for Storwize storage and SVC driver ..... 73
- 2.12. Description of IBM Storwise driver configuration options ..... 74

---

2.13. Description of IBM XIV and DS8000 volume driver configuration options .....	77
2.14. List of configuration flags for IBM FlashSystem FC driver .....	79
2.15. Description of LVM configuration options .....	80
2.16. Description of NetApp cDOT iSCSI driver configuration options .....	82
2.17. Description of NetApp cDOT NFS driver configuration options .....	83
2.18. Description of extra specs options for NetApp Unified Driver with Clustered Data ONTAP .....	86
2.19. Description of NetApp 7-Mode iSCSI driver configuration options .....	87
2.20. Description of NetApp 7-Mode NFS driver configuration options .....	89
2.21. Description of NetApp E-Series driver configuration options .....	91
2.22. Description of Nimble driver configuration options .....	93
2.23. Description of NFS storage configuration options .....	94
2.24. Description of ProphetStor Fibre Channel and iSCSI drivers configuration options .....	97
2.25. Description of Quobyte USP volume driver configuration options .....	100
2.26. Description of Scality SOFS volume driver configuration options .....	102
2.27. Description of Samba volume driver configuration options .....	103
2.28. Description of SolidFire driver configuration options .....	104
2.29. Description of Tintri driver configuration options .....	105
2.30. Description of VMware configuration options .....	106
2.31. Extra spec entry to VMDK disk file type mapping .....	107
2.32. Extra spec entry to clone type mapping .....	107
2.33. Description of Windows configuration options .....	113
2.34. Description of X-IO volume driver configuration options .....	114
2.35. Extra specs .....	115
2.36. Description of ZFS Storage Appliance NFS driver configuration options .....	119
2.37. Description of Ceph backup driver configuration options .....	120
2.38. Description of IBM Tivoli Storage Manager backup driver configuration options .....	121
2.39. Description of Swift backup driver configuration options .....	121
2.40. Description of NFS backup driver configuration options .....	123
2.41. Log files used by Block Storage services .....	177
2.42. Description of zoning configuration options .....	177
2.43. Description of zoning manager configuration options .....	178
2.44. Description of zoning fabrics configuration options .....	178
2.45. Description of cisco zoning manager configuration options .....	179
2.46. Description of cisco zoning fabrics configuration options .....	179
2.47. Description of API configuration options .....	183
2.48. Description of AMQP configuration options .....	184
2.49. Description of authorization configuration options .....	185
2.50. Description of authorization token configuration options .....	185
2.51. Description of backups configuration options .....	187
2.52. Description of block device configuration options .....	187
2.53. Description of CA and SSL configuration options .....	187
2.54. Description of CloudByte volume driver configuration options .....	188
2.55. Description of common configuration options .....	188
2.56. Description of Compute configuration options .....	190
2.57. Description of database configuration options .....	191
2.58. Description of logging configuration options .....	192
2.59. Description of EMC configuration options .....	192
2.60. Description of IBM FlashSystem volume driver configuration options .....	193

---

2.61. Description of HP 3PAR Fibre Channel and iSCSI drivers configuration options.....	193
2.62. Description of HP LeftHand/StoreVirtual driver configuration options .....	193
2.63. Description of Huawei storage driver configuration options .....	193
2.64. Description of IBM NAS volume driver configuration options .....	194
2.65. Description of images configuration options .....	194
2.66. Description of key manager configuration options .....	194
2.67. Description of logging configuration options .....	195
2.68. Description of NAS configuration options .....	196
2.69. Description of Open vStorage driver configuration options .....	197
2.70. Description of oslo_middleware configuration options .....	197
2.71. Description of profiler configuration options .....	197
2.72. Description of Pure Storage driver configuration options .....	197
2.73. Description of Qpid configuration options .....	197
2.74. Description of quota configuration options .....	198
2.75. Description of RabbitMQ configuration options .....	198
2.76. Description of Redis configuration options .....	199
2.77. Description of RPC configuration options .....	199
2.78. Description of SAN configuration options .....	200
2.79. Description of scheduler configuration options .....	200
2.80. Description of SCST volume driver configuration options .....	201
2.81. Description of Scality REST Block storage driver configuration options .....	201
2.82. Description of storage configuration options .....	201
2.83. Description of Violin volume driver configuration options .....	203
2.84. Description of ZeroMQ configuration options .....	203
2.85. Description of zones configuration options .....	203
2.86. New options .....	203
2.87. New default values .....	209
2.88. Deprecated options .....	210
3.1. Description of RabbitMQ configuration options .....	215
3.2. Description of Qpid configuration options .....	216
3.3. Description of ZeroMQ configuration options .....	217
3.4. Description of AMQP configuration options .....	217
3.5. Description of RPC configuration options .....	217
3.6. Default API rate limits .....	219
3.7. vCenter permissions tree .....	241
3.8. OpenStack Image service disk type settings .....	244
3.9. Host weighting options .....	269
3.10. Cell weighting options .....	270
3.11. Log files used by Compute services .....	284
3.12. Description of API configuration options .....	284
3.13. Description of API v3 configuration options .....	285
3.14. Description of authentication configuration options .....	286
3.15. Description of authorization token configuration options .....	286
3.16. Description of availability zones configuration options .....	288
3.17. Description of Barbican configuration options .....	288
3.18. Description of CA and SSL configuration options .....	288
3.19. Description of cell configuration options .....	289
3.20. Description of common configuration options .....	289
3.21. Description of Compute configuration options .....	290
3.22. Description of conductor configuration options .....	292
3.23. Description of config drive configuration options .....	292



3.24. Description of console configuration options .....	293
3.25. Description of database configuration options .....	293
3.26. Description of logging configuration options .....	294
3.27. Description of EC2 configuration options .....	295
3.28. Description of ephemeral storage encryption configuration options .....	295
3.29. Description of fping configuration options .....	295
3.30. Description of glance configuration options .....	295
3.31. Description of HyperV configuration options .....	296
3.32. Description of hypervisor configuration options .....	297
3.33. Description of bare metal configuration options .....	297
3.34. Description of IPv6 configuration options .....	297
3.35. Description of key manager configuration options .....	297
3.36. Description of LDAP configuration options .....	298
3.37. Description of Libvirt configuration options .....	298
3.38. Description of live migration configuration options .....	300
3.39. Description of logging configuration options .....	300
3.40. Description of metadata configuration options .....	302
3.41. Description of network configuration options .....	302
3.42. Description of neutron configuration options .....	305
3.43. Description of oslo_middlewares configuration options .....	306
3.44. Description of PCI configuration options .....	306
3.45. Description of periodic configuration options .....	306
3.46. Description of policy configuration options .....	306
3.47. Description of Quobyte USP volume driver configuration options .....	307
3.48. Description of quota configuration options .....	307
3.49. Description of RDP configuration options .....	308
3.50. Description of Redis configuration options .....	308
3.51. Description of S3 configuration options .....	308
3.52. Description of scheduler configuration options .....	309
3.53. Description of serial console configuration options .....	311
3.54. Description of SPICE configuration options .....	311
3.55. Description of testing configuration options .....	311
3.56. Description of trusted computing configuration options .....	311
3.57. Description of upgrade levels configuration options .....	312
3.58. Description of VMware configuration options .....	312
3.59. Description of VNC configuration options .....	313
3.60. Description of volumes configuration options .....	314
3.61. Description of VPN configuration options .....	315
3.62. Description of Xen configuration options .....	315
3.63. Description of XCP VNC proxy configuration options .....	317
3.64. Description of Zookeeper configuration options .....	317
3.65. New options .....	328
3.66. New default values .....	333
3.67. Deprecated options .....	334
4.1. Dashboard/httpd log files .....	352
5.1. Description of API configuration options .....	353
5.2. Description of authorization token configuration options .....	354
5.3. Description of backup configuration options .....	356
5.4. Description of CA and SSL configuration options .....	356
5.5. Description of clients configuration options .....	356
5.6. Description of cluster configuration options .....	357

---

5.7. Description of common configuration options .....	357
5.8. Description of Compute configuration options .....	357
5.9. Description of logging configuration options .....	358
5.10. Description of DNS configuration options .....	358
5.11. Description of guest agent configuration options .....	359
5.12. Description of Orchestration module configuration options .....	359
5.13. Description of logging configuration options .....	360
5.14. Description of network configuration options .....	361
5.15. Description of nova configuration options .....	361
5.16. Description of quota configuration options .....	361
5.17. Description of Redis configuration options .....	361
5.18. Description of swift configuration options .....	362
5.19. Description of taskmanager configuration options .....	362
5.20. Description of upgrades configuration options .....	362
5.21. Description of volume configuration options .....	363
5.22. Description of database configuration options .....	363
5.23. Description of Cassandra database configuration options .....	363
5.24. Description of Couchbase database configuration options .....	364
5.25. Description of DB2 database configuration options .....	365
5.26. Description of MongoDB database configuration options .....	365
5.27. Description of MySQL database configuration options .....	366
5.28. Description of Percona database configuration options .....	367
5.29. Description of PostgreSQL database configuration options .....	367
5.30. Description of Redis database configuration options .....	368
5.31. Description of Vertica database configuration options .....	369
5.32. Description of RabbitMQ configuration options .....	369
5.33. Description of Qpid configuration options .....	370
5.34. Description of ZeroMQ configuration options .....	371
5.35. Description of AMQP configuration options .....	371
5.36. Description of RPC configuration options .....	372
5.37. New options .....	373
5.38. New default values .....	376
5.39. Deprecated options .....	377
6.1. Description of AMQP configuration options .....	378
6.2. Description of authorization token configuration options .....	378
6.3. Description of CA and SSL configuration options .....	380
6.4. Description of clients configuration options .....	380
6.5. Description of common configuration options .....	381
6.6. Description of database configuration options .....	383
6.7. Description of domain configuration options .....	384
6.8. Description of logging configuration options .....	384
6.9. Description of oslo_middleware configuration options .....	385
6.10. Description of policy configuration options .....	385
6.11. Description of Qpid configuration options .....	386
6.12. Description of RabbitMQ configuration options .....	386
6.13. Description of Redis configuration options .....	387
6.14. Description of RPC configuration options .....	387
6.15. Description of timeouts configuration options .....	388
6.16. Description of ZeroMQ configuration options .....	388
6.17. New options .....	389
6.18. New default values .....	392

---

6.19. Deprecated options .....	393
7.1. Description of cache configuration options .....	394
7.2. Description of API configuration options .....	396
7.3. Description of assignment configuration options .....	398
7.4. Description of authorization configuration options .....	398
7.5. Description of authorization token configuration options .....	398
7.6. Description of CA and SSL configuration options .....	400
7.7. Description of catalog configuration options .....	401
7.8. Description of common configuration options .....	401
7.9. Description of credential configuration options .....	401
7.10. Description of database configuration options .....	401
7.11. Description of logging configuration options .....	402
7.12. Description of domain configuration options .....	402
7.13. Description of EC2 configuration options .....	403
7.14. Description of federation configuration options .....	403
7.15. Description of Fernet tokens configuration options .....	403
7.16. Description of identity configuration options .....	403
7.17. Description of KVS configuration options .....	404
7.18. Description of LDAP configuration options .....	404
7.19. Description of logging configuration options .....	407
7.20. Description of mapping configuration options .....	409
7.21. Description of memcache configuration options .....	409
7.22. Description of OAuth configuration options .....	409
7.23. Description of os_inherit configuration options .....	409
7.24. Description of oslo_middleware configuration options .....	410
7.25. Description of policy configuration options .....	410
7.26. Description of revoke configuration options .....	410
7.27. Description of role configuration options .....	410
7.28. Description of SAML configuration options .....	411
7.29. Description of security configuration options .....	411
7.30. Description of token configuration options .....	411
7.31. Description of trust configuration options .....	412
7.32. Description of RPC configuration options .....	412
7.33. Description of AMQP configuration options .....	413
7.34. Description of Qpid configuration options .....	413
7.35. Description of RabbitMQ configuration options .....	414
7.36. Description of ZeroMQ configuration options .....	415
7.37. Description of Redis configuration options .....	415
7.38. New options .....	452
7.39. New default values .....	457
7.40. Deprecated options .....	457
8.1. Description of authorization token configuration options .....	458
8.2. Description of common configuration options .....	460
8.3. Description of database configuration options .....	462
8.4. Description of logging configuration options .....	463
8.5. Description of Elasticsearch configuration options .....	463
8.6. Description of flagmappings configuration options .....	463
8.7. Description of logging configuration options .....	464
8.8. Description of policy configuration options .....	465
8.9. Description of profiler configuration options .....	466
8.10. Description of Redis configuration options .....	466

---

8.11. Description of registry configuration options .....	466
8.12. Description of replicator configuration options .....	467
8.13. Description of scrubber configuration options .....	467
8.14. Description of TaskFlow configuration options .....	467
8.15. Description of testing configuration options .....	467
8.16. Description of API configuration options .....	468
8.17. Description of CA and SSL configuration options .....	469
8.18. Description of ZeroMQ configuration options .....	470
8.19. Description of AMQP configuration options .....	470
8.20. Description of RPC configuration options .....	471
8.21. Description of RabbitMQ configuration options .....	471
8.22. Description of Qpid configuration options .....	472
8.23. Description of cinder configuration options .....	474
8.24. Description of filesystem configuration options .....	474
8.25. Description of GridFS configuration options .....	474
8.26. Description of RBD configuration options .....	474
8.27. Description of S3 configuration options .....	475
8.28. Description of Sheepdog configuration options .....	475
8.29. Description of swift configuration options .....	475
8.30. Description of VMware configuration options .....	477
8.31. New options .....	507
8.32. New default values .....	512
8.33. Deprecated options .....	512
9.1. Description of common configuration options .....	514
9.2. Description of BigSwitch configuration options .....	517
9.3. Description of Brocade configuration options .....	518
9.4. Description of Brocade MLX L3 plug-in configuration options .....	519
9.5. Description of Brocade Vyatta L3 plug-in configuration options .....	519
9.6. Description of Cisco configuration options .....	519
9.7. Description of HyperV agent configuration options .....	521
9.8. Description of Embrane configuration options .....	521
9.9. Description of SDN-VE configuration options .....	522
9.10. Description of Layer 2 Gateway configuration options .....	522
9.11. Description of Linux Bridge agent configuration options .....	523
9.12. Description of meta configuration options .....	523
9.13. Description of ML2 configuration options .....	524
9.14. Description of ML2 Flat mechanism driver configuration options .....	525
9.15. Description of ML2 GRE configuration options .....	525
9.16. Description of ML2 VLAN configuration options .....	525
9.17. Description of ML2 VXLN configuration options .....	525
9.18. Description of ML2 Arista mechanism driver configuration options .....	525
9.19. Description of Arista layer-3 service plug-in configuration options .....	526
9.20. Description of ML2 BigSwitch mechanism driver configuration options .....	526
9.21. Description of ML2 Brocade mechanism driver configuration options .....	528
9.22. Description of ML2 Brocade MLX ICX mechanism driver configuration options.....	529
9.23. Description of ML2 Cisco mechanism driver configuration options .....	529
9.24. Description of ML2 Freescale SDN mechanism driver configuration options .....	530
9.25. Description of ML2 OpenDaylight mechanism driver configuration options .....	531
9.26. Description of ML2 ofagent mechanism driver configuration options .....	531
9.27. Description of ML2 L2 population configuration options .....	531
9.28. Description of ML2 NCS mechanism driver configuration options .....	532

---

9.29. Description of ML2 ML2 SR-IOV driver configuration options .....	532
9.30. Description of Midonet configuration options .....	532
9.31. Description of Nec configuration options .....	532
9.32. Description of NVSD driver configuration options .....	533
9.33. Description of OpenContrail configuration options .....	533
9.34. Description of Open vSwitch agent configuration options .....	533
9.35. Description of PLUMgrid configuration options .....	534
9.36. Description of SR-IOV configuration options .....	535
9.37. Description of VMware configuration options .....	535
9.38. Description of VMware NSX configuration options .....	536
9.39. Description of Load-Balancer-as-a-Service agent configuration options .....	538
9.40. Description of RabbitMQ configuration options .....	539
9.41. Description of Qpid configuration options .....	540
9.42. Description of ZeroMQ configuration options .....	541
9.43. Description of RPC configuration options .....	542
9.44. Description of Redis configuration options .....	542
9.45. Description of AMQP configuration options .....	543
9.46. Description of agent configuration options .....	543
9.47. Description of API configuration options .....	543
9.48. Description of authorization token configuration options .....	544
9.49. Description of Compute configuration options .....	546
9.50. Description of database configuration options .....	546
9.51. Description of logging configuration options .....	547
9.52. Description of DHCP agent configuration options .....	548
9.53. Description of DVR configuration options .....	548
9.54. Description of FwaaS configuration options .....	548
9.55. Description of IPv6 router advertisement configuration options .....	548
9.56. Description of L3 agent configuration options .....	549
9.57. Description of logging configuration options .....	550
9.58. Description of metadata configuration options .....	552
9.59. Description of metering agent configuration options .....	553
9.60. Description of nova configuration options .....	553
9.61. Description of oslo_middleware configuration options .....	553
9.62. Description of policy configuration options .....	554
9.63. Description of quotas configuration options .....	554
9.64. Description of scheduler configuration options .....	555
9.65. Description of security groups configuration options .....	555
9.66. Description of CA and SSL configuration options .....	556
9.67. Log files used by Networking services .....	556
9.68. New options .....	584
9.69. New default values .....	591
9.70. Deprecated options .....	592
10.1. Description of configuration options for [swift-hash] in swift.conf .....	595
10.2. Description of configuration options for [DEFAULT] in object-server.conf .....	596
10.3. Description of configuration options for [app-object-server] in object-server.conf .....	597
10.4. Description of configuration options for [pipeline-main] in object-server.conf .....	598
10.5. Description of configuration options for [object-replicator] in object-server.conf .....	598

---

10.6. Description of configuration options for [object-updater] in object-server.conf .....	598
10.7. Description of configuration options for [object-auditor] in object-server.conf .....	599
10.8. Description of configuration options for [filter-healthcheck] in object-server.conf .....	599
10.9. Description of configuration options for [filter-recon] in object-server.conf .....	600
10.10. Description of configuration options for [filter-xprofile] in object-server.conf .....	600
10.11. Description of configuration options for [DEFAULT] in object-expirer.conf .....	605
10.12. Description of configuration options for [app-proxy-server] in object-expirer.conf .....	606
10.13. Description of configuration options for [filter-cache] in object-expirer.conf .....	606
10.14. Description of configuration options for [filter-catch_errors] in object-expirer.conf .....	606
10.15. Description of configuration options for [filter-proxy-logging] in object-expirer.conf .....	606
10.16. Description of configuration options for [object-expirer] in object-expirer.conf .....	607
10.17. Description of configuration options for [pipeline-main] in object-expirer.conf .....	607
10.18. Description of configuration options for [DEFAULT] in container-server.conf .....	609
10.19. Description of configuration options for [app-container-server] in container-server.conf .....	611
10.20. Description of configuration options for [pipeline-main] in container-server.conf .....	611
10.21. Description of configuration options for [container-replicator] in container-server.conf .....	611
10.22. Description of configuration options for [container-updater] in container-server.conf .....	612
10.23. Description of configuration options for [container-auditor] in container-server.conf .....	612
10.24. Description of configuration options for [container-sync] in container-server.conf .....	612
10.25. Description of configuration options for [filter-healthcheck] in container-server.conf .....	613
10.26. Description of configuration options for [filter-recon] in container-server.conf .....	613
10.27. Description of configuration options for [filter-xprofile] in container-server.conf .....	613
10.28. Description of configuration options for [DEFAULT] in container-sync-realms.conf .....	617
10.29. Description of configuration options for [realm1] in container-sync-realms.conf .....	617
10.30. Description of configuration options for [realm2] in container-sync-realms.conf .....	617

---

10.31. Description of configuration options for [DEFAULT] in container-reconciler.conf .....	618
10.32. Description of configuration options for [app-proxy-server] in container-reconciler.conf .....	619
10.33. Description of configuration options for [container-reconciler] in container-reconciler.conf .....	619
10.34. Description of configuration options for [filter-cache] in container-reconciler.conf .....	619
10.35. Description of configuration options for [filter-catch_errors] in container-reconciler.conf .....	619
10.36. Description of configuration options for [filter-proxy-logging] in container-reconciler.conf .....	619
10.37. Description of configuration options for [pipeline-main] in container-reconciler.conf .....	620
10.38. Description of configuration options for [DEFAULT] in account-server.conf .....	621
10.39. Description of configuration options for [app-account-server] in account-server.conf .....	622
10.40. Description of configuration options for [pipeline-main] in account-server.conf .....	622
10.41. Description of configuration options for [account-replicator] in account-server.conf .....	622
10.42. Description of configuration options for [account-auditor] in account-server.conf .....	623
10.43. Description of configuration options for [account-reaper] in account-server.conf .....	623
10.44. Description of configuration options for [filter-healthcheck] in account-server.conf .....	624
10.45. Description of configuration options for [filter-recon] in account-server.conf .....	624
10.46. Description of configuration options for [filter-xprofile] in account-server.conf .....	624
10.47. Description of configuration options for [DEFAULT] in proxy-server.conf .....	628
10.48. Description of configuration options for [app-proxy-server] in proxy-server.conf .....	629
10.49. Description of configuration options for [pipeline-main] in proxy-server.conf .....	630
10.50. Description of configuration options for [filter-account-quotas] in proxy-server.conf .....	631
10.51. Description of configuration options for [filter-authtoken] in proxy-server.conf .....	631
10.52. Description of configuration options for [filter-cache] in proxy-server.conf .....	631
10.53. Description of configuration options for [filter-catch_errors] in proxy-server.conf .....	631
10.54. Description of configuration options for [filter-container_sync] in proxy-server.conf .....	632
10.55. Description of configuration options for [filter-dlo] in proxy-server.conf .....	632

---

10.56. Description of configuration options for [filter-gatekeeper] in proxy-server.conf .....	632
10.57. Description of configuration options for [filter-healthcheck] in proxy-server.conf .....	632
10.58. Description of configuration options for [filter-keystoneauth] in proxy-server.conf .....	632
10.59. Description of configuration options for [filter-list-endpoints] in proxy-server.conf .....	633
10.60. Description of configuration options for [filter-proxy-logging] in proxy-server.conf .....	633
10.61. Description of configuration options for [filter-tempauth] in proxy-server.conf .....	634
10.62. Description of configuration options for [filter-xprofile] in proxy-server.conf .....	634
10.63. Description of configuration options for [memcache] in memcache.conf .....	647
10.64. Description of configuration options for [account] in rsyncd.conf .....	647
10.65. Description of configuration options for [container] in rsyncd.conf .....	647
10.66. Description of configuration options for [object] in rsyncd.conf .....	648
10.67. Description of configuration options for [filter-ratelimit] in proxy-server.conf .....	649
10.68. Values for Rate Limiting with Sample Configuration Settings .....	650
10.69. Description of configuration options for [filter-healthcheck] in account-server.conf .....	651
10.70. Description of configuration options for [filter-domain_remap] in proxy-server.conf .....	651
10.71. Description of configuration options for [filter-cname_lookup] in proxy-server.conf .....	651
10.72. Description of configuration options for [filter-tempurl] in proxy-server.conf .....	654
10.73. Description of configuration options for [filter-name_check] in proxy-server.conf .....	654
10.74. Description of configuration options for [swift-constraints] in swift.conf .....	655
10.75. Description of configuration options for [dispersion] in dispersion.conf .....	657
10.76. Description of configuration options for [filter-slo] in proxy-server.conf .....	657
10.77. Description of configuration options for [filter-container-quotas] in proxy-server.conf .....	658
10.78. Description of configuration options for [filter-bulk] in proxy-server.conf .....	659
10.79. Description of configuration options for [drive-audit] in drive-audit.conf .....	661
10.80. Description of configuration options for [filter-formpost] in proxy-server.conf .....	663
10.81. Description of configuration options for [filter-staticweb] in proxy-server.conf .....	663
10.82. New options .....	664
10.83. New default values .....	666
11.1. Description of authorization token configuration options .....	667
11.2. Description of common configuration options .....	669



---

11.3. Description of crypt configuration options .....	670
11.4. Description of database configuration options .....	670
11.5. Description of logging configuration options .....	671
11.6. Description of load balancer configuration options .....	671
11.7. Description of logging configuration options .....	671
11.8. Description of oslo_middleware configuration options .....	672
11.9. Description of quota configuration options .....	673
11.10. Description of Redis configuration options .....	673
11.11. Description of testing configuration options .....	673
11.12. Description of API configuration options .....	673
11.13. Description of Cloudformation-compatible API configuration options .....	675
11.14. Description of CloudWatch API configuration options .....	675
11.15. Description of metadata API configuration options .....	676
11.16. Description of waitcondition API configuration options .....	676
11.17. Description of clients configuration options .....	676
11.18. Description of client backends configuration options .....	676
11.19. Description of ceilometer clients configuration options .....	677
11.20. Description of cinder clients configuration options .....	677
11.21. Description of glance clients configuration options .....	677
11.22. Description of heat clients configuration options .....	677
11.23. Description of keystone clients configuration options .....	678
11.24. Description of neutron clients configuration options .....	678
11.25. Description of nova clients configuration options .....	678
11.26. Description of sahara clients configuration options .....	678
11.27. Description of swift clients configuration options .....	679
11.28. Description of trove clients configuration options .....	679
11.29. Description of RabbitMQ configuration options .....	680
11.30. Description of Qpid configuration options .....	681
11.31. Description of ZeroMQ configuration options .....	682
11.32. Description of AMQP configuration options .....	682
11.33. Description of RPC configuration options .....	683
11.34. Description of notification configuration options .....	683
11.35. New options .....	683
11.36. New default values .....	687
11.37. Deprecated options .....	687
12.1. Description of alarm configuration options .....	688
12.2. Description of alarms configuration options .....	688
12.3. Description of AMQP configuration options .....	689
12.4. Description of API configuration options .....	689
12.5. Description of authorization configuration options .....	689
12.6. Description of authorization token configuration options .....	690
12.7. Description of collector configuration options .....	691
12.8. Description of common configuration options .....	692
12.9. Description of concurrency configuration options .....	693
12.10. Description of database configuration options .....	693
12.11. Description of logging configuration options .....	694
12.12. Description of HTTP dispatcher configuration options .....	694
12.13. Description of events configuration options .....	695
12.14. Description of exchange configuration options .....	695
12.15. Description of glance configuration options .....	695
12.16. Description of inspector configuration options .....	696

---

12.17. Description of IPMI configuration options .....	696
12.18. Description of oslo_middleware configuration options .....	696
12.19. Description of logging configuration options .....	696
12.20. Description of MagnetoDB configuration options .....	697
12.21. Description of notification configuration options .....	697
12.22. Description of policy configuration options .....	697
12.23. Description of Qpid configuration options .....	698
12.24. Description of RabbitMQ configuration options .....	698
12.25. Description of Redis configuration options .....	699
12.26. Description of Rados gateway configuration options .....	699
12.27. Description of RPC configuration options .....	699
12.28. Description of service types configuration options .....	700
12.29. Description of swift configuration options .....	701
12.30. Description of TripleO configuration options .....	701
12.31. Description of VMware configuration options .....	701
12.32. Description of XenAPI configuration options .....	701
12.33. Description of Zaqar configuration options .....	701
12.34. Description of ZeroMQ configuration options .....	702
12.35. New options .....	733
12.36. New default values .....	737
12.37. Deprecated options .....	737
B.1. Default ports that OpenStack components use .....	742
B.2. Default ports that secondary services related to OpenStack components use .....	742

# List of Examples

- 2.1. Default (single-instance) configuration ..... 22
- 2.2. Multi back-end Dell EqualLogic configuration ..... 23
- 2.3. Sample iSCSI Configuration ..... 24
- 2.4. Sample FC configuration ..... 25
- 2.5. Sample Block Storage Configuration ..... 102
- 2.6. Sample Compute Configuration ..... 102
- 4.1. Before ..... 338
- 4.2. After ..... 338

# OpenStack configuration overview

OpenStack is a collection of open source project components that enable setting up cloud services. Each component uses similar configuration techniques and a common framework for INI file options.

This guide pulls together multiple references and configuration options for the following OpenStack components:

- Bare metal service
- OpenStack Block Storage
- OpenStack Compute
- OpenStack dashboard
- Database service for OpenStack
- Data processing service
- OpenStack Identity
- OpenStack Image service
- OpenStack Networking
- OpenStack Object Storage
- Orchestration
- Telemetry

## Conventions

The OpenStack documentation uses several typesetting conventions.

## Notices

Notices take these forms:



### Note

A handy tip or reminder.



### Important

Something you must be aware of before proceeding.



## Warning

Critical information about the risk of data loss or security issues.

## Command prompts

- \$ prompt** Any user, including the `root` user, can run commands that are prefixed with the `$` prompt.
- # prompt** The `root` user must run commands that are prefixed with the `#` prompt. You can also prefix these commands with the `sudo` command, if available, to run them.

## Document change history

This version of the guide replaces and obsoletes all earlier versions.

The following table describes the most recent changes:

Revision Date	Summary of Changes
October 15, 2014	<ul style="list-style-type: none"> <li>Updates for Juno: updated all configuration tables, include sample configuration files, add chapter for Data processing service, update and enhance driver configuration.</li> </ul>
April 16, 2014	<ul style="list-style-type: none"> <li>Update for Icehouse: Updated all configuration tables, include sample configuration files, add chapters for Database service, Orchestration, and Telemetry.</li> </ul>
March 11, 2014	<ul style="list-style-type: none"> <li>Sorted component listing. Moved procedures to the <a href="#">Cloud Administrator Guide</a></li> </ul>
January 9, 2014	<ul style="list-style-type: none"> <li>Removes content addressed in installation, merges duplicated content, and revises legacy references.</li> </ul>
October 17, 2013	<ul style="list-style-type: none"> <li>Havana release.</li> </ul>
August 16, 2013	<ul style="list-style-type: none"> <li>Moves Block Storage driver configuration information from the <i>Block Storage Administration Guide</i> to this reference.</li> </ul>
June 10, 2013	<ul style="list-style-type: none"> <li>Initial creation of Configuration Reference.</li> </ul>

## Configuration file format

OpenStack uses the *INI* file format for configuration files. An INI file is a simple text file that specifies options as `key=value` pairs, grouped into sections. The `DEFAULT` section contains most of the configuration options. Lines starting with a hash sign (`#`) are comment lines. For example:

```
[DEFAULT]
# Print debugging output (set logging level to DEBUG instead
# of default WARNING level). (boolean value)
debug = true
# Print more verbose output (set logging level to INFO instead
# of default WARNING level). (boolean value)
verbose = true

[database]
# The SQLAlchemy connection string used to connect to the
# database (string value)
connection = mysql://keystone:KEYSTONE_DBPASS@controller/keystone
```



```
# onready allows you to send a notification when the
# process
# is ready to serve. For example, to have it notify
# using
# systemd, one could set shell command: "onready =
# systemd-
# notify --ready" or a module with notify() method:
# "onready =
# keystone.common.systemd". (string value)
onready = systemd-notify --ready

# If an instance is passed with the log message,
# format it
# like this (string value)
instance_format = "[instance: %(uuid)s] "
```

## Sections

Configuration options are grouped by section. Most configuration files support at least the following sections:

- [DEFAULT]** Contains most configuration options. If the documentation for a configuration option does not specify its section, assume that it appears in this section.
- [database]** Configuration options for the database that stores the state of the OpenStack service.

## Substitution

The configuration file supports variable substitution. After you set a configuration option, it can be referenced in later configuration values when you precede it with a `$`, like `$OPTION`.

The following example uses the values of `rabbit_host` and `rabbit_port` to define the value of the `rabbit_hosts` option, in this case as `controller:5672`.

```
# The RabbitMQ broker address where a single node is used.
# (string value)
rabbit_host = controller

# The RabbitMQ broker port where a single node is used.
# (integer value)
rabbit_port = 5672

# RabbitMQ HA cluster host:port pairs. (list value)
rabbit_hosts = $rabbit_host:$rabbit_port
```

To avoid substitution, use `$$`, it is replaced by a single `$`. For example, if your LDAP DNS password is `$xkj432`, specify it, as follows:

```
ldap_dns_password = $$xkj432
```

The code uses the Python `string.Template.safe_substitute()` method to implement variable substitution. For more details on how variable substitution is resolved, see <http://docs.python.org/2/library/string.html#template-strings> and [PEP 292](#).

---

## Whitespace

To include whitespace in a configuration value, use a quoted string. For example:

```
ldap_dns_password='a password with spaces'
```

## Define an alternate location for a config file

Most services and the **\*-manage** command-line clients load the configuration file. To define an alternate location for the configuration file, pass the `--config-file CONFIG_FILE` parameter when you start a service or call a **\*-manage** command.



# 1. Bare metal

The Bare metal service is capable of managing and provisioning physical machines. The configuration file of this module is `/etc/ironic/ironic.conf`.

The following tables provide a comprehensive list of the Bare metal service configuration options.

**Table 1.1. Description of agent configuration options**

Configuration option = Default value	Description
[agent]	
<code>agent_api_version = v1</code>	(StrOpt) API version to use for communicating with the ramdisk agent.
<code>agent_erase_devices_priority = None</code>	(IntOpt) 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.
<code>agent_pxe_append_params = nofb nomodeset vga=normal</code>	(StrOpt) Additional append parameters for baremetal PXE boot.
<code>agent_pxe_bootfile_name = pxelinux.0</code>	(StrOpt) Neutron bootfile DHCP parameter.
<code>agent_pxe_config_template = \$pybasedir/drivers/modules/agent_config.template</code>	(StrOpt) Template file for PXE configuration.
<code>heartbeat_timeout = 300</code>	(IntOpt) Maximum interval (in seconds) for agent heartbeats.
<code>manage_tftp = True</code>	(BoolOpt) Whether Ironic will manage TFTP files for the deploy ramdisks. If set to False, you will need to configure your own TFTP server that allows booting the deploy ramdisks.

**Table 1.2. Description of AMQP configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>control_exchange = openstack</code>	(StrOpt) The default exchange under which topics are scoped. May be overridden by an exchange name specified in the <code>transport_url</code> option.
<code>notification_driver = []</code>	(MultiStrOpt) Driver or drivers to handle sending notifications.
<code>notification_topics = notifications</code>	(ListOpt) AMQP topic used for OpenStack notifications.
<code>transport_url = None</code>	(StrOpt) A URL representing the messaging driver to use and its full configuration. If not set, we fall back to the <code>rpc_backend</code> option and driver specific configuration.

**Table 1.3. Description of AMT configuration options**

Configuration option = Default value	Description
[amt]	
<code>action_wait = 10</code>	(IntOpt) Amount of time (in seconds) to wait, before retrying an AMT operation
<code>max_attempts = 3</code>	(IntOpt) Maximum number of times to attempt an AMT operation, before failing
<code>protocol = http</code>	(StrOpt) Protocol used for AMT endpoint, support http/https









Configuration option = Default value	Description
	(True) or should the database be updated based on the hardware state (False).
<code>heartbeat_interval = 10</code>	(IntOpt) Seconds between conductor heart beats.
<code>heartbeat_timeout = 60</code>	(IntOpt) Maximum time (in seconds) since the last check-in of a conductor. A conductor is considered inactive when this time has been exceeded.
<code>inspect_timeout = 1800</code>	(IntOpt) Timeout (seconds) for waiting for node inspection. 0 - unlimited.
<code>node_locked_retry_attempts = 3</code>	(IntOpt) Number of attempts to grab a node lock.
<code>node_locked_retry_interval = 1</code>	(IntOpt) Seconds to sleep between node lock attempts.
<code>periodic_max_workers = 8</code>	(IntOpt) Maximum number of worker threads that can be started simultaneously by a periodic task. Should be less than RPC thread pool size.
<code>power_state_sync_max_retries = 3</code>	(IntOpt) During <code>sync_power_state</code> failures, limit the number of times Ironic should try syncing the hardware node power state with the node power state in DB
<code>send_sensor_data = False</code>	(BoolOpt) Enable sending sensor data message via the notification bus
<code>send_sensor_data_interval = 600</code>	(IntOpt) Seconds between conductor sending sensor data message to ceilometer via the notification bus.
<code>send_sensor_data_types = ALL</code>	(ListOpt) 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.
<code>sync_local_state_interval = 180</code>	(IntOpt) 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.
<code>sync_power_state_interval = 60</code>	(IntOpt) Interval between syncing the node power state to the database, in seconds.
<code>workers_pool_size = 100</code>	(IntOpt) The size of the workers greenthread pool.

**Table 1.10. Description of console configuration options**

Configuration option = Default value	Description
[console]	
<code>subprocess_checking_interval = 1</code>	(IntOpt) Time interval (in seconds) for checking the status of console subprocess.
<code>subprocess_timeout = 10</code>	(IntOpt) Time (in seconds) to wait for the console subprocess to start.
<code>terminal = shellinaboxd</code>	(StrOpt) Path to serial console terminal program
<code>terminal_cert_dir = None</code>	(StrOpt) Directory containing the terminal SSL cert(PEM) for serial console access
<code>terminal_pid_dir = None</code>	(StrOpt) Directory for holding terminal pid files. If not specified, the temporary directory will be used.

**Table 1.11. Description of database configuration options**

Configuration option = Default value	Description
[database]	
<code>backend = sqlalchemy</code>	(StrOpt) The back end to use for the database.
<code>connection = None</code>	(StrOpt) The SQLAlchemy connection string to use to connect to the database.

























## 2. Block Storage

### Table of Contents

Introduction to the Block Storage service .....	17
Volume drivers .....	18
Backup drivers .....	119
Block Storage sample configuration files .....	123
Log files used by Block Storage .....	176
Fibre Channel Zone Manager .....	177
Volume encryption with static key .....	180
Additional options .....	183
New, updated and deprecated options in Kilo for OpenStack Block Storage .....	203

The OpenStack Block Storage service works with many different storage drivers that you can configure by using these instructions.

### Introduction to the Block Storage service

The OpenStack Block Storage service provides persistent block storage resources that OpenStack Compute instances can consume. This includes secondary attached storage similar to the Amazon Elastic Block Storage (EBS) offering. In addition, you can write images to a Block Storage device for Compute to use as a bootable persistent instance.

The Block Storage service differs slightly from the Amazon EBS offering. The Block Storage service does not provide a shared storage solution like NFS. With the Block Storage service, you can attach a device to only one instance.

The Block Storage service provides:

- `cinder-api`. A WSGI app that authenticates and routes requests throughout the Block Storage service. It supports the OpenStack APIs only, although there is a translation that can be done through Compute's EC2 interface, which calls in to the Block Storage client.
- `cinder-scheduler`. Schedules and routes requests to the appropriate volume service. Depending upon your configuration, this may be simple round-robin scheduling to the running volume services, or it can be more sophisticated through the use of the Filter Scheduler. The Filter Scheduler is the default and enables filters on things like Capacity, Availability Zone, Volume Types, and Capabilities as well as custom filters.
- `cinder-volume`. Manages Block Storage devices, specifically the back-end devices themselves.
- `cinder-backup`. Provides a means to back up a Block Storage volume to OpenStack Object Storage (swift).

The Block Storage service contains the following components:

- **Back-end Storage Devices**. The Block Storage service requires some form of back-end storage that the service is built on. The default implementation is to use LVM on a local




## Ceph RADOS Block Device (RBD)

If you use KVM or QEMU as your hypervisor, you can configure the Compute service to use [Ceph RADOS block devices \(RBD\)](#) for volumes.

Ceph is a massively scalable, open source, distributed storage system. It is comprised of an object store, block store, and a POSIX-compliant distributed file system. The platform can auto-scale to the exabyte level and beyond. It runs on commodity hardware, is self-healing and self-managing, and has no single point of failure. Ceph is in the Linux kernel and is integrated with the OpenStack cloud operating system. Due to its open-source nature, you can install and use this portable storage platform in public or private clouds.

**Figure 2.1. Ceph architecture**



## RADOS

Ceph is based on *RADOS: Reliable Autonomic Distributed Object Store*. RADOS distributes objects across the storage cluster and replicates objects for fault tolerance. RADOS contains the following major components:

- *Object Storage Device (OSD) Daemon*. The storage daemon for the RADOS service, which interacts with the OSD (physical or logical storage unit for your data).

You must run this daemon on each server in your cluster. For each OSD, you can have an associated hard drive disk. For performance purposes, pool your hard drive disk with raid arrays, logical volume management (LVM), or B-tree file system (`Btrfs`) pooling. By default, the following pools are created: data, metadata, and RBD.

- *Meta-Data Server (MDS)*. Stores metadata. MDSs build a POSIX file system on top of objects for Ceph clients. However, if you do not use the Ceph file system, you do not need a metadata server.
- *Monitor (MON)*. A lightweight daemon that handles all communications with external applications and clients. It also provides a consensus for distributed decision making in a Ceph/RADOS cluster. For instance, when you mount a Ceph shared on a client, you point to the address of a MON server. It checks the state and the consistency of the data. In an ideal setup, you must run at least three `ceph-mon` daemons on separate servers.





























































































































































```

volume_driver = cinder.volume.drivers.netapp.common.NetAppDriver
netapp_storage_family = ontap_7mode
netapp_storage_protocol = nfs
netapp_server_hostname = myhostname
netapp_server_port = 80
netapp_login = username
netapp_password = password
nfs_shares_config = /etc/cinder/nfs_shares

```

**Table 2.20. Description of NetApp 7-Mode NFS driver configuration options**

Configuration option = Default value	Description
[DEFAULT]	
expiry_thres_minutes = 720	(IntOpt) 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.
netapp_login = None	(StrOpt) Administrative user account name used to access the storage system or proxy server.
netapp_partner_backend_name = None	(StrOpt) The name of the config.conf stanza for a Data ONTAP (7-mode) HA partner. This option is only used by the driver when connecting to an instance with a storage family of Data ONTAP operating in 7-Mode, and it is required if the storage protocol selected is FC.
netapp_password = None	(StrOpt) Password for the administrative user account specified in the netapp_login option.
netapp_server_hostname = None	(StrOpt) The hostname (or IP address) for the storage system or proxy server.
netapp_server_port = None	(IntOpt) 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.
netapp_storage_family = ontap_cluster	(StrOpt) The storage family type used on the storage system; valid values are ontap_7mode for using Data ONTAP operating in 7-Mode, ontap_cluster for using clustered Data ONTAP, or eseries for using E-Series.
netapp_storage_protocol = None	(StrOpt) The storage protocol to be used on the data path with the storage system.
netapp_transport_type = http	(StrOpt) The transport protocol used when communicating with the storage system or proxy server.
netapp_vfiler = None	(StrOpt) The vFiler unit on which provisioning of block storage volumes will be done. This option is only used by the driver when connecting to an instance with a storage family of Data ONTAP operating in 7-Mode. Only use this option when utilizing the MultiStore feature on the NetApp storage system.
thres_avl_size_perc_start = 20	(IntOpt) 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.
thres_avl_size_perc_stop = 60	(IntOpt) 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 expiry_thres_minutes configuration option.



### Note

Additional NetApp NFS configuration options are shared with the generic NFS driver. For a description of these, see [Table 2.23, “Description of NFS storage configuration options” \[94\]](#).



### Tip

For more information on these options and other deployment and operational scenarios, visit the [NetApp OpenStack Deployment and Operations Guide](#).

## NetApp E-Series storage family

The NetApp E-Series storage family represents a configuration group which provides OpenStack compute instances access to E-Series storage systems. At present it can be configured in OpenStack Block Storage to work with the iSCSI storage protocol.

### NetApp iSCSI configuration for E-Series

The NetApp iSCSI configuration for E-Series is an interface from OpenStack to E-Series storage systems for provisioning and managing the SAN block storage entity; that is, a NetApp LUN which can be accessed using the iSCSI protocol.

The iSCSI configuration for E-Series is an interface from OpenStack Block Storage to the E-Series proxy instance and as such requires the deployment of the proxy instance in order to achieve the desired functionality. The driver uses REST APIs to interact with the E-Series proxy instance, which in turn interacts directly with the E-Series controllers.

The use of multipath and DM-MP are required when using the OpenStack Block Storage driver for E-Series. In order for OpenStack Block Storage and OpenStack Compute to take advantage of multiple paths, the following configuration options must be correctly configured:

- The `use_multipath_for_image_xfer` option should be set to `True` in the `cinder.conf` file within the driver-specific stanza (for example, `[myDriver]`).
- The `iscsi_use_multipath` option should be set to `True` in the `nova.conf` file within the `[libvirt]` stanza.

### Configuration options for E-Series storage family with iSCSI protocol

Configure the volume driver, storage family, and storage protocol to the NetApp unified driver, E-Series, and iSCSI respectively by setting the `volume_driver`, `netapp_storage_family` and `netapp_storage_protocol` options in `cinder.conf` as follows:



Configuration option = Default value	Description
	comma separated list of disk pool names to be used for provisioning.
<code>netapp_transport_type = http</code>	(StrOpt) The transport protocol used when communicating with the storage system or proxy server.
<code>netapp_webservice_path = /devmgr/v2</code>	(StrOpt) 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 <code>netapp_transport_type</code> , <code>netapp_server_hostname</code> , and <code>netapp_server_port</code> options to create the URL used by the driver to connect to the proxy application.



### Tip

For more information on these options and other deployment and operational scenarios, visit the [NetApp OpenStack Deployment and Operations Guide](#).

## Upgrading prior NetApp drivers to the NetApp unified driver

NetApp introduced a new unified block storage driver in Havana for configuring different storage families and storage protocols. This requires defining upgrade path for NetApp drivers which existed in releases prior to Havana. This section covers the upgrade configuration for NetApp drivers to the new unified configuration and a list of deprecated NetApp drivers.

### Upgraded NetApp drivers

This section describes how to update OpenStack Block Storage configuration from a pre-Havana release to the unified driver format.

#### Driver upgrade configuration

1. NetApp iSCSI direct driver for Clustered Data ONTAP in Grizzly (or earlier).

```

volume_driver = cinder.volume.drivers.netapp.iscsi.
NetAppDirectCmodeISCSIDriver
  
```

NetApp unified driver configuration.

```

volume_driver = cinder.volume.drivers.netapp.common.NetAppDriver
netapp_storage_family = ontap_cluster
netapp_storage_protocol = iscsi
  
```

2. NetApp NFS direct driver for Clustered Data ONTAP in Grizzly (or earlier).

```

volume_driver = cinder.volume.drivers.netapp.nfs.NetAppDirectCmodeNfsDriver
  
```

NetApp unified driver configuration.

```

volume_driver = cinder.volume.drivers.netapp.common.NetAppDriver
netapp_storage_family = ontap_cluster
netapp_storage_protocol = nfs
  
```

3. NetApp iSCSI direct driver for Data ONTAP operating in 7-Mode storage controller in Grizzly (or earlier)

```

volume_driver = cinder.volume.drivers.netapp.iscsi.
NetAppDirect7modeISCSIDriver
  
```



## NFS driver

The Network File System (NFS) is a distributed file system protocol originally developed by Sun Microsystems in 1984. An NFS server *exports* one or more of its file systems, known as *shares*. An NFS client can mount these exported shares on its own file system. You can perform file actions on this mounted remote file system as if the file system were local.

### How the NFS driver works

The NFS driver, and other drivers based on it, work quite differently than a traditional block storage driver.

The NFS driver does not actually allow an instance to access a storage device at the block level. Instead, files are created on an NFS share and mapped to instances, which emulates a block device. This works in a similar way to QEMU, which stores instances in the `/var/lib/nova/instances` directory.

### Enable the NFS driver and related options

To use Cinder with the NFS driver, first set the `volume_driver` in `cinder.conf`:

```
volume_driver=cinder.volume.drivers.nfs.NfsDriver
```

The following table contains the options supported by the NFS driver.

**Table 2.23. Description of NFS storage configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>nfs_mount_attempts = 3</code>	(IntOpt) 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.
<code>nfs_mount_options = None</code>	(StrOpt) Mount options passed to the nfs client. See section of the nfs man page for details.
<code>nfs_mount_point_base = \$state_path/mnt</code>	(StrOpt) Base dir containing mount points for nfs shares.
<code>nfs_oversub_ratio = 1.0</code>	(FloatOpt) This will compare the allocated to available space on the volume destination. If the ratio exceeds this number, the destination will no longer be valid.
<code>nfs_shares_config = /etc/cinder/nfs_shares</code>	(StrOpt) File with the list of available nfs shares
<code>nfs_sparsed_volumes = True</code>	(BoolOpt) 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.
<code>nfs_used_ratio = 0.95</code>	(FloatOpt) Percent of ACTUAL usage of the underlying volume before no new volumes can be allocated to the volume destination.



#### Note

As of the Icehouse release, the NFS driver (and other drivers based off it) will attempt to mount shares using version 4.1 of the NFS protocol (including pNFS). If the mount attempt is unsuccessful due to a lack of client or server support, a subsequent mount attempt that requests the default behavior of the `mount.nfs` command will be performed. On most distributions, the default be-



## NFS driver notes

- `cinder-volume` manages the mounting of the NFS shares as well as volume creation on the shares. Keep this in mind when planning your OpenStack architecture. If you have one master NFS server, it might make sense to only have one `cinder-volume` service to handle all requests to that NFS server. However, if that single server is unable to handle all requests, more than one `cinder-volume` service is needed as well as potentially more than one NFS server.
- Because data is stored in a file and not actually on a block storage device, you might not see the same IO performance as you would with a traditional block storage driver. Please test accordingly.
- Despite possible IO performance loss, having volume data stored in a file might be beneficial. For example, backing up volumes can be as easy as copying the volume files.



### Note

Regular IO flushing and syncing still stands.

## ProphetStor Fibre Channel and iSCSI drivers

ProphetStor Fibre Channel and iSCSI drivers add support for ProphetStor Flexvisor through OpenStack Block Storage. ProphetStor Flexvisor enables commodity x86 hardware as software-defined storage leveraging well-proven ZFS for disk management to provide enterprise grade storage services such as snapshots, data protection with different RAID levels, replication, and deduplication.

The `DPLFCDriver` and `DPLISCSIDriver` drivers run volume operations by communicating with the ProphetStor storage system over HTTPS.

## Supported operations

- Create, delete, attach, and detach volumes.
- Create, list, and delete volume snapshots.
- Create a volume from a snapshot.
- Copy an image to a volume.
- Copy a volume to an image.
- Clone a volume.
- Extend a volume.

## Enable the Fibre Channel or iSCSI drivers

The `DPLFCDriver` and `DPLISCSIDriver` are installed with the OpenStack software.

1. Query storage pool id for configure `dpl_pool` of the `cinder.conf`.
  - a. Logon onto the storage system with administrator access.





















```
# API version for the storage system (string value)
#tintri_api_version=v310

# Following options needed for NFS configuration
# File with the list of available nfs shares (string value)
#nfs_shares_config=/etc/cinder/nfs_shares
```

Following changes are needed in `/etc/cinder/nfs_shares`

```
# Edit /etc/cinder/nfs_shares to add one or more Tintri VMstore mount points
# associated with already configured VMstore management IP in cinder.conf
{vmstore_data_ip}:/tintri/{submount1}
{vmstore_data_ip}:/tintri/{submount2}
```

Following changes are needed in `nova.conf` file:

```
# Edit /etc/nova/nova.conf
nfs_mount_options=vers=3
```

**Table 2.29. Description of Tintri driver configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>tintri_server_hostname = None</code>	(StrOpt) Tintri VMstore management IP address
<code>tintri_server_username = None</code>	(StrOpt) Tintri VMstore username
<code>tintri_server_password = None</code>	(StrOpt) Tintri VMstore password for the user
<code>tintri_api_version = v310</code>	(StrOpt) Tintri VMstore API version

## VMware VMDK driver

Use the VMware VMDK driver to enable management of the OpenStack Block Storage volumes on vCenter-managed data stores. Volumes are backed by VMDK files on data stores that use any VMware-compatible storage technology such as NFS, iSCSI, FiberChannel, and vSAN.



### Warning

The VMware ESX VMDK driver is deprecated as of the Icehouse release and might be removed in Juno or a subsequent release. The VMware vCenter VMDK driver continues to be fully supported.

## Functional context

The VMware VMDK driver connects to vCenter, through which it can dynamically access all the data stores visible from the ESX hosts in the managed cluster.

When you create a volume, the VMDK driver creates a VMDK file on demand. The VMDK file creation completes only when the volume is subsequently attached to an instance. The

reason for this requirement is that data stores visible to the instance determine where to place the volume. Before the service creates the VMDK file, attach a volume to the target instance.

The running vSphere VM is automatically reconfigured to attach the VMDK file as an extra disk. Once attached, you can log in to the running vSphere VM to rescan and discover this extra disk.

With the update to ESX version 6.0, the VMDK driver now supports NFS version 4.1.

## Configuration

The recommended volume driver for OpenStack Block Storage is the VMware vCenter VMDK driver. When you configure the driver, you must match it with the appropriate OpenStack Compute driver from VMware and both drivers must point to the same server.

In the `nova.conf` file, use this option to define the Compute driver:

```
compute_driver=vmwareapi.VMwareVCDriver
```

In the `cinder.conf` file, use this option to define the volume driver:

```
volume_driver=cinder.volume.drivers.vmware.vmdk.VMwareVcVmdkDriver
```

The following table lists various options that the drivers support for the OpenStack Block Storage configuration (`cinder.conf`):

**Table 2.30. Description of VMware configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>vmware_api_retry_count = 10</code>	(IntOpt) Number of times VMware ESX/VC server API must be retried upon connection related issues.
<code>vmware_host_ip = None</code>	(StrOpt) IP address for connecting to VMware ESX/VC server.
<code>vmware_host_password = None</code>	(StrOpt) Password for authenticating with VMware ESX/VC server.
<code>vmware_host_username = None</code>	(StrOpt) Username for authenticating with VMware ESX/VC server.
<code>vmware_host_version = None</code>	(StrOpt) Optional string specifying the VMware VC server version. The driver attempts to retrieve the version from VMware VC server. Set this configuration only if you want to override the VC server version.
<code>vmware_image_transfer_timeout_secs = 7200</code>	(IntOpt) Timeout in seconds for VMDK volume transfer between Cinder and Glance.
<code>vmware_max_objects_retrieval = 100</code>	(IntOpt) 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.
<code>vmware_task_poll_interval = 0.5</code>	(FloatOpt) The interval (in seconds) for polling remote tasks invoked on VMware ESX/VC server.
<code>vmware_tmp_dir = /tmp</code>	(StrOpt) Directory where virtual disks are stored during volume backup and restore.
<code>vmware_volume_folder = cinder-volumes</code>	(StrOpt) Name for the folder in the VC datacenter that will contain cinder volumes.





2. Set the extra spec key `vmware:storage_profile` in the desired Block Storage volume types to the policy name that you created in the previous step.
3. Optionally, for the `vmware_host_version` parameter, enter the version number of your vSphere platform. For example, `5.5`.

This setting overrides the default location for the corresponding WSDL file. Among other scenarios, you can use this setting to prevent WSDL error messages during the development phase or to work with a newer version of vCenter.

4. Complete the other vCenter configuration parameters as appropriate.



### Note

The following considerations apply to configuring SPBM for the Block Storage service:

- Any volume that is created without an associated policy (that is to say, without an associated volume type that specifies `vmware:storage_profile` extra spec), there is no policy-based placement for that volume.

## Supported operations

The VMware vCenter and ESX VMDK drivers support these operations:

- Create, delete, attach, and detach volumes.



### Note

When a volume is attached to an instance, a reconfigure operation is performed on the instance to add the volume's VMDK to it. The user must manually rescan and mount the device from within the guest operating system.

- Create, list, and delete volume snapshots.



### Note

Allowed only if volume is not attached to an instance.

- Create a volume from a snapshot.
- Copy an image to a volume.



### Note






Only images in `vmdk` disk format with `bare` container format are supported. The `vmware_disktype` property of the image can be `preallocated`, `sparse`, `streamOptimized` or `thin`.

- Copy a volume to an image.



### Note

- Allowed only if the volume is not attached to an instance.

- This operation creates a `streamOptimized` disk image.
- Clone a volume.
  - 
**Note**  
 Supported only if the source volume is not attached to an instance.
- Backup a volume.
  - 
**Note**  
 This operation creates a backup of the volume in `streamOptimized` disk format.
- Restore backup to new or existing volume.
  - 
**Note**  
 Supported only if the existing volume doesn't contain snapshots.
- Change the type of a volume.
  - 
**Note**  
 This operation is supported only if the volume state is `available`.
- Extend a volume.
  - 
**Note**  
 Although the VMware ESX VMDK driver supports these operations, it has not been extensively tested.

## Storage policy-based configuration in vCenter

You can configure Storage Policy-Based Management (SPBM) profiles for vCenter data stores supporting the Compute, Image Service, and Block Storage components of an OpenStack implementation.

In a vSphere OpenStack deployment, SPBM enables you to delegate several data stores for storage, which reduces the risk of running out of storage space. The policy logic selects the data store based on accessibility and available storage space.

## Prerequisites

- Determine the data stores to be used by the SPBM policy.
- Determine the tag that identifies the data stores in the OpenStack component configuration.
- Create separate policies or sets of data stores for separate OpenStack components.

## Create storage policies in vCenter

### Procedure 2.3. To create storage policies in vCenter

1. In vCenter, create the tag that identifies the data stores:
  - a. From the Home screen, click **Tags**.
  - b. Specify a name for the tag.
  - c. Specify a tag category. For example, `spbm-cinder`.
2. Apply the tag to the data stores to be used by the SPBM policy.



#### Note

For details about creating tags in vSphere, see the [vSphere documentation](#).

3. In vCenter, create a tag-based storage policy that uses one or more tags to identify a set of data stores.



#### Note

For details about creating storage policies in vSphere, see the [vSphere documentation](#).

## Data store selection

If storage policy is enabled, the driver initially selects all the data stores that match the associated storage policy.

If two or more data stores match the storage policy, the driver chooses a data store that is connected to the maximum number of hosts.

In case of ties, the driver chooses the data store with lowest space utilization, where space utilization is defined by the  $(1 - \text{freespace} / \text{totalspace})$  meters.

These actions reduce the number of volume migrations while attaching the volume to instances.

The volume must be migrated if the ESX host for the instance cannot access the data store that contains the volume.

## Windows iSCSI volume driver

Windows Server 2012 and Windows Storage Server 2012 offer an integrated iSCSI Target service that can be used with OpenStack Block Storage in your stack. Being entirely a software solution, consider it in particular for mid-sized networks where the costs of a SAN might be excessive.

The Windows `cinder-volume` driver works with OpenStack Compute on any hypervisor. It includes snapshotting support and the "boot from volume" feature.

This driver creates volumes backed by fixed-type VHD images on Windows Server 2012 and dynamic-type VHDX on Windows Server 2012 R2, stored locally on a user-specified path. The system uses those images as iSCSI disks and exports them through iSCSI targets. Each volume has its own iSCSI target.

This driver has been tested with Windows Server 2012 and Windows Server R2 using the Server and Storage Server distributions.

Install the `cinder-volume` service as well as the required Python components directly on to the Windows node.

You may install and configure `cinder-volume` and its dependencies manually using the following guide or you may use the `Cinder Volume Installer`, presented below.

## Installing using the OpenStack cinder volume installer

In case you want to avoid all the manual setup, you can use Cloudbase Solutions' installer. You can find it at [https://www.cloudbase.it/downloads/CinderVolumeSetup\\_Beta.msi](https://www.cloudbase.it/downloads/CinderVolumeSetup_Beta.msi). It installs an independent Python environment, in order to avoid conflicts with existing applications, dynamically generates a `cinder.conf` file based on the parameters provided by you.

`cinder-volume` will be configured to run as a Windows Service, which can be restarted using:

```
PS C:\> net stop cinder-volume ; net start cinder-volume
```

The installer can also be used in unattended mode. More details about how to use the installer and its features can be found at <https://www.cloudbase.it>

## Windows Server configuration

The required service in order to run `cinder-volume` on Windows is `wintarget`. This will require the iSCSI Target Server Windows feature to be installed. You can install it by running the following command:

```
PS C:\> Add-WindowsFeature  
FS-iSCSITarget-ServerAdd-WindowsFeatureFS-iSCSITarget-Server
```



### Note

The Windows Server installation requires at least 16 GB of disk space. The volumes hosted by this node need the extra space.

For `cinder-volume` to work properly, you must configure NTP as explained in [the section called "Configure NTP" \[250\]](#).

Next, install the requirements as described in [the section called "Requirements" \[253\]](#).

## Getting the code

Git can be used to download the necessary source code. The installer to run Git on Windows can be downloaded here:

<https://github.com/msysgit/msysgit/releases/download/Git-1.9.2-preview20140411/Git-1.9.2-preview20140411.exe>





## Supported operations

- Create, delete, attach, detach, retype, clone, and extend volumes.
- Create a volume from snapshot.
- Create, list, and delete volume snapshots.
- Manage and unmanage a volume.
- Get volume statistics.
- Create a thin provisioned volume.
- Create volumes with QoS specifications.

## Configure X-IO Volume driver

To configure the use of an ISE product with OpenStack Block Storage, modify your `cinder.conf` file as follows. Be careful to use the one that matches the storage protocol in use:

### Fibre Channel

```
volume_driver = cinder.volume.drivers.xio.XIOISEFCDriver  
san_ip = 1.2.3.4 # the address of your ISE REST management interface  
san_login = administrator # your ISE management admin login  
san_password = password # your ISE management admin password
```

### iSCSI

```
volume_driver = cinder.volume.drivers.xio.XIOISEISCSIDriver  
san_ip = 1.2.3.4 # the address of your ISE REST management interface  
san_login = administrator # your ISE management admin login  
san_password = password # your ISE management admin password  
iscsi_ip_address = ionet_ip # ip address to one ISE port connected to the IONET
```

## Optional configuration parameters

Table 2.34. Description of X-IO volume driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
<code>driver_use_ssl = False</code>	(BoolOpt) Tell driver to use SSL for connection to backend storage if the driver supports it.
<code>ise_completion_retries = 30</code>	(IntOpt) Number on retries to get completion status after issuing a command to ISE.
<code>ise_connection_retries = 5</code>	(IntOpt) Number of retries (per port) when establishing connection to ISE management port.
<code>ise_raid = 1</code>	(IntOpt) Raid level for ISE volumes.
<code>ise_retry_interval = 1</code>	(IntOpt) Interval (secs) between retries.
<code>ise_storage_pool = 1</code>	(IntOpt) Default storage pool for volumes.

## Multipath

The X-IO ISE supports a multipath configuration, but multipath must be enabled on the compute node (see *ISE Storage Blade Best Practices Guide*). For more information, see [www.openstack.org](http://www.openstack.org).

## Volume types

OpenStack Block Storage uses volume types to help the administrator specify attributes for volumes. These attributes are called extra-specs. The X-IO volume driver support the following extra-specs.

Table 2.35. Extra specs

Extra-specs name	Valid values	Description
Feature:Raid	1, 5	RAID level for volume.
Feature:Pool	1 - n (n being number of pools on ISE)	Pool to create volume in.
Affinity:Type	cadp, flash, hdd	Volume media affinity type.
Alloc:Type	0 (thick), 1 (thin)	Allocation type for volume. Thick or thin
QoS:minIOPS	n (value less than maxIOPS)	Minimum IOPS setting for volume.
QoS:maxIOPS	n (value bigger than minIOPS)	Maximum IOPS setting for volume.
QoS:burstIOPS	n (value bigger than minIOPS)	Burst IOPS setting for volume.

### Examples

Create a volume type called xio1-flash for volumes that should reside on ssd storage:

```
$ cinder type-create xio1-flash
$ cinder type-key xio1-flash set Affinity:Type=flash
```

Create a volume type called xio1 and set QoS min and max:

```
$ cinder type-create xio1
$ cinder type-key xio1 set QoS:minIOPS=20
$ cinder type-key xio1 set QoS:maxIOPS=5000
```

## Oracle ZFS Storage Appliance NFS driver

The Oracle ZFS Storage Appliance (ZFSSA) NFS driver enables the ZFSSA to be used seamlessly as a block storage resource. The driver enables you to create volumes on a ZFS share that is NFS mounted.

### Requirements

Oracle ZFS Storage Appliance Software version 2013.1.2.0 or later

### Supported operations

- Create, extend, delete volumes
- Attach and detach volumes
- Create, delete snapshots

- Create a volume from a snapshot
- Copy an image to a volume
- Copy a volume to an image
- Clone a volume

## Appliance configuration

Appliance configuration using the command line interface (CLI) is described below. To access the CLI, ensure SSH remote access is enabled, which is the default. You can also perform configuration using the browser user interface (BUI) or the RESTful API. Please refer to the [Oracle ZFS Storage Appliance documentation](#) for details on how to configure the Oracle ZFS Storage Appliance using the BUI, CLI and RESTful API.

1. Log in to the Oracle ZFS Storage Appliance CLI and enable the REST service. REST service needs to stay online for this driver to function.

```
zfssa:>configuration services rest enable
```

2. Create a new storage pool on the appliance if you do not want to use an existing one. This storage pool is named `'mypool'` for the sake of this documentation.
3. Create a new project and share in the storage pool (`mypool`) if you do not want to use existing ones. This driver will create a project and share by the names specified in `cinder.conf`, if the a project or share by that name doesnt already exist in the storage pool (`mypool`). The project and share are named `'NFSPProject'` and `'nfs_share'` in the sample `cinder.conf` entries below.

4. To perform driver operations, create a role with the following authorizations:
  - `scope=svc - allow_administer=true, allow_restart=true, allow_configure=true`
  - `scope=nas - pool=pool_name, project=project_name, share=share_name, allow_clone=true, allow_createProject=true, allow_createShare=true, allow_changeSpaceProps=true, allow_changeGeneralProps=true, allow_destroy=true, allow_rollback=true, allow_takeSnap=true`

The following examples show how to create a role with authorizations.

```
zfssa:> configuration roles
zfssa:configuration roles> role OpenStackRole
zfssa:configuration roles OpenStackRole (uncommitted)> set description=
"OpenStack NFS Cinder Driver"
zfssa:configuration roles OpenStackRole (uncommitted)> commit
zfssa:configuration roles> select OpenStackRole
zfssa:configuration roles OpenStackRole> authorizations create
zfssa:configuration roles OpenStackRole auth (uncommitted)> set scope=svc
zfssa:configuration roles OpenStackRole auth (uncommitted)> set
allow_administer=true
zfssa:configuration roles OpenStackRole auth (uncommitted)> set
allow_restart=true
```

```
zfssa:configuration roles OpenStackRole auth (uncommitted)> set
allow_configure=true
zfssa:configuration roles OpenStackRole auth (uncommitted)> commit
```

```
zfssa:> configuration roles OpenStackRole authorizations> set scope=nas
```

The following properties need to be set when the scope of this role needs to be limited to a pool (`mypool`), a project (`NFSProject`) and a share (`nfs_share`) created in the steps above. This will prevent the user assigned to this role from being used to modify other pools, projects and shares.

```
zfssa:configuration roles OpenStackRole auth (uncommitted)> set pool=
mypool
zfssa:configuration roles OpenStackRole auth (uncommitted)> set project=
NFSProject
zfssa:configuration roles OpenStackRole auth (uncommitted)> set share=
nfs_share
```

The following properties only need to be set when a share or a project has not been created following the steps above and wish to allow the driver to create them for you.

```
zfssa:configuration roles OpenStackRole auth (uncommitted)> set
allow_createProject=true
zfssa:configuration roles OpenStackRole auth (uncommitted)> set
allow_createShare=true
```

```
zfssa:configuration roles OpenStackRole auth (uncommitted)> set
allow_clone=true
zfssa:configuration roles OpenStackRole auth (uncommitted)> set
allow_changeSpaceProps=true
zfssa:configuration roles OpenStackRole auth (uncommitted)> set
allow_destroy=true
zfssa:configuration roles OpenStackRole auth (uncommitted)> set
allow_rollback=true
zfssa:configuration roles OpenStackRole auth (uncommitted)> set
allow_takeSnap=true
zfssa:configuration roles OpenStackRole auth (uncommitted)> commit
```

5. Create a new user or modify an existing one and assign the new role to the user.

The following example shows how to create a new user and assign the new role to the user.

```
zfssa:> configuration users
zfssa:configuration users> user cinder
zfssa:configuration users cinder (uncommitted)> set fullname="OpenStack
Cinder Driver"
zfssa:configuration users cinder (uncommitted)> set initial_password=12345
zfssa:configuration users cinder (uncommitted)> commit
zfssa:configuration users> select cinder set roles=OpenStackRole
```

6. Ensure that NFS and HTTP services on the appliance are online. Note the HTTPS port number for later entry in the cinder service configuration file (`cinder.conf`). This driver uses WebDAV over HTTPS to create snapshots and clones of volumes, and therefore needs to have the HTTP service online.

The following example illustrates enabling the services and showing their properties.

```
zfssa:> configuration services nfs
```

```
zfssa:configuration services nfs> enable
zfssa:configuration services nfs> show
Properties:
<status>= online
...
```

```
zfssa:configuration services http> enable
zfssa:configuration services http> show
Properties:
<status>= online
require_login = true
protocols = http/https
listen_port = 80
https_port = 443
```

7. Create a network interface to be used exclusively for data. An existing network interface may also be used. The following example illustrates how to make a network interface for data traffic flow only.



**Note**

For better performance and reliability, it is recommended to configure a separate subnet exclusively for data traffic in your cloud environment.

```
zfssa:> configuration net interfaces
zfssa:configuration net interfaces> select igbx
zfssa:configuration net interfaces igbx> set admin=false
zfssa:configuration net interfaces igbx> commit
```

8. For clustered controller systems, the following verification is required in addition to the above steps. Skip this step if a standalone system is used.

```
zfssa:> configuration cluster resources list
```

Verify that both the newly created pool and the network interface are of type "singleton" and are not locked to the current controller. This approach ensures that the pool and the interface used for data always belong to the active controller, regardless of the current state of the cluster. Verify that both the network interface used for management and data, and the storage pool belong to the same head.



**Note**

There will be a short service interruption during failback/takeover, but once the process is complete, the driver should be able to access the ZFSSA for data as well as for management.

**Cinder service configuration**

1. Define the following required properties in the `cinder.conf` configuration file:

```
volume_driver = cinder.volume.drivers.zfssa.zfssanfs.ZFSSANFSDriver
san_ip = myhost
san_login = username
san_password = password
zfssa_data_ip = mydata
zfssa_nfs_pool = mypool
```



### Note

Management interface `san_ip` can be used instead of `zfssa_data_ip`, but it is not recommended.

- You can also define the following additional properties in the `cinder.conf` configuration file:

```
zfssa_nfs_project = NFSProject
zfssa_nfs_share = nfs_share
zfssa_nfs_mount_options =
zfssa_nfs_share_compression = off
zfssa_nfs_share_logbias = latency
zfssa_https_port = 443
```



### Note

The driver does not use the file specified in the `nfs_shares_config` option.

## Driver options

The Oracle ZFS Storage Appliance NFS driver supports these options:

**Table 2.36. Description of ZFS Storage Appliance NFS driver configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>zfssa_data_ip = None</code>	(StrOpt) Data path IP address
<code>zfssa_https_port = 443</code>	(StrOpt) HTTPS port number
<code>zfssa_nfs_mount_options =</code>	(StrOpt) Options to be passed while mounting share over nfs
<code>zfssa_nfs_pool =</code>	(StrOpt) Storage pool name.
<code>zfssa_nfs_project = NFSProject</code>	(StrOpt) Project name.
<code>zfssa_nfs_share = nfs_share</code>	(StrOpt) Share name.
<code>zfssa_nfs_share_compression = off</code>	(StrOpt) Data compression.
<code>zfssa_nfs_share_logbias = latency</code>	(StrOpt) Synchronous write bias-latency, throughput.
<code>zfssa_rest_timeout = None</code>	(IntOpt) REST connection timeout. (seconds)

This driver shares additional NFS configuration options with the generic NFS driver. For a description of these, see [Table 2.23, “Description of NFS storage configuration options” \[94\]](#).

## Backup drivers

This section describes how to configure the `cinder-backup` service and its drivers.

The volume drivers are included with the Block Storage repository (<https://git.openstack.org/cgit/openstack/cinder/>). To set a backup driver, use the `backup_driver` flag. By default there is no backup driver enabled.

## Ceph backup driver

The Ceph backup driver backs up volumes of any type to a Ceph back-end store. The driver can also detect whether the volume to be backed up is a Ceph RBD volume, and if so, it tries to perform incremental and differential backups.

For source Ceph RBD volumes, you can perform backups within the same Ceph pool (not recommended). You can also perform backups between different Ceph pools and between different Ceph clusters.

At the time of writing, differential backup support in Ceph/librbd was quite new. This driver attempts a differential backup in the first instance. If the differential backup fails, the driver falls back to full backup/copy.

If incremental backups are used, multiple backups of the same volume are stored as snapshots so that minimal space is consumed in the backup store. It takes far less time to restore a volume than to take a full copy.



### Note

Block Storage enables you to:

- Restore to a new volume, which is the default and recommended action.
- Restore to the original volume from which the backup was taken. The restore action takes a full copy because this is the safest action.

To enable the Ceph backup driver, include the following option in the `cinder.conf` file:

```
backup_driver = cinder.backup.drivers.ceph
```

The following configuration options are available for the Ceph backup driver.

**Table 2.37. Description of Ceph backup driver configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>backup_ceph_chunk_size = 134217728</code>	(IntOpt) The chunk size, in bytes, that a backup is broken into before transfer to the Ceph object store.
<code>backup_ceph_conf = /etc/ceph/ceph.conf</code>	(StrOpt) Ceph configuration file to use.
<code>backup_ceph_pool = backups</code>	(StrOpt) The Ceph pool where volume backups are stored.
<code>backup_ceph_stripe_count = 0</code>	(IntOpt) RBD stripe count to use when creating a backup image.
<code>backup_ceph_stripe_unit = 0</code>	(IntOpt) RBD stripe unit to use when creating a backup image.
<code>backup_ceph_user = cinder</code>	(StrOpt) 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.
<code>restore_discard_excess_bytes = True</code>	(BoolOpt) If True, always discard excess bytes when restoring volumes i.e. pad with zeroes.

This example shows the default options for the Ceph backup driver.























```
# Swift authentication mechanism (string value)
#backup_swift_auth=per_user

# Swift authentication version. Specify "1" for auth 1.0, or
# "2" for auth 2.0 (string value)
#backup_swift_auth_version=1

# Swift tenant/account name. Required when connecting to an
# auth 2.0 system (string value)
#backup_swift_tenant=<None>

# Swift user name (string value)
#backup_swift_user=<None>

# Swift key for authentication (string value)
#backup_swift_key=<None>

# The default Swift container to use (string value)
#backup_swift_container=volumebackups

# The size in bytes of Swift backup objects (integer value)
#backup_swift_object_size=52428800

# The size in bytes that changes are tracked for incremental
# backups. backup_swift_object_size has to be multiple of
# backup_swift_block_size. (integer value)
#backup_swift_block_size=32768

# The number of retries to make for Swift operations (integer
# value)
#backup_swift_retry_attempts=3

# The backoff time in seconds between Swift retries (integer
# value)
#backup_swift_retry_backoff=2

# 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. (boolean value)
#backup_swift_enable_progress_timer=true

#
# Options defined in cinder.backup.drivers.tsm
#

# Volume prefix for the backup id when backing up to TSM
# (string value)
#backup_tsm_volume_prefix=backup

# TSM password for the running username (string value)
#backup_tsm_password=password

# Enable or Disable compression for backups (boolean value)
#backup_tsm_compression=true

#
```

```
# Options defined in cinder.backup.manager
#

# Driver to use for backups. (string value)
# Deprecated group/name - [DEFAULT]/backup_service
#backup_driver=cinder.backup.drivers.swift

#

# Options defined in cinder.cmd.volume
#

# Backend override of host value. (string value)
# Deprecated group/name - [DEFAULT]/host
#backend_host=<None>

#

# Options defined in cinder.cmd.volume_usage_audit
#

# If this option is specified then the start time specified is
# used instead of the start time of the last completed audit
# period. (string value)
#start_time=<None>

# If this option is specified then the end time specified is
# used instead of the end time of the last completed audit
# period. (string value)
#end_time=<None>

# Send the volume and snapshot create and delete notifications
# generated in the specified period. (boolean value)
#send_actions=false

#

# Options defined in cinder.common.config
#

# File name for the paste.deploy config for cinder-api (string
# value)
#api_paste_config=api-paste.ini

# Top-level directory for maintaining cinder's state (string
# value)
# Deprecated group/name - [DEFAULT]/pybasedir
#state_path=/var/lib/cinder

# IP address of this host (string value)
#my_ip=10.0.0.1

# Default glance host name or IP (string value)
#glance_host=$my_ip

# Default glance port (integer value)
#glance_port=9292

# A list of the glance API servers available to cinder
# ([hostname|ip]:port) (list value)
```

```
#glance_api_servers=$glance_host:$glance_port

# Version of the glance API to use (integer value)
#glance_api_version=1

# Number retries when downloading an image from glance
# (integer value)
#glance_num_retries=0

# Allow to perform insecure SSL (https) requests to glance
# (boolean value)
#glance_api_insecure=false

# 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. (boolean
# value)
#glance_api_ssl_compression=false

# Location of ca certificates file to use for glance client
# requests. (string value)
#glance_ca_certificates_file=<None>

# http/https timeout value for glance operations. If no value
# (None) is supplied here, the glanceclient default value is
# used. (integer value)
#glance_request_timeout=<None>

# The topic that scheduler nodes listen on (string value)
#scheduler_topic=cinder-scheduler

# The topic that volume nodes listen on (string value)
#volume_topic=cinder-volume

# The topic that volume backup nodes listen on (string value)
#backup_topic=cinder-backup

# DEPRECATED: Deploy v1 of the Cinder API. (boolean value)
#enable_v1_api=true

# Deploy v2 of the Cinder API. (boolean value)
#enable_v2_api=true

# Enables or disables rate limit of the API. (boolean value)
#api_rate_limit=true

# Specify list of extensions to load when using
# osapi_volume_extension option with
# cinder.api.contrib.select_extensions (list value)
#osapi_volume_ext_list=

# osapi volume extension to load (multi valued)
#osapi_volume_extension=cinder.api.contrib.standard_extensions

# Full class name for the Manager for volume (string value)
#volume_manager=cinder.volume.manager.VolumeManager

# Full class name for the Manager for volume backup (string
# value)
```

```
#backup_manager=cinder.backup.manager.BackupManager

# Full class name for the Manager for scheduler (string value)
#scheduler_manager=cinder.scheduler.manager.SchedulerManager

# Name of this node. This can be an opaque identifier. It is
# not necessarily a host name, FQDN, or IP address. (string
# value)
#host=cinder

# Availability zone of this node (string value)
#storage_availability_zone=nova

# Default availability zone for new volumes. If not set, the
# storage_availability_zone option value is used as the
# default for new volumes. (string value)
#default_availability_zone=<None>

# Default volume type to use (string value)
#default_volume_type=<None>

# Time period for which to generate volume usages. The options
# are hour, day, month, or year. (string value)
#volume_usage_audit_period=month

# Path to the rootwrap configuration file to use for running
# commands as root (string value)
#rootwrap_config=/etc/cinder/rootwrap.conf

# Enable monkey patching (boolean value)
#monkey_patch=false

# List of modules/decorators to monkey patch (list value)
#monkey_patch_modules=

# Maximum time since last check-in for a service to be
# considered up (integer value)
#service_down_time=60

# The full class name of the volume API class to use (string
# value)
#volume_api_class=cinder.volume.api.API

# The full class name of the volume backup API class (string
# value)
#backup_api_class=cinder.backup.api.API

# The strategy to use for auth. Supports noauth, keystone, and
# deprecated. (string value)
#auth_strategy=noauth

# A list of backend names to use. These backend names should
# be backed by a unique [CONFIG] group with its options (list
# value)
#enabled_backends=<None>

# Whether snapshots count against gigabyte quota (boolean
# value)
#no_snapshot_gb_quota=false
```

```
# The full class name of the volume transfer API class (string
# value)
#transfer_api_class=cinder.transfer.api.API

# The full class name of the volume replication API class
# (string value)
#replication_api_class=cinder.replication.api.API

# The full class name of the consistencygroup API class
# (string value)
#consistencygroup_api_class=cinder.consistencygroup.api.API

# OpenStack privileged account username. Used for requests to
# other services (such as Nova) that require an account with
# special rights. (string value)
#os_privileged_user_name=<None>

# Password associated with the OpenStack privileged account.
# (string value)
#os_privileged_user_password=<None>

# Tenant name associated with the OpenStack privileged
# account. (string value)
#os_privileged_user_tenant=<None>

#
# Options defined in cinder.compute
#

# The full class name of the compute API class to use (string
# value)
#compute_api_class=cinder.compute.nova.API

#
# Options defined in cinder.compute.nova
#

# Match this value when searching for nova in the service
# catalog. Format is: separated values of the form:
# <service_type>:<service_name>:<endpoint_type> (string value)
#nova_catalog_info=compute:Compute Service:publicURL

# Same as nova_catalog_info, but for admin endpoint. (string
# value)
#nova_catalog_admin_info=compute:Compute Service:adminURL

# Override service catalog lookup with template for nova
# endpoint e.g. http://localhost:8774/v2/%(project_id)s
# (string value)
#nova_endpoint_template=<None>

# Same as nova_endpoint_template, but for admin endpoint.
# (string value)
#nova_endpoint_admin_template=<None>

# Region name of this node (string value)
#os_region_name=<None>
```

```
# Location of ca certificates file to use for nova client
# requests. (string value)
#nova_ca_certificates_file=<None>

# Allow to perform insecure SSL requests to nova (boolean
# value)
#nova_api_insecure=false

#
# Options defined in cinder.db.api
#

# Services to be added to the available pool on create
# (boolean value)
#enable_new_services=true

# Template string to be used to generate volume names (string
# value)
#volume_name_template=volume-%s

# Template string to be used to generate snapshot names
# (string value)
#snapshot_name_template=snapshot-%s

# Template string to be used to generate backup names (string
# value)
#backup_name_template=backup-%s

#
# Options defined in cinder.db.base
#

# Driver to use for database access (string value)
#db_driver=cinder.db

#
# Options defined in cinder.image.glance
#

# Default core properties of image (list value)
#glance_core_properties=checksum,container_format,disk_format,image_name,
image_id,min_disk,min_ram,name,size

# A list of url schemes that can be downloaded directly via
# the direct_url. Currently supported schemes: [file]. (list
# value)
#allowed_direct_url_schemes=

#
# Options defined in cinder.image.image_utils
#

# Directory used for temporary storage during image conversion
# (string value)
#image_conversion_dir=$state_path/conversion
```







```
#
# Options defined in cinder.scheduler.weights.volume_number
#
# Multiplier used for weighing volume number. Negative numbers
# mean to spread vs stack. (floating point value)
#volume_number_multiplier=-1.0

#
# Options defined in cinder.transfer.api
#
# The number of characters in the salt. (integer value)
#volume_transfer_salt_length=8

# The number of characters in the autogenerated auth key.
# (integer value)
#volume_transfer_key_length=16

#
# Options defined in cinder.volume.api
#
# Cache volume availability zones in memory for the provided
# duration in seconds (integer value)
#az_cache_duration=3600

# Create volume from snapshot at the host where snapshot
# resides (boolean value)
#snapshot_same_host=true

# Ensure that the new volumes are the same AZ as snapshot or
# source volume (boolean value)
#cloned_volume_same_az=true

#
# Options defined in cinder.volume.driver
#
# The maximum number of times to rescan iSER target to find
# volume (integer value)
#num_iser_scan_tries=3

# This option is deprecated and unused. It will be removed in
# the Liberty release. (integer value)
#iser_num_targets=<None>

# Prefix for iSER volumes (string value)
#iser_target_prefix=iqn.2010-10.org.openstack:

# The IP address that the iSER daemon is listening on (string
# value)
#iser_ip_address=$my_ip

# The port that the iSER daemon is listening on (integer
```





```
# The path to the client certificate for verification, if the
# driver supports it. (string value)
#driver_client_cert=<None>

# Tell driver to use SSL for connection to backend storage if
# the driver supports it. (boolean value)
#driver_use_ssl=false

# Float representation of the over subscription ratio when
# thin provisioning is involved. 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. A ratio lower than 1.0
# will be ignored and the default value will be used instead.
# (floating point value)
#max_over_subscription_ratio=20.0

# Certain ISCSI targets have predefined target names, SCST
# target driver uses this name. (string value)
#scst_target_igsn_name=<None>

# SCST target implementation can choose from multiple SCST
# target drivers. (string value)
#scst_target_driver=iscsi

# Option to enable/disable CHAP authentication for targets.
# (boolean value)
# Deprecated group/name - [DEFAULT]/eqlx_use_chap
#use_chap_auth=false

# CHAP user name. (string value)
# Deprecated group/name - [DEFAULT]/eqlx_chap_login
#chap_username=

# Password for specified CHAP account name. (string value)
# Deprecated group/name - [DEFAULT]/eqlx_chap_password
#chap_password=

# Namespace for driver private data values to be saved in.
# (string value)
#driver_data_namespace=<None>

# 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. (string value)
#filter_function=<None>

# 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.
# (string value)
#goodness_function=<None>

#
# Options defined in cinder.volume.drivers.block_device
#
```



```
#
# Options defined in cinder.volume.drivers.dell.dell_storagecenter_common
#

# Storage Center System Serial Number (integer value)
#dell_sc_ssn=64702

# Dell API port (integer value)
#dell_sc_api_port=3033

# Name of the server folder to use on the Storage Center
# (string value)
#dell_sc_server_folder=openstack

# Name of the volume folder to use on the Storage Center
# (string value)
#dell_sc_volume_folder=openstack

#
# Options defined in cinder.volume.drivers.emc.emc_vmax_common
#

# use this file for cinder emc plugin config data (string
# value)
#cinder_emc_config_file=/etc/cinder/cinder_emc_config.xml

#
# Options defined in cinder.volume.drivers.emc.emc_vnx_cli
#

# VNX authentication scope type. (string value)
#storage_vnx_authentication_type=global

# Directory path that contains the VNX security file. Make
# sure the security file is generated first. (string value)
#storage_vnx_security_file_dir=<None>

# Naviseccli Path. (string value)
#naviseccli_path=

# Storage pool name. (string value)
#storage_vnx_pool_name=<None>

# VNX secondary SP IP Address. (string value)
#san_secondary_ip=<None>

# 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.
# (integer value)
#default_timeout=525600

# Default max number of LUNs in a storage group. By default,
# the value is 255. (integer value)
#max_luns_per_storage_group=255
```

```
# To destroy storage group when the last LUN is removed from
# it. By default, the value is False. (boolean value)
#destroy_empty_storage_group=false

# Mapping between hostname and its iSCSI initiator IP
# addresses. (string value)
#iscsi_initiators=

# Automatically register initiators. By default, the value is
# False. (boolean value)
#initiator_auto_registration=false

# Automatically deregister initiators after the related
# storage group is destroyed. By default, the value is False.
# (boolean value)
#initiator_auto_deregistration=false

# Report free_capacity_gb as 0 when the limit to maximum
# number of pool LUNs is reached. By default, the value is
# False. (boolean value)
#check_max_pool_luns_threshold=false

# Delete a LUN even if it is in Storage Groups.(boolean
# value)
#force_delete_lun_in_storagegroup=false

#
# Options defined in cinder.volume.drivers.emc.xtremio
#

# XMS cluster id in multi-cluster environment (string value)
#xtremio_cluster_name=

#
# Options defined in cinder.volume.drivers.eqlx
#

# Group name to use for creating volumes. Defaults to
# "group-0". (string value)
#eqlx_group_name=group-0

# Timeout for the Group Manager cli command execution. Default
# is 30. (integer value)
#eqlx_cli_timeout=30

# Maximum retry count for reconnection. Default is 5. (integer
# value)
#eqlx_cli_max_retries=5

# Use CHAP authentication for targets. Note that this option
# is deprecated in favour of "use_chap_auth" as specified in
# cinder/volume/driver.py and will be removed in next release.
# (boolean value)
#eqlx_use_chap=false

# Existing CHAP account name. Note that this option is
# deprecated in favour of "chap_username" as specified in
```





```
#hds_hnas_nfs_config_file=/opt/hds/hnas/cinder_nfs_conf.xml

#
# Options defined in cinder.volume.drivers.hitachi.hbsd_common
#

# Serial number of storage system (string value)
#hitachi_serial_number=<None>

# Name of an array unit (string value)
#hitachi_unit_name=<None>

# Pool ID of storage system (integer value)
#hitachi_pool_id=<None>

# Thin pool ID of storage system (integer value)
#hitachi_thin_pool_id=<None>

# Range of logical device of storage system (string value)
#hitachi_ldev_range=<None>

# Default copy method of storage system (string value)
#hitachi_default_copy_method=FULL

# Copy speed of storage system (integer value)
#hitachi_copy_speed=3

# Interval to check copy (integer value)
#hitachi_copy_check_interval=3

# Interval to check copy asynchronously (integer value)
#hitachi_async_copy_check_interval=10

# Control port names for HostGroup or iSCSI Target (string
# value)
#hitachi_target_ports=<None>

# Range of group number (string value)
#hitachi_group_range=<None>

# Request for creating HostGroup or iSCSI Target (boolean
# value)
#hitachi_group_request=false

#
# Options defined in cinder.volume.drivers.hitachi.hbsd_fc
#

# Request for FC Zone creating HostGroup (boolean value)
#hitachi_zoning_request=false

#
# Options defined in cinder.volume.drivers.hitachi.hbsd_horcm
#

# Instance numbers for HORCM (string value)
#hitachi_horcm_numbers=200,201
```

```
# Username of storage system for HORCM (string value)
#hitachi_horcm_user=<None>

# Password of storage system for HORCM (string value)
#hitachi_horcm_password=<None>

# Add to HORCM configuration (boolean value)
#hitachi_horcm_add_conf=true

#
# Options defined in cinder.volume.drivers.hitachi.hbsd_iscsi
#

# Add CHAP user (boolean value)
#hitachi_add_chap_user=false

# iSCSI authentication method (string value)
#hitachi_auth_method=<None>

# iSCSI authentication username (string value)
#hitachi_auth_user=HBSD-CHAP-user

# iSCSI authentication password (string value)
#hitachi_auth_password=HBSD-CHAP-password

#
# Options defined in cinder.volume.drivers.huawei
#

# The configuration file for the Cinder Huawei driver (string
# value)
#cinder_huawei_conf_file=/etc/cinder/cinder_huawei_conf.xml

#
# Options defined in cinder.volume.drivers.ibm.flashsystem
#

# Connection protocol should be FC. (string value)
#flashsystem_connection_protocol=FC

# Connect with multipath (FC only). (boolean value)
#flashsystem_multipath_enabled=false

# Allows vdisk to multi host mapping. (boolean value)
#flashsystem_multihostmap_enabled=true

#
# Options defined in cinder.volume.drivers.ibm.gpfs
#

# Specifies the path of the GPFS directory where Block Storage
# volume and snapshot files are stored. (string value)
#gpfs_mount_point_base=<None>

# Specifies the path of the Image service repository in GPFS.
```

```
# Leave undefined if not storing images in GPFS. (string
# value)
#gpfs_images_dir=<None>

# 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. (string value)
#gpfs_images_share_mode=<None>

# 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. (integer value)
#gpfs_max_clone_depth=0

# 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. (boolean value)
#gpfs_sparse_volumes=true

# Specifies the storage pool that volumes are assigned to. By
# default, the system storage pool is used. (string value)
#gpfs_storage_pool=system

#
# Options defined in cinder.volume.drivers.ibm.ibmnas
#
# IBMNAS platform type to be used as backend storage; valid
# values are - v7ku : for using IBM Storwize V7000 Unified,
# sonas : for using IBM Scale Out NAS, gpfs-nas : for using
# NFS based IBM GPFS deployments. (string value)
#ibmnas_platform_type=v7ku

#
# Options defined in cinder.volume.drivers.ibm.storwize_svc
#
# Storage system storage pool for volumes (string value)
#storwize_svc_volpool_name=volpool

# Storage system space-efficiency parameter for volumes
# (percentage) (integer value)
#storwize_svc_vol_rsize=2

# Storage system threshold for volume capacity warnings
# (percentage) (integer value)
#storwize_svc_vol_warning=0

# Storage system autoexpand parameter for volumes (True/False)
```



```

#xiv_ds8k_connection_type=iscsi

# CHAP authentication mode, effective only for iscsi
# (disabled|enabled) (string value)
#xiv_chap=disabled

#
# Options defined in cinder.volume.drivers.lvm
#

# Name for the VG that will contain exported volumes (string
# value)
#volume_group=cinder-volumes

# If >0, create LVs with multiple mirrors. Note that this
# requires lvm_mirrors + 2 PVs with available space (integer
# value)
#lvm_mirrors=0

# Type of LVM volumes to deploy (string value)
#lvm_type=default

# 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 'None' to not use a conf file even if one
# exists). (string value)
#lvm_conf_file=/etc/cinder/lvm.conf

#
# Options defined in cinder.volume.drivers.netapp.options
#

# The vFiler unit on which provisioning of block storage
# volumes will be done. This option is only used by the driver
# when connecting to an instance with a storage family of Data
# ONTAP operating in 7-Mode. Only use this option when
# utilizing the MultiStore feature on the NetApp storage
# system. (string value)
#netapp_vfiler=<None>

# The name of the config.conf stanza for a Data ONTAP (7-mode)
# HA partner. This option is only used by the driver when
# connecting to an instance with a storage family of Data
# ONTAP operating in 7-Mode, and it is required if the storage
# protocol selected is FC. (string value)
#netapp_partner_backend_name=<None>

# Administrative user account name used to access the storage
# system or proxy server. (string value)
#netapp_login=<None>

# Password for the administrative user account specified in
# the netapp_login option. (string value)
#netapp_password=<None>

# This option specifies the virtual storage server (Vserver)
# name on the storage cluster on which provisioning of block
# storage volumes should occur. (string value)

```

```
#netapp_vserver=<None>

# The hostname (or IP address) for the storage system or proxy
# server. (string value)
#netapp_server_hostname=<None>

# 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. (integer value)
#netapp_server_port=<None>

# 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. (string value)
#netapp_webservice_path=/devmgr/v2

# This option is only utilized when the storage family is
# configured to eseries. 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.
# (string value)
#netapp_controller_ips=<None>

# Password for the NetApp E-Series storage array. (string
# value)
#netapp_sa_password=<None>

# This option is used to restrict provisioning to the
# specified storage pools. Only dynamic disk pools are
# currently supported. Specify the value of this option to be
# a comma separated list of disk pool names to be used for
# provisioning. (string value)
#netapp_storage_pools=<None>

# This option is used to define how the controllers in the
# E-Series storage array will work with the particular
# operating system on the hosts that are connected to it.
# (string value)
#netapp_eseries_host_type=linux_dm_mp

# 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. (integer value)
#thres_avl_size_perc_start=20

# 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 expiry_thres_minutes configuration option.
# (integer value)
#thres_avl_size_perc_stop=60

# 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. (integer value)
#expiry_thres_minutes=720

# 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 cinder-volume
# process to execute the file. (string value)
#netapp_copyoffload_tool_path=<None>

# 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.
# (floating point value)
#netapp_size_multiplier=1.2

# This option is only utilized when the storage protocol is
# configured to use iSCSI or FC. This option is used to
# restrict provisioning to the specified controller volumes.
# Specify the value of this option to be a comma separated
# list of NetApp controller volume names to be used for
# provisioning. (string value)
#netapp_volume_list=<None>

# The storage family type used on the storage system; valid
# values are ontap_7mode for using Data ONTAP operating in
# 7-Mode, ontap_cluster for using clustered Data ONTAP, or
# eseries for using E-Series. (string value)
#netapp_storage_family=ontap_cluster

# The storage protocol to be used on the data path with the
# storage system. (string value)
#netapp_storage_protocol=<None>

# The transport protocol used when communicating with the
# storage system or proxy server. (string value)
#netapp_transport_type=http

#
# Options defined in cinder.volume.drivers.nfs
#

# File with the list of available nfs shares (string value)
#nfs_shares_config=/etc/cinder/nfs_shares

# 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. (boolean value)
#nfs_sparsed_volumes=true

# Percent of ACTUAL usage of the underlying volume before no
# new volumes can be allocated to the volume destination.
# (floating point value)
#nfs_used_ratio=0.95

# This will compare the allocated to available space on the
# volume destination. If the ratio exceeds this number, the
```



```
# destination will no longer be valid. (floating point value)
#nfs_oversub_ratio=1.0

# Base dir containing mount points for nfs shares. (string
# value)
#nfs_mount_point_base=$state_path/mnt

# Mount options passed to the nfs client. See section of the
# nfs man page for details. (string value)
#nfs_mount_options=<None>

# 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. (integer value)
#nfs_mount_attempts=3

#
# Options defined in cinder.volume.drivers.nimble
#

# Nimble Controller pool name (string value)
#nimble_pool_name=default

# Nimble Subnet Label (string value)
#nimble_subnet_label=*

#
# Options defined in cinder.volume.drivers.openvstorage
#

# Vpool to use for volumes - backend is defined by vpool not
# by us. (string value)
#vpool_name=

#
# Options defined in cinder.volume.drivers.prophetstor.options
#

# DPL pool uuid in which DPL volumes are stored. (string
# value)
#dpl_pool=

# DPL port number. (integer value)
#dpl_port=8357

#
# Options defined in cinder.volume.drivers.pure
#

# REST API authorization token. (string value)
#pure_api_token=<None>

#
# Options defined in cinder.volume.drivers.quobyte
#
```

```
# URL to the Quobyte volume e.g., quobyte://<DIR host>/<volume
# name> (string value)
#quobyte_volume_url=<None>

# Path to a Quobyte Client configuration file. (string value)
#quobyte_client_cfg=<None>

# Create volumes as sparse 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. (boolean value)
#quobyte_sparsed_volumes=true

# Create volumes as QCOW2 files rather than raw files.
# (boolean value)
#quobyte_qcow2_volumes=true

# Base dir containing the mount point for the Quobyte volume.
# (string value)
#quobyte_mount_point_base=$state_path/mnt

#
# Options defined in cinder.volume.drivers.rbd
#

# The RADOS pool where rbd volumes are stored (string value)
#rbd_pool=rbd

# The RADOS client name for accessing rbd volumes - only set
# when using cephx authentication (string value)
#rbd_user=<None>

# Path to the ceph configuration file (string value)
#rbd_ceph_conf=

# Flatten volumes created from snapshots to remove dependency
# from volume to snapshot (boolean value)
#rbd_flatten_volume_from_snapshot=false

# The libvirt uuid of the secret for the rbd_user volumes
# (string value)
#rbd_secret_uuid=<None>

# Directory where temporary image files are stored when the
# volume driver does not write them directly to the volume.
# Warning: this option is now deprecated, please use
# image_conversion_dir instead. (string value)
#volume_tmp_dir=<None>

# Maximum number of nested volume clones that are taken before
# a flatten occurs. Set to 0 to disable cloning. (integer
# value)
#rbd_max_clone_depth=5

# Volumes will be chunked into objects of this size (in
# megabytes). (integer value)
#rbd_store_chunk_size=4

# Timeout value (in seconds) used when connecting to ceph
```

```
# cluster. If value < 0, no timeout is set and default
# librados value is used. (integer value)
#rados_connect_timeout=-1

#
# Options defined in cinder.volume.drivers.remotefs
#

# IP address or Hostname of NAS system. (string value)
#nas_ip=

# User name to connect to NAS system. (string value)
#nas_login=admin

# Password to connect to NAS system. (string value)
#nas_password=

# SSH port to use to connect to NAS system. (integer value)
#nas_ssh_port=22

# Filename of private key to use for SSH authentication.
# (string value)
#nas_private_key=

# 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. (string value)
#nas_secure_file_operations=auto

# 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. (string value)
#nas_secure_file_permissions=auto

# Path to the share to use for storing Cinder volumes. For
# example: "/srv/export1" for an NFS server export available
# at 10.0.5.10:/srv/export1 . (string value)
#nas_share_path=

# Options used to mount the storage backend file system where
# Cinder volumes are stored. (string value)
#nas_mount_options=<None>

#
# Options defined in cinder.volume.drivers.san.hp.hp_3par_common
#

# 3PAR WSAPI Server Url like https://<3par ip>:8080/api/v1
# (string value)
#hp3par_api_url=
```



```
# Use thin provisioning for SAN volumes? (boolean value)
#san_thin_provision=true

# IP address of SAN controller (string value)
#san_ip=

# Username for SAN controller (string value)
#san_login=admin

# Password for SAN controller (string value)
#san_password=

# Filename of private key to use for SSH authentication
# (string value)
#san_private_key=

# Cluster name to use for creating volumes (string value)
#san_clustername=

# SSH port to use with SAN (integer value)
#san_ssh_port=22

# Execute commands locally instead of over SSH; use if the
# volume service is running on the SAN device (boolean value)
#san_is_local=false

# SSH connection timeout in seconds (integer value)
#ssh_conn_timeout=30

# Minimum ssh connections in the pool (integer value)
#ssh_min_pool_conn=1

# Maximum ssh connections in the pool (integer value)
#ssh_max_pool_conn=5

#
# Options defined in cinder.volume.drivers.scality
#
# Path or URL to Scality SOFS configuration file (string
# value)
#scality_sofs_config=<None>

# Base dir where Scality SOFS shall be mounted (string value)
#scality_sofs_mount_point=$state_path/scality

# Path from Scality SOFS root to volume dir (string value)
#scality_sofs_volume_dir=cinder/volumes

#
# Options defined in cinder.volume.drivers.smbfs
#
# File with the list of available smbfs shares. (string value)
#smbfs_shares_config=/etc/cinder/smbfs_shares

# Default format that will be used when creating volumes if no
```



```
#
# Comma-separated list of REST servers IP to connect to. (eg
# http://IP1/,http://IP2:81/path (string value)
#srb_base_urls=<None>

#
# Options defined in cinder.volume.drivers.violin.v6000_common
#

# IP address or hostname of mg-a (string value)
#gateway_mga=<None>

# IP address or hostname of mg-b (string value)
#gateway_mgb=<None>

# Use igroups to manage targets and initiators (boolean value)
#use_igroups=false

# Global backend request timeout, in seconds (integer value)
#request_timeout=300

#
# Options defined in cinder.volume.drivers.vmware.vmdk
#

# IP address for connecting to VMware ESX/VC server. (string
# value)
#vmware_host_ip=<None>

# Username for authenticating with VMware ESX/VC server.
# (string value)
#vmware_host_username=<None>

# Password for authenticating with VMware ESX/VC server.
# (string value)
#vmware_host_password=<None>

# Optional VIM service WSDL Location e.g
# http://<server>/vimService.wsdl. Optional over-ride to
# default location for bug work-arounds. (string value)
#vmware_wsdl_location=<None>

# Number of times VMware ESX/VC server API must be retried
# upon connection related issues. (integer value)
#vmware_api_retry_count=10

# The interval (in seconds) for polling remote tasks invoked
# on VMware ESX/VC server. (floating point value)
#vmware_task_poll_interval=0.5

# Name for the folder in the VC datacenter that will contain
# cinder volumes. (string value)
#vmware_volume_folder=cinder-volumes

# Timeout in seconds for VMDK volume transfer between Cinder
# and Glance. (integer value)
#vmware_image_transfer_timeout_secs=7200
```

```
# 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. (integer value)
#vmware_max_objects_retrieval=100

# Optional string specifying the VMware VC server version. The
# driver attempts to retrieve the version from VMware VC
# server. Set this configuration only if you want to override
# the VC server version. (string value)
#vmware_host_version=<None>

# Directory where virtual disks are stored during volume
# backup and restore. (string value)
#vmware_tmp_dir=/tmp

#
# Options defined in cinder.volume.drivers.windows.windows
#

# Path to store VHD backed volumes (string value)
#windows_iscsi_lun_path=C:\iSCSIVirtualDisks

#
# Options defined in cinder.volume.drivers.xio
#

# Default storage pool for volumes. (integer value)
#ise_storage_pool=1

# Raid level for ISE volumes. (integer value)
#ise_raid=1

# Number of retries (per port) when establishing connection to
# ISE management port. (integer value)
#ise_connection_retries=5

# Interval (secs) between retries. (integer value)
#ise_retry_interval=1

# Number on retries to get completion status after issuing a
# command to ISE. (integer value)
#ise_completion_retries=30

#
# Options defined in cinder.volume.drivers.zfssa.zfssanfs
#

# Data path IP address (string value)
#zfssa_data_ip=<None>

# HTTPS port number (string value)
#zfssa_https_port=443

# Options to be passed while mounting share over nfs (string
# value)
```







```

#use_tpool=false

[fc-zone-manager]

#
# Options defined in cinder.zonemanager.drivers.brocade.brcd_fc_zone_driver
#

# Southbound connector for zoning operation (string value)
#brcd_sb_connector=cinder.zonemanager.drivers.brocade.brcd_fc_zone_client_cli.
BrcdFCZoneClientCLI

#
# Options defined in cinder.zonemanager.drivers.cisco.cisco_fc_zone_driver
#

# Southbound connector for zoning operation (string value)
#cisco_sb_connector=cinder.zonemanager.drivers.cisco.cisco_fc_zone_client_cli.
CiscoFCZoneClientCLI

#
# Options defined in cinder.zonemanager.fc_zone_manager
#

# FC Zone Driver responsible for zone management (string
# value)
#zone_driver=cinder.zonemanager.drivers.brocade.brcd_fc_zone_driver.
BrcdFCZoneDriver

# Zoning policy configured by user; valid values include
# "initiator-target" or "initiator" (string value)
#zoning_policy=initiator-target

# 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 (string value)
#fc_fabric_names=<None>

# FC SAN Lookup Service (string value)
#fc_san_lookup_service=cinder.zonemanager.drivers.brocade.
brcd_fc_san_lookup_service.BrcdFCSanLookupService

[keymgr]

#
# Options defined in cinder.keymgr
#

# The full class name of the key manager API class (string
# value)
#api_class=cinder.keymgr.conf_key_mgr.ConfKeyManager

#
# Options defined in cinder.keymgr.conf_key_mgr
#

```

```

# Fixed key returned by key manager, specified in hex (string
# value)
#fixed_key=<None>

#
# Options defined in cinder.keymgr.key_mgr
#

# Authentication url for encryption service. (string value)
#encryption_auth_url=http://localhost:5000/v3

# Url for encryption service. (string value)
#encryption_api_url=http://localhost:9311/v1

[keystone_authtoken]

#
# Options defined in keystonemiddleware.auth_token
#

# Complete public Identity API endpoint. (string value)
#auth_uri=<None>

# API version of the admin Identity API endpoint. (string
# value)
#auth_version=<None>

# Do not handle authorization requests within the middleware,
# but delegate the authorization decision to downstream WSGI
# components. (boolean value)
#delay_auth_decision=false

# Request timeout value for communicating with Identity API
# server. (integer value)
#http_connect_timeout=<None>

# How many times are we trying to reconnect when communicating
# with Identity API Server. (integer value)
#http_request_max_retries=3

# Env key for the swift cache. (string value)
#cache=<None>

# Required if identity server requires client certificate
# (string value)
#certfile=<None>

# Required if identity server requires client certificate
# (string value)
#keyfile=<None>

# A PEM encoded Certificate Authority to use when verifying
# HTTPS connections. Defaults to system CAs. (string value)
#cafile=<None>

# Verify HTTPS connections. (boolean value)
#insecure=false

```

```
# Directory used to cache files related to PKI tokens. (string
# value)
#signing_dir=<None>

# Optionally specify a list of memcached server(s) to use for
# caching. If left undefined, tokens will instead be cached
# in-process. (list value)
# Deprecated group/name - [DEFAULT]/memcache_servers
#memcached_servers=<None>

# 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. (integer value)
#token_cache_time=300

# 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. (integer
# value)
#revocation_cache_time=10

# (Optional) If defined, indicate whether token data should be
# authenticated or authenticated and encrypted. Acceptable
# values are MAC or ENCRYPT. 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. (string value)
#memcache_security_strategy=<None>

# (Optional, mandatory if memcache_security_strategy is
# defined) This string is used for key derivation. (string
# value)
#memcache_secret_key=<None>

# (Optional) Number of seconds memcached server is considered
# dead before it is tried again. (integer value)
#memcache_pool_dead_retry=300

# (Optional) Maximum total number of open connections to every
# memcached server. (integer value)
#memcache_pool_maxsize=10

# (Optional) Socket timeout in seconds for communicating with
# a memcache server. (integer value)
#memcache_pool_socket_timeout=3

# (Optional) Number of seconds a connection to memcached is
# held unused in the pool before it is closed. (integer value)
#memcache_pool_unused_timeout=60

# (Optional) Number of seconds that an operation will wait to
# get a memcache client connection from the pool. (integer
# value)
#memcache_pool_conn_get_timeout=10

# (Optional) Use the advanced (eventlet safe) memcache client
```

```

# pool. The advanced pool will only work under python 2.x.
# (boolean value)
#memcache_use_advanced_pool=false

# (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. (boolean value)
#include_service_catalog=true

# 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. (string
# value)
#enforce_token_bind=permissive

# If true, the revocation list will be checked for cached
# tokens. This requires that PKI tokens are configured on the
# identity server. (boolean value)
#check_revocations_for_cached=false

# 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.
# (list value)
#hash_algorithms=md5

[matchmaker_redis]

#
# Options defined in oslo.messaging
#

# Host to locate redis. (string value)
#host=127.0.0.1

# Use this port to connect to redis host. (integer value)
#port=6379

# Password for Redis server (optional). (string value)
#password=<None>

[matchmaker_ring]

#
# Options defined in oslo.messaging
#

```

```
# Matchmaker ring file (JSON). (string value)
# Deprecated group/name - [DEFAULT]/matchmaker_ringfile
#ringfile=/etc/oslo/matchmaker_ring.json

[oslo_messaging_amqp]

#
# Options defined in oslo.messaging
#

# address prefix used when sending to a specific server
# (string value)
#server_request_prefix=exclusive

# address prefix used when broadcasting to all servers (string
# value)
#broadcast_prefix=broadcast

# address prefix when sending to any server in group (string
# value)
#group_request_prefix=unicast

# Name for the AMQP container (string value)
#container_name=<None>

# Timeout for inactive connections (in seconds) (integer
# value)
#idle_timeout=0

# Debug: dump AMQP frames to stdout (boolean value)
#trace=false

# CA certificate PEM file to verify server certificate
# (string value)
#ssl_ca_file=

# Identifying certificate PEM file to present to clients
# (string value)
#ssl_cert_file=

# Private key PEM file used to sign cert_file certificate
# (string value)
#ssl_key_file=

# Password for decrypting ssl_key_file (if encrypted) (string
# value)
#ssl_key_password=<None>

# Accept clients using either SSL or plain TCP (boolean value)
#allow_insecure_clients=false

[oslo_messaging_qpid]

#
# Options defined in oslo.messaging
#
```

```
# Use durable queues in AMQP. (boolean value)
# Deprecated group/name - [DEFAULT]/rabbit_durable_queues
#amqp_durable_queues=false

# Auto-delete queues in AMQP. (boolean value)
#amqp_auto_delete=false

# Size of RPC connection pool. (integer value)
#rpc_conn_pool_size=30

# Qpid broker hostname. (string value)
#qpid_hostname=localhost

# Qpid broker port. (integer value)
#qpid_port=5672

# Qpid HA cluster host:port pairs. (list value)
#qpid_hosts=$qpid_hostname:$qpid_port

# Username for Qpid connection. (string value)
#qpid_username=

# Password for Qpid connection. (string value)
#qpid_password=

# Space separated list of SASL mechanisms to use for auth.
# (string value)
#qpid_sasl_mechanisms=

# Seconds between connection keepalive heartbeats. (integer
# value)
#qpid_heartbeat=60

# Transport to use, either 'tcp' or 'ssl'. (string value)
#qpid_protocol=tcp

# Whether to disable the Nagle algorithm. (boolean value)
#qpid_tcp_nodelay=true

# The number of prefetched messages held by receiver. (integer
# value)
#qpid_receiver_capacity=1

# The qpid topology version to use. Version 1 is what was
# originally used by impl_qpid. Version 2 includes some
# backwards-incompatible changes that allow broker federation
# to work. Users should update to version 2 when they are
# able to take everything down, as it requires a clean break.
# (integer value)
#qpid_topology_version=1

[oslo_messaging_rabbit]

#
# Options defined in oslo.messaging
#

# Use durable queues in AMQP. (boolean value)
# Deprecated group/name - [DEFAULT]/rabbit_durable_queues
```



```
#amqp_durable_queues=false

# Auto-delete queues in AMQP. (boolean value)
#amqp_auto_delete=false

# Size of RPC connection pool. (integer value)
#rpc_conn_pool_size=30

# 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. (string value)
#kombu_ssl_version=

# SSL key file (valid only if SSL enabled). (string value)
#kombu_ssl_keyfile=

# SSL cert file (valid only if SSL enabled). (string value)
#kombu_ssl_certfile=

# SSL certification authority file (valid only if SSL
# enabled). (string value)
#kombu_ssl_ca_certs=

# How long to wait before reconnecting in response to an AMQP
# consumer cancel notification. (floating point value)
#kombu_reconnect_delay=1.0

# The RabbitMQ broker address where a single node is used.
# (string value)
#rabbit_host=localhost

# The RabbitMQ broker port where a single node is used.
# (integer value)
#rabbit_port=5672

# RabbitMQ HA cluster host:port pairs. (list value)
#rabbit_hosts=$rabbit_host:$rabbit_port

# Connect over SSL for RabbitMQ. (boolean value)
#rabbit_use_ssl=false

# The RabbitMQ userid. (string value)
#rabbit_userid=guest

# The RabbitMQ password. (string value)
#rabbit_password=guest

# The RabbitMQ login method. (string value)
#rabbit_login_method=AMQPLAIN

# The RabbitMQ virtual host. (string value)
#rabbit_virtual_host=/

# How frequently to retry connecting with RabbitMQ. (integer
# value)
#rabbit_retry_interval=1

# How long to backoff for between retries when connecting to
# RabbitMQ. (integer value)
#rabbit_retry_backoff=2
```

```

# Maximum number of RabbitMQ connection retries. Default is 0
# (infinite retry count). (integer value)
#rabbit_max_retries=0

# Use HA queues in RabbitMQ (x-ha-policy: all). If you change
# this option, you must wipe the RabbitMQ database. (boolean
# value)
#rabbit_ha_queues=false

# Number of seconds after which the Rabbit broker is
# considered down if heartbeat's keep-alive fails (0 disables
# the heartbeat, >0 enables it. Enabling heartbeats requires
# kombu>=3.0.7 and amqp>=1.4.0). EXPERIMENTAL (integer value)
#heartbeat_timeout_threshold=0

# How often times during the heartbeat_timeout_threshold we
# check the heartbeat. (integer value)
#heartbeat_rate=2

# Deprecated, use rpc_backend=kombu+memory or rpc_backend=fake
# (boolean value)
#fake_rabbit=false

[profiler]

#
# Options defined in cinder.service
#

# If False fully disable profiling feature. (boolean value)
#profiler_enabled=false

# If False doesn't trace SQL requests. (boolean value)
#trace_sqlalchemy=false

[DEFAULT]

[keystone_auth_token]

#
# From keystonemiddleware.auth_token
#

# Complete public Identity API endpoint. (string value)
#auth_uri = <None>

# API version of the admin Identity API endpoint. (string value)
#auth_version = <None>

# Do not handle authorization requests within the middleware, but
# delegate the authorization decision to downstream WSGI components.
# (boolean value)
#delay_auth_decision = false

# Request timeout value for communicating with Identity API server.
# (integer value)

```

```
#http_connect_timeout = <None>

# How many times are we trying to reconnect when communicating with
# Identity API Server. (integer value)
#http_request_max_retries = 3

# Env key for the swift cache. (string value)
#cache = <None>

# Required if identity server requires client certificate (string
# value)
#certfile = <None>

# Required if identity server requires client certificate (string
# value)
#keyfile = <None>

# A PEM encoded Certificate Authority to use when verifying HTTPS
# connections. Defaults to system CAs. (string value)
#cafile = <None>

# Verify HTTPS connections. (boolean value)
#insecure = false

# Directory used to cache files related to PKI tokens. (string value)
#signing_dir = <None>

# Optionally specify a list of memcached server(s) to use for caching.
# If left undefined, tokens will instead be cached in-process. (list
# value)
# Deprecated group/name - [DEFAULT]/memcache_servers
#memcached_servers = <None>

# 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. (integer
# value)
#token_cache_time = 300

# 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. (integer value)
#revocation_cache_time = 10

# (Optional) If defined, indicate whether token data should be
# authenticated or authenticated and encrypted. Acceptable values are
# MAC or ENCRYPT. 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. (string value)
#memcache_security_strategy = <None>

# (Optional, mandatory if memcache_security_strategy is defined) This
# string is used for key derivation. (string value)
#memcache_secret_key = <None>

# (Optional) Number of seconds memcached server is considered dead
# before it is tried again. (integer value)
#memcache_pool_dead_retry = 300
```

```
# (Optional) Maximum total number of open connections to every
# memcached server. (integer value)
#memcache_pool_maxsize = 10

# (Optional) Socket timeout in seconds for communicating with a
# memcache server. (integer value)
#memcache_pool_socket_timeout = 3

# (Optional) Number of seconds a connection to memcached is held
# unused in the pool before it is closed. (integer value)
#memcache_pool_unused_timeout = 60

# (Optional) Number of seconds that an operation will wait to get a
# memcache client connection from the pool. (integer value)
#memcache_pool_conn_get_timeout = 10

# (Optional) Use the advanced (eventlet safe) memcache client pool.
# The advanced pool will only work under python 2.x. (boolean value)
#memcache_use_advanced_pool = false

# (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. (boolean
# value)
#include_service_catalog = true

# 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. (string value)
#enforce_token_bind = permissive

# If true, the revocation list will be checked for cached tokens. This
# requires that PKI tokens are configured on the identity server.
# (boolean value)
#check_revocations_for_cached = false

# 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. (list
# value)
#hash_algorithms = md5

# Prefix to prepend at the beginning of the path. Deprecated, use
# identity_uri. (string value)
#auth_admin_prefix =

# Host providing the admin Identity API endpoint. Deprecated, use
# identity_uri. (string value)
#auth_host = 127.0.0.1
```

```
# Port of the admin Identity API endpoint. Deprecated, use
# identity_uri. (integer value)
#auth_port = 35357

# Protocol of the admin Identity API endpoint (http or https).
# Deprecated, use identity_uri. (string value)
#auth_protocol = https

# Complete admin Identity API endpoint. This should specify the
# unversioned root endpoint e.g. https://localhost:35357/ (string
# value)
#identity_uri = <None>

# This option is deprecated and may be removed in a future release.
# Single shared secret with the Keystone configuration used for
# bootstrapping a Keystone installation, or otherwise bypassing the
# normal authentication process. This option should not be used, use
# `admin_user` and `admin_password` instead. (string value)
#admin_token = <None>

# Service username. (string value)
#admin_user = <None>

# Service user password. (string value)
#admin_password = <None>

# Service tenant name. (string value)
#admin_tenant_name = admin
```

## api-paste.ini

Use the `api-paste.ini` file to configure the Block Storage API service.

```
#####
# OpenStack #
#####

[composite:osapi_volume]
use = call:cinder.api:root_app_factory
/: apiversions
/v1: openstack_volume_api_v1
/v2: openstack_volume_api_v2

[composite:openstack_volume_api_v1]
use = call:cinder.api.middleware.auth:pipeline_factory
noauth = request_id faultwrap sizelimit osprofiler noauth apiv1
keystone = request_id faultwrap sizelimit osprofiler authtoken keystonecontext
  apiv1
keystone_nolimit = request_id faultwrap sizelimit osprofiler authtoken
  keystonecontext apiv1

[composite:openstack_volume_api_v2]
use = call:cinder.api.middleware.auth:pipeline_factory
noauth = request_id faultwrap sizelimit osprofiler noauth apiv2
keystone = request_id faultwrap sizelimit osprofiler authtoken keystonecontext
  apiv2
```

```
keystone_nolimit = request_id faultwrap sizelimit osprofiler authtoken
keystonecontext apiv2

[filter:request_id]
paste.filter_factory = oslo_middleware.request_id:RequestId.factory

[filter:faultwrap]
paste.filter_factory = cinder.api.middleware.fault:FaultWrapper.factory

[filter:osprofiler]
paste.filter_factory = osprofiler.web:WsgiMiddleware.factory
hmac_keys = SECRET_KEY
enabled = yes

[filter:noauth]
paste.filter_factory = cinder.api.middleware.auth:NoAuthMiddleware.factory

[filter:sizelimit]
paste.filter_factory = cinder.api.middleware.sizelimit:RequestBodySizeLimiter.factory

[app:apiv1]
paste.app_factory = cinder.api.v1.router:APIRouter.factory

[app:apiv2]
paste.app_factory = cinder.api.v2.router:APIRouter.factory

[pipeline:apiversions]
pipeline = faultwrap osvolumeverisonapp

[app:osvolumeverisonapp]
paste.app_factory = cinder.api.versions:Versions.factory

#####
# Shared #
#####

[filter:keystonecontext]
paste.filter_factory = cinder.api.middleware.auth:CinderKeystoneContext.factory

[filter:authtoken]
paste.filter_factory = keystone.middleware.auth_token:filter_factory
```

## policy.json

The `policy.json` file defines additional access controls that apply to the Block Storage service.

```
{
  "context_is_admin": "role:admin",
  "admin_or_owner": "is_admin:True or project_id:%(project_id)s",
  "default": "rule:admin_or_owner",

  "admin_api": "is_admin:True",

  "volume:create": "",
```

```
"volume:delete": "",
"volume:get": "",
"volume:get_all": "",
"volume:get_volume_metadata": "",
"volume:get_volume_admin_metadata": "rule:admin_api",
"volume:delete_volume_admin_metadata": "rule:admin_api",
"volume:update_volume_admin_metadata": "rule:admin_api",
"volume:get_snapshot": "",
"volume:get_all_snapshots": "",
"volume:extend": "",
"volume:update_readonly_flag": "",
"volume:retype": "",

"volume_extension:types_manage": "rule:admin_api",
"volume_extension:types_extra_specs": "rule:admin_api",
"volume_extension:volume_type_access": "",
"volume_extension:volume_type_access:addProjectAccess": "rule:admin_api",
"volume_extension:volume_type_access:removeProjectAccess":
"rule:admin_api",
"volume_extension:volume_type_encryption": "rule:admin_api",
"volume_extension:volume_encryption_metadata": "rule:admin_or_owner",
"volume_extension:extended_snapshot_attributes": "",
"volume_extension:volume_image_metadata": "",

"volume_extension:quotas:show": "",
"volume_extension:quotas:update": "rule:admin_api",
"volume_extension:quota_classes": "",

"volume_extension:volume_admin_actions:reset_status": "rule:admin_api",
"volume_extension:snapshot_admin_actions:reset_status": "rule:admin_api",
"volume_extension:backup_admin_actions:reset_status": "rule:admin_api",
"volume_extension:volume_admin_actions:force_delete": "rule:admin_api",
"volume_extension:volume_admin_actions:force_detach": "rule:admin_api",
"volume_extension:snapshot_admin_actions:force_delete": "rule:admin_api",
"volume_extension:volume_admin_actions:migrate_volume": "rule:admin_api",
"volume_extension:volume_admin_actions:migrate_volume_completion":
"rule:admin_api",

"volume_extension:volume_host_attribute": "rule:admin_api",
"volume_extension:volume_tenant_attribute": "rule:admin_or_owner",
"volume_extension:volume_mig_status_attribute": "rule:admin_api",
"volume_extension:hosts": "rule:admin_api",
"volume_extension:services": "rule:admin_api",

"volume_extension:volume_manage": "rule:admin_api",
"volume_extension:volume_unmanage": "rule:admin_api",

"volume:services": "rule:admin_api",

"volume:create_transfer": "",
"volume:accept_transfer": "",
"volume:delete_transfer": "",
"volume:get_all_transfers": "",

"volume_extension:replication:promote": "rule:admin_api",
"volume_extension:replication:reenable": "rule:admin_api",

"backup:create": "",
"backup:delete": "",
"backup:get": "",
```

```

"backup:get_all": "",
"backup:restore": "",
"backup:backup-import": "rule:admin_api",
"backup:backup-export": "rule:admin_api",

"snapshot_extension:snapshot_actions:update_snapshot_status": "",

"consistencygroup:create" : "group:nobody",
"consistencygroup:delete": "group:nobody",
"consistencygroup:update": "group:nobody",
"consistencygroup:get": "group:nobody",
"consistencygroup:get_all": "group:nobody",

"consistencygroup:create_cgsnapshot" : "group:nobody",
"consistencygroup:delete_cgsnapshot": "group:nobody",
"consistencygroup:get_cgsnapshot": "group:nobody",
"consistencygroup:get_all_cgsnapshots": "group:nobody",

"scheduler_extension:scheduler_stats:get_pools" : "rule:admin_api"
}

```

## rootwrap.conf

The `rootwrap.conf` file defines configuration values used by the `rootwrap` script when the Block Storage service must escalate its privileges to those of the root user.

```

# Configuration for cinder-rootwrap
# This file should be owned by (and only-writeable by) the root user

[DEFAULT]
# List of directories to load filter definitions from (separated by ',').
# These directories MUST all be only writeable by root !
filters_path=/etc/cinder/rootwrap.d,/usr/share/cinder/rootwrap

# List of directories to search executables in, in case filters do not
# explicitly specify a full path (separated by ',')
# If not specified, defaults to system PATH environment variable.
# These directories MUST all be only writeable by root !
exec_dirs=/sbin,/usr/sbin,/bin,/usr/bin,/usr/local/bin,/usr/local/sbin

# Enable logging to syslog
# Default value is False
use_syslog=False

# Which syslog facility to use.
# Valid values include auth, authpriv, syslog, local0, local1...
# Default value is 'syslog'
syslog_log_facility=syslog

# Which messages to log.
# INFO means log all usage
# ERROR means only log unsuccessful attempts
syslog_log_level=ERROR

```

## Log files used by Block Storage

The corresponding log file of each Block Storage service is stored in the `/var/log/cinder/` directory of the host on which each service runs.







































**Table 2.60. Description of IBM FlashSystem volume driver configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>flashsystem_connection_protocol = FC</code>	(StrOpt) Connection protocol should be FC.
<code>flashsystem_multihostmap_enabled = True</code>	(BoolOpt) Allows vdisk to multi host mapping.
<code>flashsystem_multipath_enabled = False</code>	(BoolOpt) Connect with multipath (FC only).

**Table 2.61. Description of HP 3PAR Fibre Channel and iSCSI drivers configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>hp3par_api_url =</code>	(StrOpt) 3PAR WSAPI Server Url like <code>https://&lt;3par ip&gt;:8080/api/v1</code>
<code>hp3par_cpg = OpenStack</code>	(ListOpt) List of the CPG(s) to use for volume creation
<code>hp3par_cpg_snap =</code>	(StrOpt) The CPG to use for Snapshots for volumes. If empty the userCPG will be used.
<code>hp3par_debug = False</code>	(BoolOpt) Enable HTTP debugging to 3PAR
<code>hp3par_iscsi_chap_enabled = False</code>	(BoolOpt) Enable CHAP authentication for iSCSI connections.
<code>hp3par_iscsi_ips =</code>	(ListOpt) List of target iSCSI addresses to use.
<code>hp3par_password =</code>	(StrOpt) 3PAR Super user password
<code>hp3par_snapshot_expiration =</code>	(StrOpt) The time in hours when a snapshot expires and is deleted. This must be larger than expiration
<code>hp3par_snapshot_retention =</code>	(StrOpt) The time in hours to retain a snapshot. You can't delete it before this expires.
<code>hp3par_username =</code>	(StrOpt) 3PAR Super user username

**Table 2.62. Description of HP LeftHand/StoreVirtual driver configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>hplefthand_api_url = None</code>	(StrOpt) HP LeftHand WSAPI Server Url like <code>https://&lt;Left-Hand ip&gt;:8081/lhos</code>
<code>hplefthand_clustlename = None</code>	(StrOpt) HP LeftHand cluster name
<code>hplefthand_debug = False</code>	(BoolOpt) Enable HTTP debugging to LeftHand
<code>hplefthand_iscsi_chap_enabled = False</code>	(BoolOpt) Configure CHAP authentication for iSCSI connections (Default: Disabled)
<code>hplefthand_password = None</code>	(StrOpt) HP LeftHand Super user password
<code>hplefthand_username = None</code>	(StrOpt) HP LeftHand Super user username

**Table 2.63. Description of Huawei storage driver configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>cinder_huawei_conf_file = /etc/cinder/cinder_huawei_conf.xml</code>	(StrOpt) The configuration file for the Cinder Huawei driver

**Table 2.64. Description of IBM NAS volume driver configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>ibmnas_platform_type = v7ku</code>	(StrOpt) IBMNAS platform type to be used as backend storage; valid values are - <code>v7ku</code> : for using IBM Storwize V7000 Unified, <code>sonas</code> : for using IBM Scale Out NAS, <code>gpfnas</code> : for using NFS based IBM GPFS deployments.

**Table 2.65. Description of images configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>allowed_direct_url_schemes =</code>	(ListOpt) A list of url schemes that can be downloaded directly via the <code>direct_url</code> . Currently supported schemes: <code>[file]</code> .
<code>glance_api_insecure = False</code>	(BoolOpt) Allow to perform insecure SSL (https) requests to glance
<code>glance_api_servers = \$glance_host:\$glance_port</code>	(ListOpt) A list of the glance API servers available to cinder ( <code>[hostname ip]:port</code> )
<code>glance_api_ssl_compression = False</code>	(BoolOpt) 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 <code>qcow2</code> .
<code>glance_api_version = 1</code>	(IntOpt) Version of the glance API to use
<code>glance_ca_certificates_file = None</code>	(StrOpt) Location of ca certificates file to use for glance client requests.
<code>glance_core_properties = checksum, container_format, disk_format, image_name, image_id, min_disk, min_ram, name, size</code>	(ListOpt) Default core properties of image
<code>glance_host = \$my_ip</code>	(StrOpt) Default glance host name or IP
<code>glance_num_retries = 0</code>	(IntOpt) Number retries when downloading an image from glance
<code>glance_port = 9292</code>	(IntOpt) Default glance port
<code>glance_request_timeout = None</code>	(IntOpt) http/https timeout value for glance operations. If no value ( <code>None</code> ) is supplied here, the <code>glanceclient</code> default value is used.
<code>image_conversion_dir = \$state_path/conversion</code>	(StrOpt) Directory used for temporary storage during image conversion
<code>use_multipath_for_image_xfer = False</code>	(BoolOpt) Do we attach/detach volumes in cinder using multipath for volume to image and image to volume transfers?

**Table 2.66. Description of key manager configuration options**

Configuration option = Default value	Description
[keymgr]	
<code>api_class = cinder.keymgr.conf_key_mgr.ConfKeyManager</code>	(StrOpt) The full class name of the key manager API class
<code>encryption_api_url = http://localhost:9311/v1</code>	(StrOpt) Url for encryption service.
<code>encryption_auth_url = http://localhost:5000/v3</code>	(StrOpt) Authentication url for encryption service.
<code>fixed_key = None</code>	(StrOpt) Fixed key returned by key manager, specified in hex

**Table 2.67. Description of logging configuration options**

Configuration option = Default value	Description
[DEFAULT]	
debug = <i>False</i>	(BoolOpt) Print debugging output (set logging level to DEBUG instead of default WARNING level).
default_log_levels = <i>amqp=WARN, amqpplib=WARN, boto=WARN, qpuid=WARN, sqlalchemy=WARN, suds=INFO, oslo.messaging=INFO, iso8601=WARN, requests.packages.urllib3.connectionpool=WARN, urllib3.connectionpool=WARN, websocket=WARN, requests.packages.urllib3.util.retry=WARN, urllib3.util.retry=WARN, keystonemiddleware=WARN, routes.middleware=WARN, stevedore=WARN</i>	(ListOpt) List of logger=LEVEL pairs.
fatal_deprecations = <i>False</i>	(BoolOpt) Enables or disables fatal status of deprecations.
fatal_exception_format_errors = <i>False</i>	(BoolOpt) Make exception message format errors fatal.
instance_format = "[instance: %(uuid)s] "	(StrOpt) The format for an instance that is passed with the log message.
instance_uuid_format = "[instance: %(uuid)s] "	(StrOpt) The format for an instance UUID that is passed with the log message.
log_config_append = <i>None</i>	(StrOpt) 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.
log_date_format = %Y-%m-%d %H:%M:%S	(StrOpt) Format string for %(asctime)s in log records. Default: %(default)s .
log_dir = <i>None</i>	(StrOpt) (Optional) The base directory used for relative – log-file paths.
log_file = <i>None</i>	(StrOpt) (Optional) Name of log file to output to. If no default is set, logging will go to stdout.
log_format = <i>None</i>	(StrOpt) DEPRECATED. A logging.Formatter log message format string which may use any of the available logging.LogRecord attributes. This option is deprecated. Please use logging_context_format_string and logging_default_format_string instead.
log_config_append = <i>None</i>	(StrOpt) 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.
log_date_format = %Y-%m-%d %H:%M:%S	(StrOpt) Format string for %(asctime)s in log records. Default: %(default)s .
log_dir = <i>None</i>	(StrOpt) (Optional) The base directory used for relative – log-file paths.
log_file = <i>None</i>	(StrOpt) (Optional) Name of log file to output to. If no default is set, logging will go to stdout.
log_format = <i>None</i>	(StrOpt) DEPRECATED. A logging.Formatter log message format string which may use any of the available logging.LogRecord attributes. This option is deprecated. Please use logging_context_format_string and logging_default_format_string instead.
logging_context_format_string = %(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [%(request_id)s %(user_identity)s] %(instance)s%(message)s	(StrOpt) Format string to use for log messages with context.

Configuration option = Default value	Description
<code>logging_debug_format_suffix = %(funcName)s %(pathname)s:%(lineno)d</code>	(StrOpt) Data to append to log format when level is DEBUG.
<code>logging_default_format_string = %(asctime)s. %(msecs)03d %(process)d %(levelname)s %(name)s [-] %(instance)s%(message)s</code>	(StrOpt) Format string to use for log messages without context.
<code>logging_exception_prefix = %(asctime)s. %(msecs)03d %(process)d TRACE %(name)s %(instance)s</code>	(StrOpt) Prefix each line of exception output with this format.
<code>publish_errors = False</code>	(BoolOpt) Enables or disables publication of error events.
<code>syslog_log_facility = LOG_USER</code>	(StrOpt) Syslog facility to receive log lines.
<code>syslog_log_facility = LOG_USER</code>	(StrOpt) Syslog facility to receive log lines.
<code>use_syslog = False</code>	(BoolOpt) Use syslog for logging. Existing syslog format is DEPRECATED during I, and will change in J to honor RFC5424.
<code>use_syslog_rfc_format = False</code>	(BoolOpt) (Optional) Enables or disables syslog rfc5424 format for logging. If enabled, prefixes the MSG part of the syslog message with APP-NAME (RFC5424). The format without the APP-NAME is deprecated in I, and will be removed in J.
<code>use_stderr = True</code>	(BoolOpt) Log output to standard error.
<code>use_syslog = False</code>	(BoolOpt) Use syslog for logging. Existing syslog format is DEPRECATED during I, and will change in J to honor RFC5424.
<code>use_syslog_rfc_format = False</code>	(BoolOpt) (Optional) Enables or disables syslog rfc5424 format for logging. If enabled, prefixes the MSG part of the syslog message with APP-NAME (RFC5424). The format without the APP-NAME is deprecated in I, and will be removed in J.
<code>verbose = False</code>	(BoolOpt) Print more verbose output (set logging level to INFO instead of default WARNING level).

**Table 2.68. Description of NAS configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>nas_ip =</code>	(StrOpt) IP address or Hostname of NAS system.
<code>nas_login = admin</code>	(StrOpt) User name to connect to NAS system.
<code>nas_mount_options = None</code>	(StrOpt) Options used to mount the storage backend file system where Cinder volumes are stored.
<code>nas_password =</code>	(StrOpt) Password to connect to NAS system.
<code>nas_private_key =</code>	(StrOpt) Filename of private key to use for SSH authentication.
<code>nas_secure_file_operations = auto</code>	(StrOpt) 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.
<code>nas_secure_file_permissions = auto</code>	(StrOpt) 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	Description
<code>broadcast_prefix = broadcast</code>	(StrOpt) address prefix used when broadcasting to all servers
<code>container_name = None</code>	(StrOpt) Name for the AMQP container
<code>group_request_prefix = unicast</code>	(StrOpt) address prefix when sending to any server in group
<code>idle_timeout = 0</code>	(IntOpt) Timeout for inactive connections (in seconds)
<code>server_request_prefix = exclusive</code>	(StrOpt) address prefix used when sending to a specific server
<code>ssl_ca_file =</code>	(StrOpt) CA certificate PEM file to verify server certificate
<code>ssl_cert_file =</code>	(StrOpt) Identifying certificate PEM file to present to clients
<code>ssl_key_file =</code>	(StrOpt) Private key PEM file used to sign cert_file certificate
<code>ssl_key_password = None</code>	(StrOpt) Password for decrypting ssl_key_file (if encrypted)
<code>trace = False</code>	(BoolOpt) Debug: dump AMQP frames to stdout

**Table 2.78. Description of SAN configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>san_clustername =</code>	(StrOpt) Cluster name to use for creating volumes
<code>san_ip =</code>	(StrOpt) IP address of SAN controller
<code>san_is_local = False</code>	(BoolOpt) Execute commands locally instead of over SSH; use if the volume service is running on the SAN device
<code>san_login = admin</code>	(StrOpt) Username for SAN controller
<code>san_password =</code>	(StrOpt) Password for SAN controller
<code>san_private_key =</code>	(StrOpt) Filename of private key to use for SSH authentication
<code>san_secondary_ip = None</code>	(StrOpt) VNX secondary SP IP Address.
<code>san_ssh_port = 22</code>	(IntOpt) SSH port to use with SAN
<code>san_thin_provision = True</code>	(BoolOpt) Use thin provisioning for SAN volumes?
<code>ssh_conn_timeout = 30</code>	(IntOpt) SSH connection timeout in seconds
<code>ssh_max_pool_conn = 5</code>	(IntOpt) Maximum ssh connections in the pool
<code>ssh_min_pool_conn = 1</code>	(IntOpt) Minimum ssh connections in the pool

**Table 2.79. Description of scheduler configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>filter_function = None</code>	(StrOpt) 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.
<code>goodness_function = None</code>	(StrOpt) 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.
<code>scheduler_default_filters = AvailabilityZoneFilter, CapacityFilter, CapabilitiesFilter</code>	(ListOpt) Which filter class names to use for filtering hosts when not specified in the request.



Configuration option = Default value	Description
<code>scheduler_default_weighers = CapacityWeigher</code>	(ListOpt) Which weigher class names to use for weighing hosts.
<code>scheduler_driver = cinder.scheduler.filter_scheduler.FilterScheduler</code>	(StrOpt) Default scheduler driver to use
<code>scheduler_host_manager = cinder.scheduler.host_manager.HostManager</code>	(StrOpt) The scheduler host manager class to use
<code>scheduler_json_config_location =</code>	(StrOpt) Absolute path to scheduler configuration JSON file.
<code>scheduler_manager = cinder.scheduler.manager.SchedulerManager</code>	(StrOpt) Full class name for the Manager for scheduler
<code>scheduler_max_attempts = 3</code>	(IntOpt) Maximum number of attempts to schedule an volume
<code>scheduler_topic = cinder-scheduler</code>	(StrOpt) The topic that scheduler nodes listen on

**Table 2.80. Description of SCST volume driver configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>scst_target_driver = iscsi</code>	(StrOpt) SCST target implementation can choose from multiple SCST target drivers.
<code>scst_target_ign_name = None</code>	(StrOpt) Certain iSCSI targets have predefined target names, SCST target driver uses this name.

**Table 2.81. Description of Scality REST Block storage driver configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>srb_base_urls = None</code>	(StrOpt) Comma-separated list of REST servers IP to connect to. (eg <code>http://IP1/,http://IP2:81/path</code>

**Table 2.82. Description of storage configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>allocated_capacity_weight_multiplier = -1.0</code>	(FloatOpt) Multiplier used for weighing volume capacity. Negative numbers mean to stack vs spread.
<code>capacity_weight_multiplier = 1.0</code>	(FloatOpt) Multiplier used for weighing volume capacity. Negative numbers mean to stack vs spread.
<code>enabled_backends = None</code>	(ListOpt) A list of backend names to use. These backend names should be backed by a unique [CONFIG] group with its options
<code>iscsi_helper = tgtadm</code>	(StrOpt) 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>iseradm</code> for the ISER protocol, <code>ietadm</code> for iSCSI Enterprise Target, <code>iscsictl</code> for Chelsio iSCSI Target or fake for testing.
<code>iscsi_iotype = fileio</code>	(StrOpt) 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
<code>iscsi_ip_address = \$my_ip</code>	(StrOpt) The IP address that the iSCSI daemon is listening on
<code>iscsi_num_targets = None</code>	(IntOpt) This option is deprecated and unused. It will be removed in the Liberty release.

















Option = default value	(Type) Help string
[oslo_messaging_rabbit] kombu_ssl_ca_certs =	(StrOpt) SSL certification authority file (valid only if SSL enabled).
[oslo_messaging_rabbit] kombu_ssl_certfile =	(StrOpt) SSL cert file (valid only if SSL enabled).
[oslo_messaging_rabbit] kombu_ssl_keyfile =	(StrOpt) SSL key file (valid only if SSL enabled).
[oslo_messaging_rabbit] kombu_ssl_version =	(StrOpt) 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.
[oslo_messaging_rabbit] rabbit_ha_queues = False	(BoolOpt) Use HA queues in RabbitMQ (x-ha-policy: all). If you change this option, you must wipe the RabbitMQ database.
[oslo_messaging_rabbit] rabbit_host = localhost	(StrOpt) The RabbitMQ broker address where a single node is used.
[oslo_messaging_rabbit] rabbit_hosts = \$rabbit_host:\$rabbit_port	(ListOpt) RabbitMQ HA cluster host:port pairs.
[oslo_messaging_rabbit] rabbit_login_method = AMQ-PLAIN	(StrOpt) The RabbitMQ login method.
[oslo_messaging_rabbit] rabbit_max_retries = 0	(IntOpt) Maximum number of RabbitMQ connection retries. Default is 0 (infinite retry count).
[oslo_messaging_rabbit] rabbit_password = guest	(StrOpt) The RabbitMQ password.
[oslo_messaging_rabbit] rabbit_port = 5672	(IntOpt) The RabbitMQ broker port where a single node is used.
[oslo_messaging_rabbit] rabbit_retry_backoff = 2	(IntOpt) How long to backoff for between retries when connecting to RabbitMQ.
[oslo_messaging_rabbit] rabbit_retry_interval = 1	(IntOpt) How frequently to retry connecting with RabbitMQ.
[oslo_messaging_rabbit] rabbit_use_ssl = False	(BoolOpt) Connect over SSL for RabbitMQ.
[oslo_messaging_rabbit] rabbit_userid = guest	(StrOpt) The RabbitMQ userid.
[oslo_messaging_rabbit] rabbit_virtual_host = /	(StrOpt) The RabbitMQ virtual host.
[oslo_messaging_rabbit] rpc_conn_pool_size = 30	(IntOpt) Size of RPC connection pool.
[oslo_middleware] max_request_body_size = 114688	(IntOpt) The maximum body size for each request, in bytes.

**Table 2.87. New default values**

Option	Previous default value	New default value
[DEFAULT] backup_metadata_version	1	2
[DEFAULT] client_socket_timeout	0	900
[DEFAULT] default_log_levels	amqp=WARN, amqplib=WARN, boto=WARN, qpid=WARN, sqlalchemy=WARN, suds=INFO, oslo.messaging=INFO, iso8601=WARN, requests.packages.urllib3.connectionpool=WARN, urllib3.connectionpool=WARN, websocket=WARN, keystone.middleware=WARN, routes.middleware=WARN, stevedore=WARN	amqp=WARN, amqplib=WARN, boto=WARN, qpid=WARN, sqlalchemy=WARN, suds=INFO, oslo.messaging=INFO, iso8601=WARN, requests.packages.urllib3.connectionpool=WARN, urllib3.connectionpool=WARN, websocket=WARN, requests.packages.urllib3.util.retry=WARN, urllib3.util.retry=WARN, keystone.middleware=WARN, routes.middleware=WARN, stevedore=WARN
[DEFAULT] iscsi_num_targets	100	None
[DEFAULT] iser_num_targets	100	None
[DEFAULT] iser_target_prefix	iqn.2010-10.org.iser.openstack:	iqn.2010-10.org.openstack:

Option	Previous default value	New default value
[DEFAULT] nova_catalog_admin_info	compute:nova:adminURL	compute:Compute Service:adminURL
[DEFAULT] nova_catalog_info	compute:nova:publicURL	compute:Compute Service:publicURL
[DEFAULT] rpc_zmq_matchmaker	oslo.messaging_drivers.matchmaker.MatchmakerLocalhost	MatchmakerLocalhost
[keymgr] encryption_auth_url	http://localhost:5000/v2.0	http://localhost:5000/v3

**Table 2.88. Deprecated options**

Deprecated option	New Option
[DEFAULT] log-format	None
[DEFAULT] use-syslog	None
[DEFAULT] use_syslog	None
[DEFAULT] osapi_max_request_body_size	[oslo_middleware] max_request_body_size
[DEFAULT] eqlx_chap_password	[DEFAULT] chap_password
[DEFAULT] datera_api_token	None
[DEFAULT] eqlx_use_chap	[DEFAULT] use_chap_auth
[DEFAULT] enable_v1_api	None
[DEFAULT] db_backend	[database] backend
[DEFAULT] host	[DEFAULT] backend_host
[DEFAULT] eqlx_chap_login	[DEFAULT] chap_username
[DEFAULT] log_format	None

## 3. Compute

### Table of Contents

Overview of nova.conf .....	211
Configure logging .....	213
Configure authentication and authorization .....	213
Configure resize .....	213
Database configuration .....	214
Configure the Oslo RPC messaging system .....	214
Configure the Compute API .....	218
Configure the EC2 API .....	220
Fibre Channel support in Compute .....	221
iSCSI interface and offload support in Compute .....	221
Hypervisors .....	223
Scheduling .....	256
Cells .....	274
Conductor .....	279
Example nova.conf configuration files .....	279
Compute log files .....	284
Compute sample configuration files .....	284
New, updated and deprecated options in Kilo for OpenStack Compute .....	328

The OpenStack Compute service is a cloud computing fabric controller, which is the main part of an IaaS system. You can use OpenStack Compute to host and manage cloud computing systems. This section describes the OpenStack Compute configuration options.

To configure your Compute installation, you must define configuration options in these files:

- `nova.conf`. Contains most of the Compute configuration options. Resides in the `/etc/nova` directory.
- `api-paste.ini`. Defines Compute limits. Resides in the `/etc/nova` directory.
- Related Image service and Identity service management configuration files.

### Overview of nova.conf

The `nova.conf` configuration file is an [INI file format](#) as explained in [the section called "Configuration file format" \[xxi\]](#).

You can use a particular configuration option file by using the `option (nova.conf)` parameter when you run one of the `nova-*` services. This parameter inserts configuration option definitions from the specified configuration file name, which might be useful for debugging or performance tuning.

For a list of configuration options, see the tables in this guide.



[spice]	Configures virtual consoles using SPICE.
[ssl]	Configures certificate authority using SSL.
[trusted_computing]	Configures the trusted computing pools functionality and how to connect to a remote attestation service.
[upgrade_levels]	Configures version locking on the RPC (message queue) communications between the various Compute services to allow live upgrading an OpenStack installation.
[vmware]	Configures the VMWare hypervisor driver.
[xenserver]	Configures the XenServer hypervisor driver.
[zookeeper]	Configures the ZooKeeper ServiceGroup driver.

## Configure logging

You can use `nova.conf` file to configure where Compute logs events, the level of logging, and log formats.

To customize log formats for OpenStack Compute, use the configuration option settings documented in [Table 3.39, "Description of logging configuration options" \[300\]](#).

## Configure authentication and authorization

There are different methods of authentication for the OpenStack Compute project, including no authentication. The preferred system is the OpenStack Identity service, code-named Keystone.

To customize authorization settings for Compute, use the configuration options documented in [Table 3.14, "Description of authentication configuration options" \[286\]](#).

To customize certificate authority settings for Compute, use the configuration options documented in [Table 3.18, "Description of CA and SSL configuration options" \[288\]](#).

To customize Compute and the Identity service to use LDAP as a backend, refer to the configuration options documented in [Table 3.36, "Description of LDAP configuration options" \[298\]](#).

## Configure resize

Resize (or Server resize) is the ability to change the flavor of a server, thus allowing it to up-scale or downscale according to user needs. For this feature to work properly, you might need to configure some underlying virt layers.

### KVM

Resize on KVM is implemented currently by transferring the images between compute nodes over ssh. For KVM you need hostnames to resolve properly and passwordless ssh ac-

cess between your compute hosts. Direct access from one compute host to another is needed to copy the VM file across.

Cloud end users can find out how to resize a server by reading the [OpenStack End User Guide](#).

## XenServer

To get resize to work with XenServer (and XCP), you need to establish a root trust between all hypervisor nodes and provide an `/image` mount point to your hypervisors `dom0`.

## Database configuration

You can configure OpenStack Compute to use any SQLAlchemy-compatible database. The database name is `nova`. The `nova-conductor` service is the only service that writes to the database. The other Compute services access the database through the `nova-conductor` service.

To ensure that the database schema is current, run the following command:

```
# nova-manage db sync
```

If `nova-conductor` is not used, entries to the database are mostly written by the `nova-scheduler` service, although all services must be able to update entries in the database.

In either case, use the configuration option settings documented in [Table 3.25, “Description of database configuration options” \[293\]](#) to configure the connection string for the `nova` database.

## Configure the Oslo RPC messaging system

OpenStack projects use AMQP, an open standard for messaging middleware. OpenStack services that run on multiple servers to talk to each other. OpenStack Oslo RPC supports three implementations of AMQP: RabbitMQ, Qpid, and ZeroMQ.

## Configure RabbitMQ

OpenStack Oslo RPC uses RabbitMQ by default. Use these options to configure the RabbitMQ message system. The `rpc_backend` option is not required as long as RabbitMQ is the default messaging system. However, if it is included the configuration, you must set it to `rabbit`.

```
rpc_backend=rabbit
```

You can use these additional options to configure the RabbitMQ messaging system. You can configure messaging communication for different installation scenarios, tune retries for RabbitMQ, and define the size of the RPC thread pool. To monitor notifications through RabbitMQ, you must set the `notification_driver` option to `nova.openstack.common.notifier.rpc_notifier` in the `nova.conf` file. The de-

fault for sending usage data is sixty seconds plus a random number of seconds from zero to sixty.

**Table 3.1. Description of RabbitMQ configuration options**

Configuration option = Default value	Description
[oslo_messaging_rabbit]	
<code>amqp_auto_delete = False</code>	(BoolOpt) Auto-delete queues in AMQP.
<code>amqp_durable_queues = False</code>	(BoolOpt) Use durable queues in AMQP.
<code>fake_rabbit = False</code>	(BoolOpt) Deprecated, use <code>rpc_backend=kombu+memory</code> or <code>rpc_backend=fake</code>
<code>heartbeat_rate = 2</code>	(IntOpt) How often times during the <code>heartbeat_timeout_threshold</code> we check the heartbeat.
<code>heartbeat_timeout_threshold = 0</code>	(IntOpt) Number of seconds after which the Rabbit broker is considered down if heartbeat's keep-alive fails (0 disables the heartbeat, >0 enables it. Enabling heartbeats requires <code>kombu&gt;=3.0.7</code> and <code>amqp&gt;=1.4.0</code> ). EXPERIMENTAL
<code>kombu_reconnect_delay = 1.0</code>	(FloatOpt) How long to wait before reconnecting in response to an AMQP consumer cancel notification.
<code>kombu_ssl_ca_certs =</code>	(StrOpt) SSL certification authority file (valid only if SSL enabled).
<code>kombu_ssl_certfile =</code>	(StrOpt) SSL cert file (valid only if SSL enabled).
<code>kombu_ssl_keyfile =</code>	(StrOpt) SSL key file (valid only if SSL enabled).
<code>kombu_ssl_version =</code>	(StrOpt) 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.
<code>rabbit_ha_queues = False</code>	(BoolOpt) Use HA queues in RabbitMQ (x-ha-policy: all). If you change this option, you must wipe the RabbitMQ database.
<code>rabbit_host = localhost</code>	(StrOpt) The RabbitMQ broker address where a single node is used.
<code>rabbit_hosts = \$rabbit_host:\$rabbit_port</code>	(ListOpt) RabbitMQ HA cluster host:port pairs.
<code>rabbit_login_method = AMQPLAIN</code>	(StrOpt) The RabbitMQ login method.
<code>rabbit_max_retries = 0</code>	(IntOpt) Maximum number of RabbitMQ connection retries. Default is 0 (infinite retry count).
<code>rabbit_password = guest</code>	(StrOpt) The RabbitMQ password.
<code>rabbit_port = 5672</code>	(IntOpt) The RabbitMQ broker port where a single node is used.
<code>rabbit_retry_backoff = 2</code>	(IntOpt) How long to backoff for between retries when connecting to RabbitMQ.
<code>rabbit_retry_interval = 1</code>	(IntOpt) How frequently to retry connecting with RabbitMQ.
<code>rabbit_use_ssl = False</code>	(BoolOpt) Connect over SSL for RabbitMQ.
<code>rabbit_userid = guest</code>	(StrOpt) The RabbitMQ userid.
<code>rabbit_virtual_host = /</code>	(StrOpt) The RabbitMQ virtual host.
<code>rpc_conn_pool_size = 30</code>	(IntOpt) Size of RPC connection pool.

## Configure Qpid

Use these options to configure the Qpid messaging system for OpenStack Oslo RPC. Qpid is not the default messaging system, so you must enable it by setting the `rpc_backend` option in the `nova.conf` file.

```
rpc_backend=qpido
```

This critical option points the compute nodes to the Qpid broker (server). Set `qpido_hostname` to the host name where the broker runs in the `novo.conf` file.



**Note**

The `--qpido_hostname` parameter accepts a host name or IP address value.

```
qpido_hostname=hostname.example.com
```

If the Qpid broker listens on a port other than the AMQP default of 5672, you must set the `qpido_port` option to that value:

```
qpido_port=12345
```

If you configure the Qpid broker to require authentication, you must add a user name and password to the configuration:

```
qpido_username=username
qpido_password=password
```

By default, TCP is used as the transport. To enable SSL, set the `qpido_protocol` option:

```
qpido_protocol=ssl
```

This table lists additional options that you use to configure the Qpid messaging driver for OpenStack Oslo RPC. These options are used infrequently.

**Table 3.2. Description of Qpid configuration options**

Configuration option = Default value	Description
[oslo_messaging_qpido]	
amqp_auto_delete = <i>False</i>	(BoolOpt) Auto-delete queues in AMQP.
amqp_durable_queues = <i>False</i>	(BoolOpt) Use durable queues in AMQP.
qpido_heartbeat = <i>60</i>	(IntOpt) Seconds between connection keepalive heartbeats.
qpido_hostname = <i>localhost</i>	(StrOpt) Qpid broker hostname.
qpido_hosts = <i>\$qpido_hostname:\$qpido_port</i>	(ListOpt) Qpid HA cluster host:port pairs.
qpido_password =	(StrOpt) Password for Qpid connection.
qpido_port = <i>5672</i>	(IntOpt) Qpid broker port.
qpido_protocol = <i>tcp</i>	(StrOpt) Transport to use, either 'tcp' or 'ssl'.
qpido_receiver_capacity = <i>1</i>	(IntOpt) The number of prefetched messages held by receiver.
qpido_sasl_mechanisms =	(StrOpt) Space separated list of SASL mechanisms to use for auth.
qpido_tcp_nodelay = <i>True</i>	(BoolOpt) Whether to disable the Nagle algorithm.
qpido_topology_version = <i>1</i>	(IntOpt) The qpid topology version to use. Version 1 is what was originally used by impl_qpido. Version 2 includes some backwards-incompatible changes that allow broker federation to work. Users should update to version 2 when they are able to take everything down, as it requires a clean break.
qpido_username =	(StrOpt) Username for Qpid connection.







In practice, how the admin password is handled depends on the hypervisor in use and might require additional configuration of the instance. For example, you might have to install an agent to handle the password setting. If the hypervisor and instance configuration do not support setting a password at server create time, the password that is returned by the create API call is misleading because it was ignored.

To prevent this confusion, use the `enable_instance_password` configuration option to disable the return of the admin password for installations that do not support setting instance passwords.

## Configure Compute API rate limiting

OpenStack Compute supports API rate limiting for the OpenStack API. The rate limiting allows an administrator to configure limits on the type and number of API calls that can be made in a specific time interval.

When API rate limits are exceeded, HTTP requests return an error with a status code of 403 Forbidden.

Rate limiting is not available for the EC2 API.

## Define limits

To define limits, set these values:

- The **HTTP method** used in the API call, typically one of GET, PUT, POST, or DELETE.
- A **human readable URI** that is used as a friendly description of where the limit is applied.
- A **regular expression**. The limit is applied to all URIs that match the regular expression and HTTP method.
- A **limit value** that specifies the maximum count of units before the limit takes effect.
- An **interval** that specifies time frame to which the limit is applied. The interval can be SECOND, MINUTE, HOUR, or DAY.

Rate limits are applied in relative order to the HTTP method, going from least to most specific.

## Default limits

Normally, you install OpenStack Compute with the following limits enabled:

**Table 3.6. Default API rate limits**

HTTP method	API URI	API regular expression	Limit
POST	any URI (*)	.*	120 per minute
POST	/servers	^/servers	120 per minute
PUT	any URI (*)	.*	120 per minute
GET	*changes-since*	.*changes-since.*	120 per minute
DELETE	any URI (*)	.*	120 per minute

HTTP method	API URI	API regular expression	Limit
GET	*/os-fping	^/os-fping	12 per minute

## Configure and change limits

As part of the WSGI pipeline, the `etc/nova/api-paste.ini` file defines the actual limits.

To enable limits, include the `ratelimit` filter in the API pipeline specification. If the `ratelimit` filter is removed from the pipeline, limiting is disabled. You must also define the rate limit filter. The lines appear as follows:

```
[pipeline:openstack_compute_api_v2]
pipeline = faultwrap authToken keystonecontext ratelimit osapi_compute_app_v2

[pipeline:openstack_volume_api_v1]
pipeline = faultwrap authToken keystonecontext ratelimit osapi_volume_app_v1

[filter:ratelimit]
paste.filter_factory = nova.api.openstack.compute.
limits:RateLimitingMiddleware.factory
```

To modify the limits, add a `limits` specification to the `[filter:ratelimit]` section of the file. Specify the limits in this order:

1. HTTP method
2. friendly URI
3. regex
4. limit
5. interval

The following example shows the default rate-limiting values:

```
[filter:ratelimit]
paste.filter_factory = nova.api.openstack.compute.
limits:RateLimitingMiddleware.factory
limits = (POST, "*", .*, 120, MINUTE);(POST, "*/servers", ^/servers, 120,
MINUTE);(PUT, "*", .*, 120, MINUTE);(GET, "*changes-since", .*changes-since.
*, 120, MINUTE);(DELETE, "*", .*, 120, MINUTE);(GET, "*/os-fping", ^/os-fping,
12, MINUTE)
```

## Configuration reference

The Compute API configuration options are documented in [Table 3.12, "Description of API configuration options" \[284\]](#).

## Configure the EC2 API

You can set options in the `nova.conf` configuration file to control which network address and port the EC2 API listens on, the formatting of some API responses, and authentication related options.



Currently supported transports (`iface.transport_name`) are `be2iscsi`, `bnx2i`, `cxgb3i`, `cxgb4i`, `qla4xxx`, `ocs`. No configuration changes are needed outside of Compute node.

iSER is currently supported via the separate `iSER LibvirtISERVVolumeDriver` and will be rejected if used via the `iscsi_iface` parameter.

## iSCSI iface configuration

- Note the distinction between the transport name (`iface.transport_name`) and iface name (`iface.iscsi_ifacename`). The actual iface name must be specified via the `iscsi_iface` parameter to libvirt for offload to work.
- The default name for an iscsi iface (`open-iscsi` parameter `iface.iscsi_ifacename`) is in the format `transport_name.hwaddress` when generated by `iscsiadm`.
- `iscsiadm` can be used to view and generate current iface configuration. Every network interface that supports an open-iscsi transport can have one or more iscsi ifaces associated with it. If no ifaces have been configured for a network interface supported by an open-iscsi transport, this command will create a default iface configuration for that network interface. For example :

```
# iscsiadm -m iface
  default tcp,<empty>,<empty>,<empty>,<empty>
  iser iser,<empty>,<empty>,<empty>,<empty>
  bnx2i.00:05:b5:d2:a0:c2 bnx2i,00:05:b5:d2:a0:c2,5.10.10.20,<empty>,
<empty>
  cxgb4i.00:07:43:28:b2:58 cxgb4i,00:07:43:28:b2:58,102.50.50.80,<empty>,
<empty>
  qla4xxx.00:c0:dd:08:63:ea qla4xxx,00:c0:dd:08:63:ea,20.15.0.9,<empty>,
<empty>
```

The output is in the format : `iface_name transport_name,hwaddress,ipaddress,net_ifacename,initiatorname`.

- Individual iface configuration can be viewed via

```
# iscsiadm -m iface -I IFACE_NAME
# BEGIN RECORD 2.0-873
iface.iscsi_ifacename = cxgb4i.00:07:43:28:b2:58
iface.net_ifacename = <empty>
iface.ipaddress = 102.50.50.80
iface.hwaddress = 00:07:43:28:b2:58
iface.transport_name = cxgb4i
iface.initiatorname = <empty>
# END RECORD
```

Configuration can be updated as desired via

```
# iscsiadm -m iface -I IFACE_NAME--op=update -n iface.SETTING -v VALUE
```

- All iface configurations need a minimum of `iface.iface_name`, `iface.transport_name` and `iface.hwaddress` to be correctly configured to work. Some transports may require `iface.ipaddress` and `iface.net_ifacename` as well to bind correctly.







### For x86 based systems

- To determine whether the `svm` or `vmx` CPU extensions are present, run this command:

```
# grep -E 'svm|vmx' /proc/cpuinfo
```

This command generates output if the CPU is capable of hardware-virtualization. Even if output is shown, you might still need to enable virtualization in the system BIOS for full support.

If no output appears, consult your system documentation to ensure that your CPU and motherboard support hardware virtualization. Verify that any relevant hardware virtualization options are enabled in the system BIOS.

The BIOS for each manufacturer is different. If you must enable virtualization in the BIOS, look for an option containing the words `virtualization`, `VT`, `VMX`, or `SVM`.

- To list the loaded kernel modules and verify that the `kvm` modules are loaded, run this command:

```
# lsmod | grep kvm
```

If the output includes `kvm_intel` or `kvm_amd`, the `kvm` hardware virtualization modules are loaded and your kernel meets the module requirements for OpenStack Compute.

If the output does not show that the `kvm` module is loaded, run this command to load it:

```
# modprobe -a kvm
```

Run the command for your CPU. For Intel, run this command:

```
# modprobe -a kvm-intel
```

For AMD, run this command:

```
# modprobe -a kvm-amd
```

Because a KVM installation can change user group membership, you might need to log in again for changes to take effect.

If the kernel modules do not load automatically, use the procedures listed in these subsections.

If the checks indicate that required hardware virtualization support or kernel modules are disabled or unavailable, you must either enable this support on the system or find a system with this support.



**Note**

Some systems require that you enable VT support in the system BIOS. If you believe your processor supports hardware acceleration but the previous command did not produce output, reboot your machine, enter the system BIOS, and enable the VT option.

If KVM acceleration is not supported, configure Compute to use a different hypervisor, such as [QEMU](#) or [Xen](#).

These procedures help you load the kernel modules for Intel-based and AMD-based processors if they do not load automatically during KVM installation.

#### Intel-based processors

If your compute host is Intel-based, run these commands as root to load the kernel modules:

```
# modprobe kvm
# modprobe kvm-intel
```

Add these lines to the `/etc/modules` file so that these modules load on reboot:

```
kvm
kvm-intel
```

#### AMD-based processors

If your compute host is AMD-based, run these commands as root to load the kernel modules:

```
# modprobe kvm
# modprobe kvm-amd
```

Add these lines to `/etc/modules` file so that these modules load on reboot:

```
kvm
kvm-amd
```

#### For POWER based systems

KVM as a hypervisor is supported on POWER system's PowerNV platform.

1. To determine if your POWER platform supports KVM based virtualization run the following command:

```
#cat /proc/cpuinfo | grep PowerNV
```

If the previous command generates the following output, then CPU supports KVM based virtualization

```
platform: PowerNV
```

If no output is displayed, then your POWER platform does not support KVM based hardware virtualization.

2. To list the loaded kernel modules and verify that the `kvm` modules are loaded, run the following command:

```
# lsmod | grep kvm
```

If the output includes `kvm_hv`, the `kvm` hardware virtualization modules are loaded and your kernel meets the module requirements for OpenStack Compute.

If the output does not show that the `kvm` module is loaded, run the following command to load it:

```
# modprobe -a kvm
```

For PowerNV platform, run the following command:

```
# modprobe -a kvm-hv
```

Because a KVM installation can change user group membership, you might need to log in again for changes to take effect.

## Specify the CPU model of KVM guests

The Compute service enables you to control the guest CPU model that is exposed to KVM virtual machines. Use cases include:

- To maximize performance of virtual machines by exposing new host CPU features to the guest
- To ensure a consistent default CPU across all machines, removing reliance of variable QEMU defaults

In libvirt, the CPU is specified by providing a base CPU model name (which is a shorthand for a set of feature flags), a set of additional feature flags, and the topology (sockets/cores/threads). The libvirt KVM driver provides a number of standard CPU model names. These models are defined in the `/usr/share/libvirt/cpu_map.xml` file. Check this file to determine which models are supported by your local installation.

Two Compute configuration options in the `[libvirt]` group of `nova.conf` define which type of CPU model is exposed to the hypervisor when using KVM: `cpu_mode` and `cpu_model`.

The `cpu_mode` option can take one of the following values: `none`, `host-passthrough`, `host-model`, and `custom`.

### Host model (default for KVM & QEMU)

If your `nova.conf` file contains `cpu_mode=host-model`, libvirt identifies the CPU model in `/usr/share/libvirt/cpu_map.xml` file that most closely matches the host, and requests additional CPU flags to complete the match. This configuration provides the maximum functionality and performance and maintains good reliability and compatibility if the guest is migrated to another host with slightly different host CPUs.

### Host pass through

If your `nova.conf` file contains `cpu_mode=host-passthrough`, libvirt tells KVM to pass through the host CPU with no modifications. The difference to `host-model`, instead of just matching feature flags, every last detail of the host CPU is matched. This gives the best performance, and can be important to some apps which check low level CPU details, but it comes at a cost with respect to migration. The guest can only be migrated to a matching host CPU.

## Custom

If your `nova.conf` file contains `cpu_mode=custom`, you can explicitly specify one of the supported named models using the `cpu_model` configuration option. For example, to configure the KVM guests to expose Nehalem CPUs, your `nova.conf` file should contain:

```
[libvirt]
cpu_mode = custom
cpu_model = Nehalem
```

## None (default for all libvirt-driven hypervisors other than KVM & QEMU)

If your `nova.conf` file contains `cpu_mode=none`, libvirt does not specify a CPU model. Instead, the hypervisor chooses the default model.

## Guest agent support

Use guest agents to enable optional access between compute nodes and guests through a socket, using the QMP protocol.

To enable this feature, you must set `hw_qemu_guest_agent=yes` as a metadata parameter on the image you wish to use to create the guest-agent-capable instances from. You can explicitly disable the feature by setting `hw_qemu_guest_agent=no` in the image metadata.

## KVM performance tweaks

The [VHostNet](#) kernel module improves network performance. To load the kernel module, run the following command as root:

```
# modprobe vhost_net
```

## Troubleshoot KVM

Trying to launch a new virtual machine instance fails with the `ERRORstate`, and the following error appears in the `/var/log/nova/nova-compute.log` file:

```
libvirtError: internal error no supported architecture for os type 'hvm'
```

This message indicates that the KVM kernel modules were not loaded.

If you cannot start VMs after installation without rebooting, the permissions might not be set correctly. This can happen if you load the KVM module before you install `nova-compute`. To check whether the group is set to `kvm`, run:

```
# ls -l /dev/kvm
```

If it is not set to `kvm`, run:

```
# udevadm trigger
```

## QEMU

From the perspective of the Compute service, the QEMU hypervisor is very similar to the KVM hypervisor. Both are controlled through libvirt, both support the same feature set,

and all virtual machine images that are compatible with KVM are also compatible with QEMU. The main difference is that QEMU does not support native virtualization. Consequently, QEMU has worse performance than KVM and is a poor choice for a production deployment.

The typical uses cases for QEMU are

- Running on older hardware that lacks virtualization support.
- Running the Compute service inside of a virtual machine for development or testing purposes, where the hypervisor does not support native virtualization for guests.

To enable QEMU, add these settings to `nova.conf`:

```
compute_driver = libvirt.LibvirtDriver

[libvirt]
virt_type = qemu
```

For some operations you may also have to install the **guestmount** utility:

On Ubuntu:

```
# apt-get install guestmount
```

On Red Hat Enterprise Linux, Fedora, or CentOS:

```
# yum install libguestfs-tools
```

On openSUSE:

```
# zypper install guestfs-tools
```

The QEMU hypervisor supports the following virtual machine image formats:

- Raw
- QEMU Copy-on-write (qcow2)
- VMware virtual machine disk format (vmdk)

## XenServer (and other XAPI based Xen variants)

This section describes XAPI managed hypervisors, and how to use them with OpenStack.

### Terminology

#### Xen

A hypervisor that provides the fundamental isolation between virtual machines. Xen is open source (GPLv2) and is managed by [XenProject.org](http://XenProject.org), a cross-industry organization and a Linux Foundation Collaborative project.

Xen is a component of many different products and projects. The hypervisor itself is very similar across all these projects, but the way that it is managed can be different, which can



the OpenStack software (much of which is customer-facing). This architecture is described in more detail later.

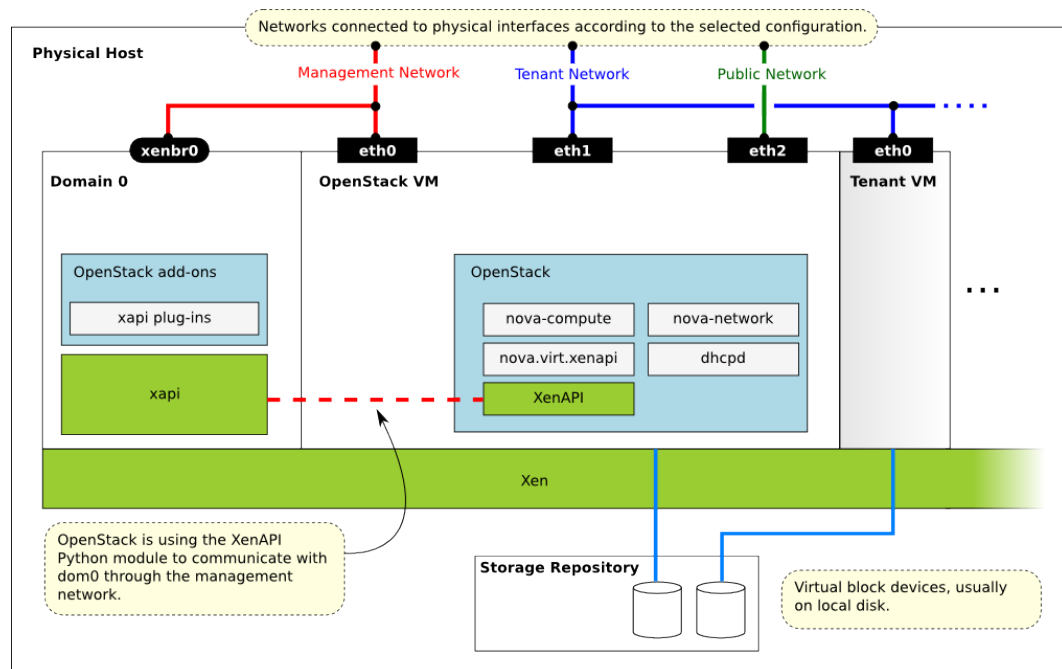
### Paravirtualized versus hardware virtualized domains

A Xen virtual machine can be paravirtualized (PV) or hardware virtualized (HVM). This refers to the interaction between Xen, domain 0, and the guest VM's kernel. PV guests are aware of the fact that they are virtualized and will co-operate with Xen and domain 0; this gives them better performance characteristics. HVM guests are not aware of their environment, and the hardware has to pretend that they are running on an unvirtualized machine. HVM guests do not need to modify the guest operating system, which is essential when running Windows.

In OpenStack, customer VMs may run in either PV or HVM mode. However, the OpenStack domU (that's the one running `nova-compute`) must be running in PV mode.

### XenAPI deployment architecture

A basic OpenStack deployment on a XAPI-managed server, assuming that the network provider is nova-network, looks like this:



Key things to note:

- The hypervisor: Xen
- Domain 0: runs XAPI and some small pieces from OpenStack, the XAPI plug-ins.
- OpenStack VM: The `Compute` service runs in a paravirtualized virtual machine, on the host under management. Each host runs a local instance of `Compute`. It is also running an instance of `nova-network`.

- OpenStack Compute uses the XenAPI Python library to talk to XAPI, and it uses the Management Network to reach from the OpenStack VM to Domain 0.

Some notes on the networking:

- The above diagram assumes FlatDHCP networking.
- There are three main OpenStack networks:
  - Management network: RabbitMQ, MySQL, inter-host communication, and compute-XAPI communication. Please note that the VM images are downloaded by the XenAPI plug-ins, so make sure that the OpenStack Image service is accessible through this network. It usually means binding those services to the management interface.
  - Tenant network: controlled by nova-network, this is used for tenant traffic.
  - Public network: floating IPs, public API endpoints.
- The networks shown here must be connected to the corresponding physical networks within the data center. In the simplest case, three individual physical network cards could be used. It is also possible to use VLANs to separate these networks. Please note, that the selected configuration must be in line with the networking model selected for the cloud. (In case of VLAN networking, the physical channels have to be able to forward the tagged traffic.)

## Further reading

Here are some of the resources available to learn more about Xen:

- Citrix XenServer official documentation: <http://docs.vmd.citrix.com/XenServer>
- What is Xen? by XenProject.org: [XenProject.org > Users > Cloud](#)
- Xen Hypervisor project: <http://www.xenproject.org/developers/teams/hypervisor.html>
- Xapi project: <http://www.xenproject.org/developers/teams/xapi.html>
- Further XenServer and OpenStack information: <http://wiki.openstack.org/XenServer>

## Install XenServer

Before you can run OpenStack with XenServer, you must install the hypervisor on [an appropriate server](#) .



### Note

Xen is a type 1 hypervisor: When your server starts, Xen is the first software that runs. Consequently, you must install XenServer before you install the operating system where you want to run OpenStack code. You then install `nova-compute` into a dedicated virtual machine on the host.

Use the following link to download XenServer's installation media:























### Note

In addition to the default VNC port numbers (5900 to 6000) specified in the above document, the following ports are also used: 6101, 6102, and 6105.

You must modify the ESXi firewall configuration to allow the VNC ports. Additionally, for the firewall modifications to persist after a reboot, you must create a custom vSphere Installation Bundle (VIB) which is then installed onto the running ESXi host or added to a custom image profile used to install ESXi hosts. For details about how to create a VIB for persisting the firewall configuration modifications, see [http://kb.vmware.com/selfservice/microsites/search.do?language=en\\_US&cmd=displayKC&externalId=2007381](http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2007381).



### Note

The VIB can be downloaded from <https://github.com/openstack-vmwareapi-team/Tools>.

8. To use multiple vCenter installations with OpenStack, each vCenter must be assigned to a separate availability zone. This is required as the OpenStack Block Storage VMDK driver does not currently work across multiple vCenter installations.

## VMware vCenter service account

OpenStack integration requires a vCenter service account with the following minimum permissions. Apply the permissions to the `Datacenter` root object, and select the **Propagate to Child Objects** option.

**Table 3.7. vCenter permissions tree**

All Privileges		
	Datastore	
		Allocate space
		Browse datastore
		Low level file operation
		Remove file
	Extension	
		Register extension
	Folder	
		Create folder
	Host	
		Configuration
	Network	
		Assign network
	Resource	
		Assign virtual machine to resource pool

		Migrate powered off virtual machine	
		Migrate powered on virtual machine	
	Virtual Machine		
		Configuration	
		Interaction	
		Inventory	
		Provisioning	
		Sessions	
		Snapshot management	
	vApp		
		Export	
		Import	

**VMware vCenter driver**

Use the VMware vCenter driver (VMwareVCDriver) to connect OpenStack Compute with vCenter. This recommended configuration enables access through vCenter to advanced vSphere features like vMotion, High Availability, and Dynamic Resource Scheduling (DRS).







type to `vmware`. Other valid hypervisor types include: `xen`, `qemu`, `lxc`, `uml`, and `hyperv`. Note that `qemu` is used for both QEMU and KVM hypervisor types.

```
$ glance image-create --name "ubuntu-thick-scsi" --disk-format vmdk \
--container-format bare \
--property vmware_adaptype="lsiLogic" \
--property vmware_disktype="preallocated" \
--property hypervisor_type="vmware" \
--property vmware_ostype="ubuntu64Guest" < ubuntuLTS-flat.vmdk
```

## Optimize images

Monolithic Sparse disks are considerably faster to download but have the overhead of an additional conversion step. When imported into ESX, sparse disks get converted to VMFS flat thin provisioned disks. The download and conversion steps only affect the first launched instance that uses the sparse disk image. The converted disk image is cached, so subsequent instances that use this disk image can simply use the cached version.

To avoid the conversion step (at the cost of longer download times) consider converting sparse disks to thin provisioned or preallocated disks before loading them into the OpenStack Image service.

Use one of the following tools to pre-convert sparse disks.

### vSphere CLI tools

Sometimes called the remote CLI or rCLI.

Assuming that the sparse disk is made available on a data store accessible by an ESX host, the following command converts it to preallocated format:

```
vmkfstools --server=ip_of_some_ESX_host -i /
vmfs/volumes/datastore1/sparse.vmdk /vmfs/
volumes/datastore1/converted.vmdk
```

Note that the `vifs` tool from the same CLI package can be used to upload the disk to be converted. The `vifs` tool can also be used to download the converted disk if necessary.

### vmkfstools directly on the ESX host

If the SSH service is enabled on an ESX host, the sparse disk can be uploaded to the ESX data store through `scp` and the `vmkfstools` local to the ESX host can use used to perform the conversion. After you log in to the host through `ssh`, run this command:

```
vmkfstools -i /vmfs/volumes/datastore1/sparse.
vmdk /vmfs/volumes/datastore1/converted.vmdk
```

### vmware-vdiskmanager

`vmware-vdiskmanager` is a utility that comes bundled with VMware Fusion and VMware Workstation. The following example converts a sparse disk to preallocated format:

```
'/Applications/VMware Fusion.app/Contents/
Library/vmware-vdiskmanager' -r sparse.vmdk -t 4
converted.vmdk
```







```
$ mkdir -p /opt/stack/vmware/wsd1/5.0
```

3. Change into the new directory.

```
$ cd /opt/stack/vmware/wsd1/5.0
```

4. Use your OS-specific tools to install a command-line tool that can download files like **wget**.
5. Download the files to the local file cache:

```
wget --no-check-certificate https://$VMWAREAPI_IP/sdk/vimService.wsdl
wget --no-check-certificate https://$VMWAREAPI_IP/sdk/vim.wsdl
wget --no-check-certificate https://$VMWAREAPI_IP/sdk/core-types.xsd
wget --no-check-certificate https://$VMWAREAPI_IP/sdk/query-messagetypes.xsd
wget --no-check-certificate https://$VMWAREAPI_IP/sdk/query-types.xsd
wget --no-check-certificate https://$VMWAREAPI_IP/sdk/vim-messagetypes.xsd
wget --no-check-certificate https://$VMWAREAPI_IP/sdk/vim-types.xsd
wget --no-check-certificate https://$VMWAREAPI_IP/sdk/reflect-messagetypes.xsd
wget --no-check-certificate https://$VMWAREAPI_IP/sdk/reflect-types.xsd
```

Because the `reflect-types.xsd` and `reflect-messagetypes.xsd` files do not fetch properly, you must stub out these files. Use the following XML listing to replace the missing file content. The XML parser underneath Python can be very particular and if you put a space in the wrong place, it can break the parser. Copy the following contents and formatting carefully.

```
<?xml version="1.0" encoding="UTF-8"?>
<schema
  targetNamespace="urn:reflect"
  xmlns="http://www.w3.org/2001/XMLSchema"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema"
  elementFormDefault="qualified">
</schema>
```

6. Now that the files are locally present, tell the driver to look for the SOAP service WSDLs in the local file system and not on the remote vSphere server. Add the following setting to the `nova.conf` file for your `nova-compute` node:

```
[vmware]
wsdl_location=file:///opt/stack/vmware/wsd1/5.0/vimService.wsdl
```

Alternatively, download the version appropriate SDK from <http://www.vmware.com/support/developer/vc-sdk/> and copy it to the `/opt/stack/vmware` file. Make sure that the WSDL is available, in for example `/opt/stack/vmware/SDK/wsd1/vim25/vimService.wsdl`. You must point `nova.conf` to fetch this WSDL file from the local file system by using a URL.

When using the VMwareVCDriver (vCenter) with OpenStack Compute with vSphere version 5.0 or earlier, `nova.conf` must include the following extra config option:

```
[vmware]
wsdl_location=file:///opt/stack/vmware/SDK/wsd1/vim25/vimService.wsdl
```



## Configure Hyper-V virtual switching

Information regarding the Hyper-V virtual Switch can be located here: <http://technet.microsoft.com/en-us/library/hh831823.aspx>

To quickly enable an interface to be used as a Virtual Interface the following PowerShell may be used:

```
PS C:\>$if = Get-NetIPAddress -IPAddress 192* | Get-NetIPInterface
PS C:\>New-VMSwitch -NetAdapterName $if.ifAlias -Name YOUR_BRIDGE_NAME -
AllowManagementOS $false
```



### Note

It is very important to make sure that when you are using an Hyper-V node with only 1 NIC the `-AllowManagementOS` option is set on `True`, otherwise you will lose connectivity to the Hyper-V node.

## Enable iSCSI initiator service

To prepare the Hyper-V node to be able to attach to volumes provided by cinder you must first make sure the Windows iSCSI initiator service is running and started automatically.

```
PS C:\>Set-Service -Name MSiSCSI -StartupType Automatic
PS C:\>Start-Service MSiSCSI
```

## Configure shared nothing live migration

Detailed information on the configuration of live migration can be found here: <http://technet.microsoft.com/en-us/library/jj134199.aspx>

The following outlines the steps of shared nothing live migration.

1. The target hosts ensures that live migration is enabled and properly configured in Hyper-V.
2. The target hosts checks if the image to be migrated requires a base VHD and pulls it from the Image service if not already available on the target host.
3. The source hosts ensures that live migration is enabled and properly configured in Hyper-V.
4. The source hosts initiates a Hyper-V live migration.
5. The source hosts communicates to the manager the outcome of the operation.

The following two configuration options/flags are needed in order to support Hyper-V live migration and must be added to your `nova.conf` on the Hyper-V compute node:

- `instances_shared_storage = False`

This needed to support "shared nothing" Hyper-V live migrations. It is used in `nova/compute/manager.py`

- `limit_cpu_features = True`



## Requirements

### Python

Python 2.7 32bit must be installed as most of the libraries are not working properly on the 64bit version.

#### Procedure 3.2. Setting up Python prerequisites

1. Download and then install it using the MSI installer from here:

<http://www.python.org/ftp/python/2.7.3/python-2.7.3.msi>

```
PS C:\> $src = "http://www.python.org/ftp/python/2.7.3/python-2.7.3.msi"
PS C:\> $dest = "$env:temp\python-2.7.3.msi"
PS C:\> Invoke-WebRequest -Uri $src -OutFile $dest
PS C:\> Unblock-File $dest
PS C:\> Start-Process $dest
```

2. Make sure that the Python and Python\Scripts paths are set up in the PATH environment variable.

```
PS C:\>$oldPath = [System.Environment]::GetEnvironmentVariable("Path")
PS C:\>$newPath = $oldPath + ";C:\python27\C:\python27\Scripts\"
PS C:\>[System.Environment]::SetEnvironmentVariable("Path", $newPath,
[System.EnvironmentVariableTarget]::User
```

### Python dependencies

The following packages need to be downloaded and manually installed:

<b>setuptools</b>	<a href="http://pypi.python.org/packages/2.7/s/setup-tools/setuptools-0.6c11.win32-py2.7.exe">http://pypi.python.org/packages/2.7/s/setup-tools/setuptools-0.6c11.win32-py2.7.exe</a>
<b>pip</b>	<a href="http://pip.readthedocs.org/en/latest/installing.html">http://pip.readthedocs.org/en/latest/installing.html</a>
<b>MySQL-python</b>	<a href="http://codegood.com/download/10/">http://codegood.com/download/10/</a>
<b>PyWin32</b>	<a href="http://sourceforge.net/projects/pywin32/files/pywin32/Build%20217/pywin32-217.win32-py2.7.exe">http://sourceforge.net/projects/pywin32/files/pywin32/Build%20217/pywin32-217.win32-py2.7.exe</a>
<b>Greenlet</b>	<a href="http://www.lfd.uci.edu/~gohlke/pythonlibs/#greenlet">http://www.lfd.uci.edu/~gohlke/pythonlibs/#greenlet</a>
<b>PyCrypto</b>	<a href="http://www.voidspace.org.uk/downloads/pycrypt-to26/pycrypto-2.6.win32-py2.7.exe">http://www.voidspace.org.uk/downloads/pycrypt-to26/pycrypto-2.6.win32-py2.7.exe</a>

The following packages must be installed with pip:

- ecdsa
- amqp
- wmi

```
PS C:\> pip install ecdsa
PS C:\> pip install amqp
PS C:\> pip install wmi
```

## Other dependencies

`qemu-img` is required for some of the image related operations. You can get it from here: <http://qemu.weinnetz.de/>. You must make sure that the `qemu-img` path is set in the `PATH` environment variable.

Some Python packages need to be compiled, so you may use MinGW or Visual Studio. You can get MinGW from here: <http://sourceforge.net/projects/mingw/>. You must configure which compiler to be used for this purpose by using the `distutils.cfg` file in `$Python27\Lib\distutils`, which can contain:

```
[build]
compiler = mingw32
```

As a last step for setting up MinGW, make sure that the MinGW binaries' directories are set up in `PATH`.

## Install Nova-compute

### Download the nova code

1. Use Git to download the necessary source code. The installer to run Git on Windows can be downloaded here:
 

<https://github.com/msysgit/msysgit/releases/download/Git-1.9.2-preview20140411/Git-1.9.2-preview20140411.exe>
2. Download the installer. Once the download is complete, run the installer and follow the prompts in the installation wizard. The default should be acceptable for the needs of the document.

```
PS C:\>$src = "https://github.com/msysgit/msysgit/releases/download/Git-1.9.2-preview20140411/Git-1.9.2-preview20140411.exe"
PS C:\>$dest = "$env:temp\Git-1.9.2-preview20140411.exe"
PS C:\>Invoke-WebRequest -Uri $src -OutFile $dest
PS C:\>Unblock-File $dest
PS C:\>Start-Process $dest
```

3. Run the following to clone the Nova code.

```
PS C:\>git.exe clone https://github.com/openstack/nova.git
```

### Install nova-compute service

To install `Nova-compute`, run:

```
PS C:\>cd c:\Nova
PS C:\>python setup.py install
```

### Configure nova-compute

The `nova.conf` file must be placed in `C:\etc\nova` for running OpenStack on Hyper-V. Below is a sample `nova.conf` for Windows:

```
[DEFAULT]
[DEFAULT]
auth_strategy = keystone
image_service = nova.image.glance.GlanceImageService
```

```

compute_driver = nova.virt.hyperv.driver.HyperVDriver
volume_api_class = nova.volume.cinder.API
fake_network = true
instances_path = C:\Program Files (x86)\OpenStack\Instances
glance_api_servers = IP_ADDRESS:9292
use_cow_images = true
force_config_drive = false
injected_network_template = C:\Program Files (x86)\OpenStack\Nova\etc\
interfaces.template
policy_file = C:\Program Files (x86)\OpenStack\Nova\etc\policy.json
mkisofs_cmd = C:\Program Files (x86)\OpenStack\Nova\bin\mkisofs.exe
verbose = false
allow_resize_to_same_host = true
running_deleted_instance_action = reap
running_deleted_instance_poll_interval = 120
resize_confirm_window = 5
resume_guests_state_on_host_boot = true
rpc_response_timeout = 1800
lock_path = C:\Program Files (x86)\OpenStack\Log\
rpc_backend = nova.openstack.common.rpc.impl_kombu
rabbit_host = IP_ADDRESS
rabbit_port = 5672
rabbit_userid = guest
rabbit_password = Passw0rd
logdir = C:\Program Files (x86)\OpenStack\Log\
logfile = nova-compute.log
instance_usage_audit = true
instance_usage_audit_period = hour
network_api_class = nova.network.neutronv2.api.API
[neutron]
url = http://IP_ADDRESS:9696
auth_strategy = keystone
admin_tenant_name = service
admin_username = neutron
admin_password = Passw0rd
admin_auth_url = http://IP_ADDRESS:35357/v2.0
[hyperv]
vswitch_name = newVSwitch0
limit_cpu_features = false
config_drive_inject_password = false
qemu_img_cmd = C:\Program Files (x86)\OpenStack\Nova\bin\qemu-img.exe
config_drive_cdrom = true
dynamic_memory_ratio = 1
enable_instance_metrics_collection = true
[rdp]
enabled = true
html5_proxy_base_url = https://IP_ADDRESS:4430

```

[Table 3.31, "Description of HyperV configuration options" \[296\]](#) contains a reference of all options for hyper-v.

### Prepare images for use with Hyper-V

Hyper-V currently supports only the VHD and VHDX file format for virtual machine instances. Detailed instructions for installing virtual machines on Hyper-V can be found here:

<http://technet.microsoft.com/en-us/library/cc772480.aspx>

Once you have successfully created a virtual machine, you can then upload the image to glance using the native glance-client:





By default, the `scheduler_driver` is configured as a filter scheduler, as described in the next section. In the default configuration, this scheduler considers hosts that meet all the following criteria:

- Have not been attempted for scheduling purposes (`RetryFilter`).
- Are in the requested availability zone (`AvailabilityZoneFilter`).
- Have sufficient RAM available (`RamFilter`).
- Can service the request (`ComputeFilter`).
- Satisfy the extra specs associated with the instance type (`ComputeCapabilitiesFilter`).
- Satisfy any architecture, hypervisor type, or virtual machine mode properties specified on the instance's image properties (`ImagePropertiesFilter`).
- Are on a different host than other instances of a group (if requested) (`ServerGroupAntiAffinityFilter`).
- Are in a set of group hosts (if requested) (`ServerGroupAffinityFilter`).

The scheduler caches its list of available hosts; use the `scheduler_driver_task_period` option to specify how often the list is updated.



### Note

Do not configure `service_down_time` to be much smaller than `scheduler_driver_task_period`; otherwise, hosts appear to be dead while the host list is being cached.

For information about the volume scheduler, see the Block Storage section of [OpenStack Cloud Administrator Guide](#).

The scheduler chooses a new host when an instance is migrated.

When evacuating instances from a host, the scheduler service honors the target host defined by the administrator on the `evacuate` command. If a target is not defined by the administrator, the scheduler determines the target host. For information about instance evacuation, see [Evacuate instances](#) section of the [OpenStack Cloud Administrator Guide](#).

## Filter scheduler

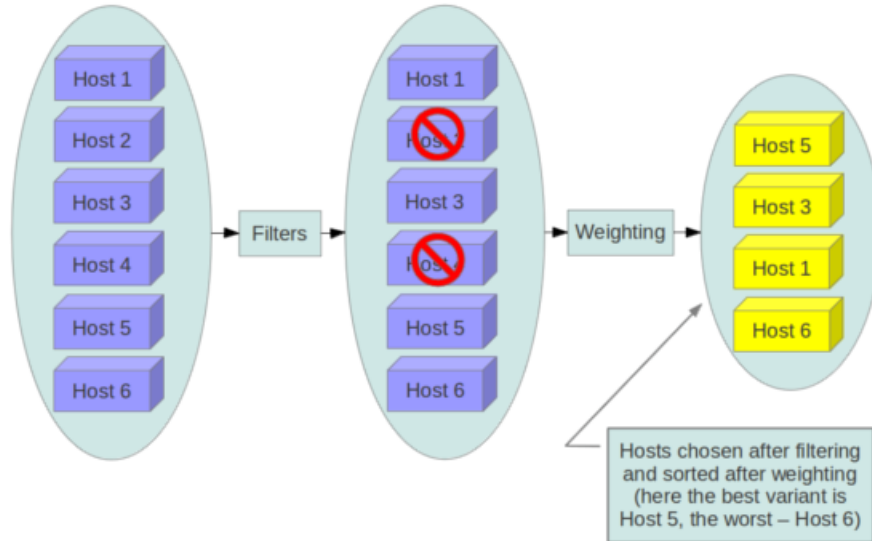
The filter scheduler (`nova.scheduler.filter_scheduler.FilterScheduler`) is the default scheduler for scheduling virtual machine instances. It supports filtering and weighting to make informed decisions on where a new instance should be created.

## Filters

When the filter scheduler receives a request for a resource, it first applies filters to determine which hosts are eligible for consideration when dispatching a resource. Filters are binary: either a host is accepted by the filter, or it is rejected. Hosts that are accepted by the

filter are then processed by a different algorithm to decide which hosts to use for that request, described in the [Weights](#) section.

**Figure 3.2. Filtering**



The `scheduler_available_filters` configuration option in `nova.conf` provides the Compute service with the list of the filters that are used by the scheduler. The default setting specifies all of the filter that are included with the Compute service:

```
scheduler_available_filters = nova.scheduler.filters.all_filters
```

This configuration option can be specified multiple times. For example, if you implemented your own custom filter in Python called `myfilter.MyFilter` and you wanted to use both the built-in filters and your custom filter, your `nova.conf` file would contain:

```
scheduler_available_filters = nova.scheduler.filters.all_filters
scheduler_available_filters = myfilter.MyFilter
```

The `scheduler_default_filters` configuration option in `nova.conf` defines the list of filters that are applied by the `nova-scheduler` service. The default filters are:

```
scheduler_default_filters = RetryFilter, AvailabilityZoneFilter, RamFilter,
ComputeFilter, ComputeCapabilitiesFilter, ImagePropertiesFilter,
ServerGroupAntiAffinityFilter, ServerGroupAffinityFilter
```

The following sections describe the available filters.

### AggregateCoreFilter

Filters host by CPU core numbers with a per-aggregate `cpu_allocation_ratio` value. If the per-aggregate value is not found, the value falls back to the global setting. If the host is in more than one aggregate and more than one value is found, the minimum value will be used. For information about how to use this filter, see [the section called "Host aggregates and availability zones"](#) [271]. See also [the section called "CoreFilter"](#) [261].

## AggregateDiskFilter

Filters host by disk allocation with a per-aggregate `disk_allocation_ratio` value. If the per-aggregate value is not found, the value falls back to the global setting. If the host is in more than one aggregate and more than one value is found, the minimum value will be used. For information about how to use this filter, see [the section called “Host aggregates and availability zones” \[271\]](#). See also [the section called “DiskFilter” \[262\]](#).

## AggregateImagePropertiesIsolation

Matches properties defined in an image's metadata against those of aggregates to determine host matches:

- If a host belongs to an aggregate and the aggregate defines one or more metadata that matches an image's properties, that host is a candidate to boot the image's instance.
- If a host does not belong to any aggregate, it can boot instances from all images.

For example, the following aggregate `myWinAgg` has the Windows operating system as metadata (named 'windows'):

```
$ nova aggregate-details MyWinAgg
+-----+-----+-----+-----+-----+
| Id | Name      | Availability Zone | Hosts      | Metadata  |
+-----+-----+-----+-----+-----+
| 1  | MyWinAgg | None              | 'sf-devel' | 'os=windows' |
+-----+-----+-----+-----+-----+
```

In this example, because the following Win-2012 image has the windows property, it boots on the `sf-devel` host (all other filters being equal):

```
$ glance image-show Win-2012
+-----+-----+
| Property          | Value                               |
+-----+-----+
| Property 'os'     | windows                             |
| checksum           | f8a2eeee2dc65b3d9b6e63678955bd83   |
| container_format  | ami                                  |
| created_at        | 2013-11-14T13:24:25                 |
| ...               | ...                                 |
```

You can configure the `AggregateImagePropertiesIsolation` filter by using the following options in the `nova.conf` file:

```
# Considers only keys matching the given namespace (string).
aggregate_image_properties_isolation_namespace = <None>

# Separator used between the namespace and keys (string).
aggregate_image_properties_isolation_separator = .
```

## AggregateInstanceExtraSpecsFilter

Matches properties defined in extra specs for an instance type against admin-defined properties on a host aggregate. Works with specifications that are scoped with `aggregate_instance_extra_specs`. For backward compatibility, also works with non-



## ComputeCapabilitiesFilter

Matches properties defined in extra specs for an instance type against compute capabilities.

If an extra specs key contains a colon (:), anything before the colon is treated as a namespace and anything after the colon is treated as the key to be matched. If a namespace is present and is not `capabilities`, the filter ignores the namespace. For backward compatibility, also treats the extra specs key as the key to be matched if no namespace is present; this action is highly discouraged because it conflicts with [AggregateInstanceExtraSpecsFilter](#) filter when you enable both filters.

## ComputeFilter

Passes all hosts that are operational and enabled.

In general, you should always enable this filter.

## CoreFilter

Only schedules instances on hosts if sufficient CPU cores are available. If this filter is not set, the scheduler might over-provision a host based on cores. For example, the virtual cores running on an instance may exceed the physical cores.

You can configure this filter to enable a fixed amount of vCPU overcommitment by using the `cpu_allocation_ratio` configuration option in `nova.conf`. The default setting is:

```
cpu_allocation_ratio = 16.0
```

With this setting, if 8 vCPUs are on a node, the scheduler allows instances up to 128 vCPU to be run on that node.

To disallow vCPU overcommitment set:

```
cpu_allocation_ratio = 1.0
```



### Note

The Compute API always returns the actual number of CPU cores available on a compute node regardless of the value of the `cpu_allocation_ratio` configuration key. As a result changes to the `cpu_allocation_ratio` are not reflected via the command line clients or the dashboard. Changes to this configuration key are only taken into account internally in the scheduler.

## NUMATopologyFilter

Filters hosts based on the NUMA topology that was specified for the instance through the use of flavor `extra_specs` in combination with the image properties, as described in detail in the related nova-spec document: Filter will try to match the exact NUMA cells of the instance to those of the host. It will consider the standard over-subscription limits each cell, and provide limits to the compute host accordingly.





the key and an arbitrary name as the value. Using the `nova` command-line tool, use the `--hint` flag. For example:

```
$ nova boot --image IMAGE_ID --flavor 1 --hint group=foo server-1
```

This filter should not be enabled at the same time as [GroupAffinityFilter](#) or neither filter will work properly.

## ImagePropertiesFilter

Filters hosts based on properties defined on the instance's image. It passes hosts that can support the specified image properties contained in the instance. Properties include the architecture, hypervisor type, hypervisor version (for Xen hypervisor type only), and virtual machine mode.

For example, an instance might require a host that runs an ARM-based processor, and QEMU as the hypervisor. You can decorate an image with these properties by using:

```
$ glance image-update img-uuid --property architecture=arm --property hypervisor_type=qemu
```

The image properties that the filter checks for are:

- `architecture`: describes the machine architecture required by the image. Examples are `i686`, `x86_64`, `arm`, and `ppc64`.
- `hypervisor_type`: describes the hypervisor required by the image. Examples are `xen`, `qemu`, and `xenapi`.



### Note

`qemu` is used for both QEMU and KVM hypervisor types.

- `hypervisor_version_requires`: describes the hypervisor version required by the image. The property is supported for Xen hypervisor type only. It can be used to enable support for multiple hypervisor versions, and to prevent instances with newer Xen tools from being provisioned on an older version of a hypervisor. If available, the property value is compared to the hypervisor version of the compute host.

To filter the hosts by the hypervisor version, add the `hypervisor_version_requires` property on the image as metadata and pass an operator and a required hypervisor version as its value:

```
$ glance image-update img-uuid --property hypervisor_type=xen --property hypervisor_version_requires=">=4.3"
```

- `vm_mode`: describes the hypervisor application binary interface (ABI) required by the image. Examples are `xen` for Xen 3.0 paravirtual ABI, `hvm` for native ABI, `uml` for User Mode Linux paravirtual ABI, `exe` for container virt executable ABI.

## IsolatedHostsFilter

Allows the admin to define a special (isolated) set of images and a special (isolated) set of hosts, such that the isolated images can only run on the iso-





Using the `nova` command-line tool, use the `--hint` flag:

```
$ nova boot --image 827d564a-e636-4fc4-a376-d36f7ebe1747 \
--flavor 1 --hint query='[">=", "$free_ram_mb", 1024]' server1
```

With the API, use the `os:scheduler_hints` key:

```
{
  "server": {
    "name": "server-1",
    "imageRef": "cedef40a-ed67-4d10-800e-17455edce175",
    "flavorRef": "1"
  },
  "os:scheduler_hints": {
    "query": "[>=, $free_ram_mb, 1024]"
  }
}
```

## MetricsFilter

Filters hosts based on meters `weight_setting`. Only hosts with the available meters are passed so that the metrics weigher will not fail due to these hosts.

## NumInstancesFilter

Hosts that have more instances running than specified by the `max_instances_per_host` option are filtered out when this filter is in place.

## PciPassthroughFilter

The filter schedules instances on a host if the host has devices that meet the device requests in the `extra_specs` attribute for the flavor.

## RamFilter

Only schedules instances on hosts that have sufficient RAM available. If this filter is not set, the scheduler may over provision a host based on RAM (for example, the RAM allocated by virtual machine instances may exceed the physical RAM).

You can configure this filter to enable a fixed amount of RAM overcommitment by using the `ram_allocation_ratio` configuration option in `nova.conf`. The default setting is:

```
ram_allocation_ratio = 1.5
```

This setting enables 1.5 GB instances to run on any compute node with 1 GB of free RAM.

## RetryFilter

Filters out hosts that have already been attempted for scheduling purposes. If the scheduler selects a host to respond to a service request, and the host fails to respond to the request, this filter prevents the scheduler from retrying that host for the service request.

This filter is only useful if the `scheduler_max_attempts` configuration option is set to a value greater than zero.

## SameHostFilter

Schedules the instance on the same host as another instance in a set of instances. To take advantage of this filter, the requester must pass a scheduler hint, using `same_host` as the key and a list of instance UUIDs as the value. This filter is the opposite of the `DifferentHostFilter`. Using the `nova` command-line tool, use the `--hint` flag:

```
$ nova boot --image cedef40a-ed67-4d10-800e-17455edce175 --flavor 1 \
  --hint same_host=a0cf03a5-d921-4877-bb5c-86d26cf818e1 \
  --hint same_host=8c19174f-4220-44f0-824a-cd1eeef10287 server-1
```

With the API, use the `os:scheduler_hints` key:

```
{
  "server": {
    "name": "server-1",
    "imageRef": "cedef40a-ed67-4d10-800e-17455edce175",
    "flavorRef": "1"
  },
  "os:scheduler_hints": {
    "same_host": [
      "a0cf03a5-d921-4877-bb5c-86d26cf818e1",
      "8c19174f-4220-44f0-824a-cd1eeef10287"
    ]
  }
}
```

## ServerGroupAffinityFilter

The `ServerGroupAffinityFilter` ensures that an instance is scheduled on to a host from a set of group hosts. To take advantage of this filter, the requester must create a server group with an affinity policy, and pass a scheduler hint, using `group` as the key and the server group UUID as the value. Using the `nova` command-line tool, use the `--hint` flag. For example:

```
$ nova server-group-create --policy affinity group-1
$ nova boot --image IMAGE_ID --flavor 1 --hint group=SERVER_GROUP_UUID
server-1
```

## ServerGroupAntiAffinityFilter

The `ServerGroupAntiAffinityFilter` ensures that each instance in a group is on a different host. To take advantage of this filter, the requester must create a server group with an `anti-affinity` policy, and pass a scheduler hint, using `group` as the key and the server group UUID as the value. Using the `nova` command-line tool, use the `--hint` flag. For example:

```
$ nova server-group-create --policy anti-affinity group-1
$ nova boot --image IMAGE_ID --flavor 1 --hint group=SERVER_GROUP_UUID
server-1
```

## SimpleCIDRAffinityFilter

Schedules the instance based on host IP subnet range. To take advantage of this filter, the requester must specify a range of valid IP address in CIDR format, by passing two scheduler hints:

**build\_near\_host\_ip** The first IP address in the subnet (for example, 192.168.1.1)

**cidr** The CIDR that corresponds to the subnet (for example, /24)

Using the **nova** command-line tool, use the `--hint` flag. For example, to specify the IP subnet 192.168.1.1/24

```
$ nova boot --image cedef40a-ed67-4d10-800e-17455edce175 --flavor 1 \
--hint build_near_host_ip=192.168.1.1 --hint cidr=/24 server-1
```

With the API, use the `os:scheduler_hints` key:

```
{
  "server": {
    "name": "server-1",
    "imageRef": "cedef40a-ed67-4d10-800e-17455edce175",
    "flavorRef": "1"
  },
  "os:scheduler_hints": {
    "build_near_host_ip": "192.168.1.1",
    "cidr": "24"
  }
}
```

## TrustedFilter

Filters hosts based on their trust. Only passes hosts that meet the trust requirements specified in the instance properties.

## TypeAffinityFilter

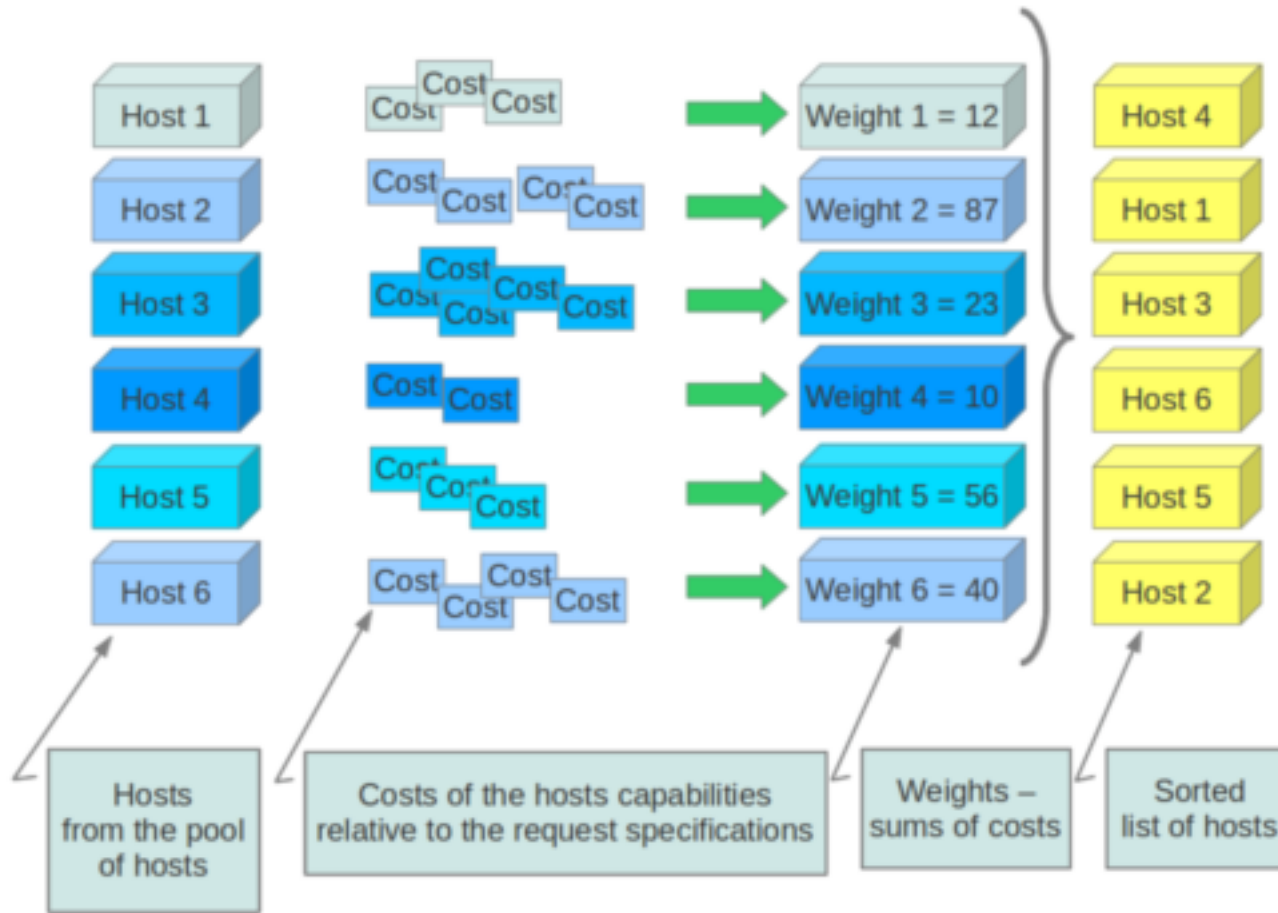
Dynamically limits hosts to one instance type. An instance can only be launched on a host, if no instance with different instances types are running on it, or if the host has no running instances at all.

## Weights

When resourcing instances, the filter scheduler filters and weights each host in the list of acceptable hosts. Each time the scheduler selects a host, it virtually consumes resources on it, and subsequent selections are adjusted accordingly. This process is useful when the customer asks for the same large amount of instances, because weight is computed for each requested instance.

All weights are normalized before being summed up; the host with the largest weight is given the highest priority.

Figure 3.3. Weighting hosts



If cells are used, cells are weighted by the scheduler in the same manner as hosts.

Hosts and cells are weighted based on the following options in the `/etc/nova/nova.conf` file:

Table 3.9. Host weighting options

Section	Option	Description
[DEFAULT]	<code>ram_weight_multiplier</code>	By default, the scheduler spreads instances across all hosts evenly. Set the <code>ram_weight_multiplier</code> option to a negative number if you prefer stacking instead of spreading. Use a floating-point value.
[DEFAULT]	<code>scheduler_host_subset_size</code>	Now instances are scheduled on a host that is chosen randomly from a subset of the N best hosts. This property defines the subset size from which a host is chosen. A value of 1 chooses the first host returned by the weighting functions. This value must be at least 1. A value less than 1 is ignored, and 1 is used instead. Use an integer value.
[DEFAULT]	<code>scheduler_weight_class</code>	Defaults to <code>nova.scheduler.weights.all_weighters</code> , which selects the <code>RamWeigher</code> and <code>MetricsWeigher</code> . Hosts are then weighted and sorted with the largest weight winning.
[metrics]	<code>weight_multiplier</code>	Multiplier for weighting meters. Use a floating-point value.



## Host aggregates and availability zones

Host aggregates are a mechanism for partitioning hosts in an OpenStack cloud, or a region of an OpenStack cloud, based on arbitrary characteristics. Examples where an administrator may want to do this include where a group of hosts have additional hardware or performance characteristics.

Host aggregates are not explicitly exposed to users. Instead administrators map flavors to host aggregates. Administrators do this by setting metadata on a host aggregate, and matching flavor extra specifications. The scheduler then endeavors to match user requests for instance of the given flavor to a host aggregate with the same key-value pair in its metadata. Compute nodes can be in more than one host aggregate.

Administrators are able to optionally expose a host aggregate as an availability zone. Availability zones are different from host aggregates in that they are explicitly exposed to the user, and hosts can only be in a single availability zone. Administrators can configure a default availability zone where instances will be scheduled when the user fails to specify one.

### Command-line interface

The **nova** command-line tool supports the following aggregate-related commands.

<b>nova aggregate-list</b>	Print a list of all aggregates.
<b>nova aggregate-create</b> <i>&lt;name&gt;</i> <i>[availability-zone]</i>	Create a new aggregate named <i>&lt;name&gt;</i> , and optionally in availability zone <i>[availability-zone]</i> if specified. The command returns the ID of the newly created aggregate. Hosts can be made available to multiple host aggregates. Be careful when adding a host to an additional host aggregate when the host is also in an availability zone. Pay attention when using the <b>aggregate-set-metadata</b> and <b>aggregate-update</b> commands to avoid user confusion when they boot instances in different availability zones. An error occurs if you cannot add a particular host to an aggregate zone for which it is not intended.
<b>nova aggregate-delete</b> <i>&lt;id&gt;</i>	Delete an aggregate with id <i>&lt;id&gt;</i> .
<b>nova aggregate-details</b> <i>&lt;id&gt;</i>	Show details of the aggregate with id <i>&lt;id&gt;</i> .
<b>nova aggregate-add-host</b> <i>&lt;id&gt;</i> <i>&lt;host&gt;</i>	Add host with name <i>&lt;host&gt;</i> to aggregate with id <i>&lt;id&gt;</i> .
<b>nova aggregate-remove-host</b> <i>&lt;id&gt;</i> <i>&lt;host&gt;</i>	Remove the host with name <i>&lt;host&gt;</i> from the aggregate with id <i>&lt;id&gt;</i> .
<b>nova aggregate-set-metadata</b> <i>&lt;id&gt;</i> <i>&lt;key=value&gt;</i> <i>[&lt;key=value&gt; ...]</i>	Add or update metadata (key-value pairs) associated with the aggregate with id <i>&lt;id&gt;</i> .
<b>nova aggregate-update</b> <i>&lt;id&gt;</i> <i>&lt;name&gt;</i> <i>[availability_zone]</i>	Update the name and availability zone (optional) for the aggregate.

<b>nova host-list</b>	List all hosts by service.
<b>nova host-update <code>--maintenance [enable   disable]</code></b>	Put/resume host into/from maintenance.



### Note

Only administrators can access these commands. If you try to use these commands and the user name and tenant that you use to access the Compute service do not have the admin role or the appropriate privileges, these errors occur:

```
ERROR: Policy doesn't allow compute_extension:aggregates to be performed. (HTTP 403) (Request-ID: req-299fbff6-6729-4cef-93b2-e7e1f96b4864)
```

```
ERROR: Policy doesn't allow compute_extension:hosts to be performed. (HTTP 403) (Request-ID: req-ef2400f6-6776-4ea3-b6f1-7704085c27d1)
```

## Configure scheduler to support host aggregates

One common use case for host aggregates is when you want to support scheduling instances to a subset of compute hosts because they have a specific capability. For example, you may want to allow users to request compute hosts that have SSD drives if they need access to faster disk I/O, or access to compute hosts that have GPU cards to take advantage of GPU-accelerated code.

To configure the scheduler to support host aggregates, the `scheduler_default_filters` configuration option must contain the `AggregateInstanceExtraSpecsFilter` in addition to the other filters used by the scheduler. Add the following line to `/etc/nova/nova.conf` on the host that runs the `nova-scheduler` service to enable host aggregates filtering, as well as the other filters that are typically enabled:

```
scheduler_default_filters=AggregateInstanceExtraSpecsFilter,RetryFilter,AvailabilityZoneFilter,RamFilter,ComputeFilter,ComputeCapabilitiesFilter,ImagePropertiesFilter,ServerGroupAntiAffinityFilter,ServerGroupAffinityFilter
```

### Example: Specify compute hosts with SSDs

This example configures the Compute service to enable users to request nodes that have solid-state drives (SSDs). You create a `fast-io` host aggregate in the `nova` availability zone and you add the `ssd=true` key-value pair to the aggregate. Then, you add the `node1`, and `node2` compute nodes to it.

```
$ nova aggregate-create fast-io nova
+-----+-----+-----+-----+-----+
| Id  | Name   | Availability Zone | Hosts  | Metadata |
+-----+-----+-----+-----+-----+
| 1   | fast-io | nova              |        |           |
+-----+-----+-----+-----+-----+

$ nova aggregate-set-metadata 1 ssd=true
+-----+-----+-----+-----+-----+
| Id  | Name   | Availability Zone | Hosts  | Metadata          |
+-----+-----+-----+-----+-----+
| 1   | fast-io | nova              | []     | {'ssd': 'true'}  |
```



```

+-----+
$ nova aggregate-add-host 1 node1
+-----+
| Id | Name      | Availability Zone | Hosts           | Metadata          |
+-----+
| 1  | fast-io  | nova             | [u'node1']     | {u'ssd': u'true'} |
+-----+

$ nova aggregate-add-host 1 node2
+-----+
| Id | Name      | Availability Zone | Hosts           | Metadata          |
+-----+
| 1  | fast-io  | nova             | [u'node1', u'node2'] | {u'ssd': u'true'} |
+-----+

```

Use the **nova flavor-create** command to create the *ssd.large* flavor called with an ID of 6, 8 GB of RAM, 80 GB root disk, and four vCPUs.

```

$ nova flavor-create ssd.large 6 8192 80 4
+-----+
| ID | Name        | Memory_MB | Disk | Ephemeral | Swap | VCPUs | RXTX_Factor |
| Is_Public |
+-----+
| 6  | ssd.large  | 8192      | 80  | 0          |      | 4     | 1.0          |
| True |
+-----+

```

Once the flavor is created, specify one or more key-value pairs that match the key-value pairs on the host aggregates with scope `aggregate_instance_extra_specs`. In this case, that is the `aggregate_instance_extra_specs:ssd=true` key-value pair. Setting a key-value pair on a flavor is done using the **nova flavor-key** command.

```
$ nova flavor-key ssd.large set aggregate_instance_extra_specs:ssd=true
```

Once it is set, you should see the `extra_specs` property of the *ssd.large* flavor populated with a key of `ssd` and a corresponding value of `true`.

```

$ nova flavor-show ssd.large
+-----+
| Property                | Value          |
+-----+
| OS-FLV-DISABLED:disabled | False          |
| OS-FLV-EXT-DATA:ephemeral | 0              |
| disk                     | 80             |
| extra_specs              | {u'aggregate_instance_extra_specs:ssd': u'true'} |
+-----+

```

```

| id | 6
| name | ssd.large
| os-flavor-access:is_public | True
| ram | 8192
| rxtx_factor | 1.0
| swap |
| vcpus | 4
+-----+
+-----+

```

Now, when a user requests an instance with the `ssd.large` flavor, the scheduler only considers hosts with the `ssd=true` key-value pair. In this example, these are `node1` and `node2`.

## XenServer hypervisor pools to support live migration

When using the XenAPI-based hypervisor, the Compute service uses host aggregates to manage XenServer Resource pools, which are used in supporting live migration.

## Configuration reference

To customize the Compute scheduler, use the configuration option settings documented in [Table 3.52, "Description of scheduler configuration options" \[309\]](#).

## Cells

*Cells* functionality enables you to scale an OpenStack Compute cloud in a more distributed fashion without having to use complicated technologies like database and message queue clustering. It supports very large deployments.

When this functionality is enabled, the hosts in an OpenStack Compute cloud are partitioned into groups called cells. 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.

The `nova-cells` service handles communication between cells and selects cells for new instances. This service is required for every cell. Communication between cells is pluggable, and currently the only option is communication through RPC.

Cells scheduling is separate from host scheduling. `nova-cells` first picks a cell. Once a cell is selected and the new build request reaches its `nova-cells` service, it is sent over to the host scheduler in that cell and the build proceeds as it would have without cells.



### Warning

Cell functionality is currently considered experimental.

## Cell configuration options

Cells are disabled by default. All cell-related configuration options appear in the `[cells]` section in `nova.conf`. The following cell-related options are currently supported:

<b>enable</b>	Set to <code>True</code> to turn on cell functionality. Default is <code>false</code> .
<b>name</b>	Name of the current cell. Must be unique for each cell.
<b>capabilities</b>	List of arbitrary <i>key=value</i> pairs defining capabilities of the current cell. Values include <code>hypervisor=xenserver;kvm,os=linux;windows</code> .
<b>call_timeout</b>	How long in seconds to wait for replies from calls between cells.
<b>scheduler_filter_classes</b>	Filter classes that the cells scheduler should use. By default, uses <code>"nova.cells.filters.all_filters"</code> to map to all cells filters included with Compute.
<b>scheduler_weight_classes</b>	Weight classes that the scheduler for cells uses. By default, uses <code>nova.cells.weights.all_weighters</code> to map to all cells weight algorithms included with Compute.
<b>ram_weight_multiplier</b>	Multiplier used to weight RAM. Negative numbers indicate that Compute should stack VMs on one host instead of spreading out new VMs to more hosts in the cell. The default value is 10.0.

## Configure the API (top-level) cell

The cell type must be changed in the API cell so that requests can be proxied through `nova-cells` down to the correct cell properly. Edit the `nova.conf` file in the API cell, and specify `api` in the `cell_type` key:

```
[DEFAULT]
compute_api_class=nova.compute.cells_api.ComputeCellsAPI
...

[cells]
cell_type= api
```

## Configure the child cells

Edit the `nova.conf` file in the child cells, and specify `compute` in the `cell_type` key:

```
[DEFAULT]
# Disable quota checking in child cells. Let API cell do it exclusively.
quota_driver=nova.quota.NoopQuotaDriver

[cells]
cell_type = compute
```

## Configure the database in each cell

Before bringing the services online, the database in each cell needs to be configured with information about related cells. In particular, the API cell needs to know about its immediate children, and the child cells must know about their immediate agents. The information needed is the RabbitMQ server credentials for the particular cell.

Use the **nova-manage cell create** command to add this information to the database in each cell:

```
# nova-manage cell create -h
Options:
  -h, --help                show this help message and exit
  --name=<name>              Name for the new cell
  --cell_type=<parent|child>
                             Whether the cell is a parent or child
  --username=<username>     Username for the message broker in this cell
  --password=<password>     Password for the message broker in this cell
  --hostname=<hostname>     Address of the message broker in this cell
  --port=<number>           Port number of the message broker in this cell
  --virtual_host=<virtual_host>
                             The virtual host of the message broker in this cell
  --woffset=<float>         (weight offset) It might be used by some cell
                             scheduling code in the future
  --wscale=<float>         (weight scale) It might be used by some cell
                             scheduling code in the future
```

As an example, assume an API cell named `api` and a child cell named `cell1`.

Within the `api` cell, specify the following RabbitMQ server information:

```
rabbit_host=10.0.0.10
rabbit_port=5672
rabbit_username=api_user
rabbit_password=api_passwd
rabbit_virtual_host=api_vhost
```

Within the `cell1` child cell, specify the following RabbitMQ server information:

```
rabbit_host=10.0.1.10
rabbit_port=5673
rabbit_username=cell1_user
rabbit_password=cell1_passwd
rabbit_virtual_host=cell1_vhost
```

You can run this in the API cell as root:

```
# nova-manage cell create --name cell1 --cell_type child \
  --username cell1_user --password cell1_passwd --hostname 10.0.1.10 \
  --port 5673 --virtual_host cell1_vhost --woffset 1.0 --wscale 1.0
```

Repeat the previous steps for all child cells.

In the child cell, run the following, as root:



be the first cell to be scheduled for launching an instance.

Additionally, the following options are available for the cell scheduler:

**scheduler\_retries** Specifies how many times the scheduler tries to launch a new instance when no cells are available (default=10).

**scheduler\_retry\_delay** Specifies the delay (in seconds) between retries (default=2).

As an admin user, you can also add a filter that directs builds to a particular cell. The `policy.json` file must have a line with `"cells_scheduler_filter:TargetCellFilter" : "is_admin:True"` to let an admin user specify a scheduler hint to direct a build to a particular cell.

## Optional cell configuration

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 the `[cells]cells_config` 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 `transport_url` field, instead of `username`, `password`, and so on. The `transport_url` has the following form:

```
rabbit://USERNAME:PASSWORD@HOSTNAME:PORT/VIRTUAL_HOST
```

The scheme can be either `qpuid` or `rabbit`, as shown previously. The following sample shows this optional configuration:

```
{
  "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
  }
}
```

## Conductor

The `nova-conductor` service enables OpenStack to function without compute nodes accessing the database. Conceptually, it implements a new layer on top of `nova-compute`. It should not be deployed on compute nodes, or else the security benefits of removing database access from `nova-compute` are negated. Just like other nova services such as `nova-api` or `nova-scheduler`, it can be scaled horizontally. You can run multiple instances of `nova-conductor` on different machines as needed for scaling purposes.

The methods exposed by `nova-conductor` are relatively simple methods used by `nova-compute` to offload its database operations. Places where `nova-compute` previously performed database access are now talking to `nova-conductor`. However, we have plans in the medium to long term to move more and more of what is currently in `nova-compute` up to the `nova-conductor` layer. The Compute service will start to look like a less intelligent slave service to `nova-conductor`. The conductor service will implement long running complex operations, ensuring forward progress and graceful error handling. This will be especially beneficial for operations that cross multiple compute nodes, such as migrations or resizes.

To customize the Conductor, use the configuration option settings documented in [Table 3.22, "Description of conductor configuration options" \[292\]](#).

## Example `nova.conf` configuration files

The following sections describe the configuration options in the `nova.conf` file. You must copy the `nova.conf` file to each compute node. The sample `nova.conf` files show examples of specific configurations.

### Small, private cloud

This example `nova.conf` file configures a small private cloud with cloud controller services, database server, and messaging server on the same server. In this case, `CONTROLLER_IP` represents the IP address of a central server, `BRIDGE_INTERFACE` represents the bridge such as `br100`, the `NETWORK_INTERFACE` represents an interface to your VLAN setup, and passwords are represented as `DB_PASSWORD_COMPUTE` for your Compute (nova) database password, and `RABBIT_PASSWORD` represents the password to your message queue installation.

```
[DEFAULT]

# LOGS/STATE
verbose=True
logdir=/var/log/nova
state_path=/var/lib/nova
lock_path=/var/lock/nova
rootwrap_config=/etc/nova/rootwrap.conf

# SCHEDULER
compute_scheduler_driver=nova.scheduler.filter_scheduler.FilterScheduler

# VOLUMES
# configured in cinder.conf

# COMPUTE
```

```

compute_driver=libvirt.LibvirtDriver
instance_name_template=instance-%08x
api_paste_config=/etc/nova/api-paste.ini

# COMPUTE/APIS: if you have separate configs for separate services
# this flag is required for both nova-api and nova-compute
allow_resize_to_same_host=True

# APIS
osapi_compute_extension=nova.api.openstack.compute.contrib.standard_extensions
ec2_dmz_host=192.168.206.130
s3_host=192.168.206.130

# RABBITMQ
rabbit_host=192.168.206.130

# GLANCE
image_service=nova.image.glance.GlanceImageService

# NETWORK
network_manager=nova.network.manager.FlatDHCPManager
force_dhcp_release=True
dhcpbridge_flagfile=/etc/nova/nova.conf
firewall_driver=nova.virt.libvirt.firewall.IptablesFirewallDriver
# Change my_ip to match each host
my_ip=192.168.206.130
public_interface=eth0
vlan_interface=eth0
flat_network_bridge=br100
flat_interface=eth0

# NOVNC CONSOLE
novncproxy_base_url=http://192.168.206.130:6080/vnc_auto.html
# Change vncserver_proxycient_address and vncserver_listen to match each
compute host
vncserver_proxycient_address=192.168.206.130
vncserver_listen=192.168.206.130

# AUTHENTICATION
auth_strategy=keystone
[keystone_auth token]
auth_host = 127.0.0.1
auth_port = 35357
auth_protocol = http
admin_tenant_name = service
admin_user = nova
admin_password = nova
signing_dirname = /tmp/keystone-signing-nova

# GLANCE
[glance]
api_servers=192.168.206.130:9292

# DATABASE
[database]
connection=mysql://nova:yourpassword@192.168.206.130/nova

# LIBVIRT
[libvirt]
virt_type=qemu

```









## Compute log files

The corresponding log file of each Compute service is stored in the `/var/log/nova/` directory of the host on which each service runs.

**Table 3.11. Log files used by Compute services**

Log file	Service name (CentOS/Fedora/openSUSE/Red Hat Enterprise Linux/SUSE Linux Enterprise)	Service name (Ubuntu/Debian)
<code>api.log</code>	<code>openstack-nova-api</code>	<code>nova-api</code>
<code>cert.log</code> <sup>a</sup>	<code>openstack-nova-cert</code>	<code>nova-cert</code>
<code>compute.log</code>	<code>openstack-nova-compute</code>	<code>nova-compute</code>
<code>conductor.log</code>	<code>openstack-nova-conductor</code>	<code>nova-conductor</code>
<code>consoleauth.log</code>	<code>openstack-nova-consoleauth</code>	<code>nova-consoleauth</code>
<code>network.log</code> <sup>b</sup>	<code>openstack-nova-network</code>	<code>nova-network</code>
<code>nova-manage.log</code>	<code>nova-manage</code>	<code>nova-manage</code>
<code>scheduler.log</code>	<code>openstack-nova-scheduler</code>	<code>nova-scheduler</code>

<sup>a</sup>The X509 certificate service (`openstack-nova-cert/nova-cert`) is only required by the EC2 API to the Compute service.

<sup>b</sup>The nova network service (`openstack-nova-network/nova-network`) only runs in deployments that are not configured to use the Networking service (neutron).

## Compute sample configuration files

### nova.conf - configuration options

For a complete list of all available configuration options for each OpenStack Compute service, run `bin/nova-<servicename> --help`.

**Table 3.12. Description of API configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>api_paste_config = api-paste.ini</code>	(StrOpt) File name for the paste.deploy config for nova-api
<code>api_rate_limit = False</code>	(BoolOpt) Whether to use per-user rate limiting for the api. This option is only used by v2 api. Rate limiting is removed from v3 api.
<code>client_socket_timeout = 900</code>	(IntOpt) 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.
<code>enable_new_services = True</code>	(BoolOpt) Services to be added to the available pool on create
<code>enabled_apis = ec2, osapi_compute, metadata</code>	(ListOpt) A list of APIs to enable by default
<code>enabled_ssl_apis =</code>	(ListOpt) A list of APIs with enabled SSL
<code>instance_name_template = instance-%08x</code>	(StrOpt) Template string to be used to generate instance names
<code>max_header_line = 16384</code>	(IntOpt) 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).

Configuration option = Default value	Description
multi_instance_display_name_template = <code>%(name)s-%(count)d</code>	(StrOpt) 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. To restore legacy behavior of every instance having the same name, set this option to "%(name)s". Valid keys for the template are: name, uuid, count.
non_inheritable_image_properties = <code>cache_in_nova, bittorrent</code>	(ListOpt) These are image properties which a snapshot should not inherit from an instance
null_kernel = <code>nokernel</code>	(StrOpt) Kernel image that indicates not to use a kernel, but to use a raw disk image instead
osapi_compute_ext_list =	(ListOpt) Specify list of extensions to load when using osapi_compute_extension option with nova.api.openstack.compute.contrib.select_extensions
osapi_compute_extension = <code>['nova.api.openstack.compute.contrib.standard_extensions']</code>	(MultiStrOpt) osapi compute extension to load
osapi_compute_link_prefix = <code>None</code>	(StrOpt) Base URL that will be presented to users in links to the OpenStack Compute API
osapi_compute_listen = <code>0.0.0.0</code>	(StrOpt) The IP address on which the OpenStack API will listen.
osapi_compute_listen_port = <code>8774</code>	(IntOpt) The port on which the OpenStack API will listen.
osapi_compute_workers = <code>None</code>	(IntOpt) Number of workers for OpenStack API service. The default will be the number of CPUs available.
osapi_hide_server_address_states = <code>building</code>	(ListOpt) List of instance states that should hide network info
servicegroup_driver = <code>db</code>	(StrOpt) The driver for servicegroup service (valid options are: db, zk, mc)
snapshot_name_template = <code>snapshot-%s</code>	(StrOpt) Template string to be used to generate snapshot names
tcp_keepidle = <code>600</code>	(IntOpt) Sets the value of TCP_KEEPIDLE in seconds for each server socket. Not supported on OS X.
use_forwarded_for = <code>False</code>	(BoolOpt) Treat X-Forwarded-For as the canonical remote address. Only enable this if you have a sanitizing proxy.
wsgi_default_pool_size = <code>1000</code>	(IntOpt) Size of the pool of greenthreads used by wsgi
wsgi_keep_alive = <code>True</code>	(BoolOpt) If False, closes the client socket connection explicitly.
wsgi_log_format = <code>%(client_ip)s "%(request_line)s" status: %(status_code)s len: %(body_length)s time: %(wall_seconds).7f</code>	(StrOpt) 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.

**Table 3.13. Description of API v3 configuration options**

Configuration option = Default value	Description
[osapi_v3]	
enabled = <code>False</code>	(BoolOpt) Whether the V3 API is enabled or not
extensions_blacklist =	(ListOpt) A list of v3 API extensions to never load. Specify the extension aliases here.
extensions_whitelist =	(ListOpt) If the list is not empty then a v3 API extension will only be loaded if it exists in this list. Specify the extension aliases here.





Configuration option = Default value	Description
	for a configurable duration (in seconds). Set to -1 to disable caching completely.

**Table 3.16. Description of availability zones configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>default_availability_zone = nova</code>	(StrOpt) Default compute node availability_zone
<code>default_schedule_zone = None</code>	(StrOpt) Availability zone to use when user doesn't specify one
<code>internal_service_availability_zone = internal</code>	(StrOpt) The availability_zone to show internal services under

**Table 3.17. Description of Barbican configuration options**

Configuration option = Default value	Description
[barbican]	
<code>cafile = None</code>	(StrOpt) PEM encoded Certificate Authority to use when verifying HTTPs connections.
<code>catalog_info = key-manager:barbican:public</code>	(StrOpt) Info to match when looking for barbican in the service catalog. Format is: separated values of the form: <service_type>:<service_name>:<endpoint_type>
<code>certfile = None</code>	(StrOpt) PEM encoded client certificate cert file
<code>endpoint_template = None</code>	(StrOpt) Override service catalog lookup with template for barbican endpoint e.g. http://localhost:9311/v1/%(project_id)s
<code>insecure = False</code>	(BoolOpt) Verify HTTPS connections.
<code>keyfile = None</code>	(StrOpt) PEM encoded client certificate key file
<code>os_region_name = None</code>	(StrOpt) Region name of this node
<code>timeout = None</code>	(IntOpt) Timeout value for http requests

**Table 3.18. Description of CA and SSL configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>ca_file = cacert.pem</code>	(StrOpt) Filename of root CA
<code>ca_path = \$state_path/CA</code>	(StrOpt) Where we keep our root CA
<code>cert = self.pem</code>	(StrOpt) SSL certificate file
<code>cert_manager = nova.cert.manager.CertManager</code>	(StrOpt) Full class name for the Manager for cert
<code>cert_topic = cert</code>	(StrOpt) The topic cert nodes listen on
<code>crl_file = crl.pem</code>	(StrOpt) Filename of root Certificate Revocation List
<code>key_file = private/akey.pem</code>	(StrOpt) Filename of private key
<code>keys_path = \$state_path/keys</code>	(StrOpt) Where we keep our keys
<code>project_cert_subject = /C=US/ST=California/O=OpenStack/OU=NovaDev/CN=project-ca-%.16s-%s</code>	(StrOpt) Subject for certificate for projects, %s for project, timestamp
<code>ssl_ca_file = None</code>	(StrOpt) CA certificate file to use to verify connecting clients
<code>ssl_cert_file = None</code>	(StrOpt) SSL certificate of API server
<code>ssl_key_file = None</code>	(StrOpt) SSL private key of API server
<code>use_project_ca = False</code>	(BoolOpt) Should we use a CA for each project?

























Configuration option = Default value	Description
<code>disk_prefix = None</code>	(StrOpt) Override the default disk prefix for the devices attached to a server, which is dependent on <code>virt_type</code> . (valid options are: <code>sd</code> , <code>xvd</code> , <code>uvd</code> , <code>vd</code> )
<code>gid_maps =</code>	(ListOpt) List of guid targets and ranges. Syntax is <code>guest-gid:host-gid:count</code> Maximum of 5 allowed.
<code>hw_disk_discard = None</code>	(StrOpt) Discard option for nova managed disks (valid options are: <code>ignore</code> , <code>unmap</code> ). Need Libvirt(1.0.6) Qemu1.5 (raw format) Qemu1.6(qcow2 format)
<code>hw_machine_type = None</code>	(ListOpt) 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 " <code>virsh capabilities</code> " command. The format of the value for this config option is <code>host-arch=machine-type</code> . For example: <code>x86_64=machine1,armv7l=machine2</code>
<code>image_info_filename_pattern = \$instances_path/\$image_cache_subdirectory_name/%(image)s.info</code>	(StrOpt) Allows image information files to be stored in non-standard locations
<code>images_rbd_ceph_conf =</code>	(StrOpt) Path to the ceph configuration file to use
<code>images_rbd_pool = rbd</code>	(StrOpt) The RADOS pool in which rbd volumes are stored
<code>images_type = default</code>	(StrOpt) VM Images format. Acceptable values are: <code>raw</code> , <code>qcow2</code> , <code>lvm</code> , <code>rbd</code> , <code>default</code> . If default is specified, then <code>use_cow_images</code> flag is used instead of this one.
<code>images_volume_group = None</code>	(StrOpt) LVM Volume Group that is used for VM images, when you specify <code>images_type=lvm</code> .
<code>inject_key = False</code>	(BoolOpt) Inject the ssh public key at boot time
<code>inject_partition = -2</code>	(IntOpt) The partition to inject to : <code>-2 =&gt; disable</code> , <code>-1 =&gt; inspect</code> (libguestfs only), <code>0 =&gt; not partitioned</code> , <code>&gt;0 =&gt; partition number</code>
<code>inject_password = False</code>	(BoolOpt) Inject the admin password at boot time, without an agent.
<code>iscsi_iface = None</code>	(StrOpt) The iSCSI transport iface to use to connect to target in case offload support is desired. Supported transports are <code>be2iscsi</code> , <code>bnx2i</code> , <code>cxgb3i</code> , <code>cxgb4i</code> , <code>qla4xxx</code> and <code>ocs</code> . Default format is <code>transport_name.hwaddress</code> and can be generated manually or via <code>iscsiadm -m iface</code>
<code>iscsi_use_multipath = False</code>	(BoolOpt) Use multipath connection of the iSCSI volume
<code>iser_use_multipath = False</code>	(BoolOpt) Use multipath connection of the iSER volume
<code>mem_stats_period_seconds = 10</code>	(IntOpt) A number of seconds to memory usage statistics period. Zero or negative value mean to disable memory usage statistics.
<code>remove_unused_kernels = False</code>	(BoolOpt) Should unused kernel images be removed? This is only safe to enable if all compute nodes have been updated to support this option. This will be enabled by default in future.
<code>remove_unused_resized_minimum_age_seconds = 3600</code>	(IntOpt) Unused resized base images younger than this will not be removed
<code>rescue_image_id = None</code>	(StrOpt) Rescue ami image. This will not be used if an image id is provided by the user.
<code>rescue_kernel_id = None</code>	(StrOpt) Rescue aki image
<code>rescue_ramdisk_id = None</code>	(StrOpt) Rescue ari image
<code>rng_dev_path = None</code>	(StrOpt) A path to a device that will be used as source of entropy on the host. Permitted options are: <code>/dev/random</code> or <code>/dev/hwrng</code>

Configuration option = Default value	Description
<code>snapshot_compression = False</code>	(BoolOpt) Compress snapshot images when possible. This currently applies exclusively to qcow2 images
<code>snapshot_image_format = None</code>	(StrOpt) Snapshot image format (valid options are : raw, qcow2, vmdk, vdi). Defaults to same as source image
<code>snapshots_directory = \$instances_path/snapshots</code>	(StrOpt) Location where libvirt driver will store snapshots before uploading them to image service
<code>sparse_logical_volumes = False</code>	(BoolOpt) Create sparse logical volumes (with virtulsize) if this flag is set to True.
<code>sysinfo_serial = auto</code>	(StrOpt) The data source used to the populate the host "serial" UUID exposed to guest in the virtual BIOS. Permitted options are "hardware", "os", "none" or "auto" (default).
<code>uid_maps =</code>	(ListOpt) List of uid targets and ranges.Syntax is guest-uid:host-uid:countMaximum of 5 allowed.
<code>use_usb_tablet = True</code>	(BoolOpt) Sync virtual and real mouse cursors in Windows VMs
<code>use_virtio_for_bridges = True</code>	(BoolOpt) Use virtio for bridge interfaces with KVM/QEMU
<code>virt_type = kvm</code>	(StrOpt) Libvirt domain type (valid options are: kvm, lxc, qemu, uml, xen and parallels)
<code>volume_clear = zero</code>	(StrOpt) Method used to wipe old volumes (valid options are: none, zero, shred)
<code>volume_clear_size = 0</code>	(IntOpt) Size in MiB to wipe at start of old volumes. 0 => all
<code>wait_soft_reboot_seconds = 120</code>	(IntOpt) 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.

**Table 3.38. Description of live migration configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>live_migration_retry_count = 30</code>	(IntOpt) Number of 1 second retries needed in live_migration
[libvirt]	
<code>live_migration_bandwidth = 0</code>	(IntOpt) Maximum bandwidth to be used during migration, in Mbps
<code>live_migration_flag = VIR_MIGRATE_UNDEFINE_SOURCE, VIR_MIGRATE_PEER2PEER, VIR_MIGRATE_LIVE, VIR_MIGRATE_TUNNELLED</code>	(StrOpt) Migration flags to be set for live migration
<code>live_migration_uri = qemu+tcp://%s/system</code>	(StrOpt) Migration target URI (any included "%s" is replaced with the migration target hostname)

**Table 3.39. Description of logging configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>debug = False</code>	(BoolOpt) Print debugging output (set logging level to DEBUG instead of default WARNING level).
<code>default_log_levels = amqp=WARN, amqplib=WARN, boto=WARN, qpid=WARN, sqlalchemy=WARN, suds=INFO, oslo.messaging=INFO, iso8601=WARN,</code>	(ListOpt) List of logger=LEVEL pairs.

Configuration option = Default value	Description
<code>requests.packages.urllib3.connectionpool=WARN, urllib3.connectionpool=WARN, websocket=WARN, requests.packages.urllib3.util.retry=WARN, urllib3.util.retry=WARN, keystonemiddleware=WARN, routes.middleware=WARN, stevedore=WARN</code>	
<code>fatal_deprecations = False</code>	(BoolOpt) Enables or disables fatal status of deprecations.
<code>fatal_exception_format_errors = False</code>	(BoolOpt) Make exception message format errors fatal
<code>instance_format = "[instance: %(uuid)s] "</code>	(StrOpt) The format for an instance that is passed with the log message.
<code>instance_uuid_format = "[instance: %(uuid)s] "</code>	(StrOpt) The format for an instance UUID that is passed with the log message.
<code>log_config_append = None</code>	(StrOpt) 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.
<code>log_date_format = %Y-%m-%d %H:%M:%S</code>	(StrOpt) Format string for <code>%(asctime)s</code> in log records. Default: <code>%(default)s</code> .
<code>log_dir = None</code>	(StrOpt) (Optional) The base directory used for relative – log-file paths.
<code>log_file = None</code>	(StrOpt) (Optional) Name of log file to output to. If no default is set, logging will go to stdout.
<code>log_format = None</code>	(StrOpt) DEPRECATED. A logging.Formatter log message format string which may use any of the available logging.LogRecord attributes. This option is deprecated. Please use <code>logging_context_format_string</code> and <code>logging_default_format_string</code> instead.
<code>log_config_append = None</code>	(StrOpt) 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.
<code>log_date_format = %Y-%m-%d %H:%M:%S</code>	(StrOpt) Format string for <code>%(asctime)s</code> in log records. Default: <code>%(default)s</code> .
<code>log_dir = None</code>	(StrOpt) (Optional) The base directory used for relative – log-file paths.
<code>log_file = None</code>	(StrOpt) (Optional) Name of log file to output to. If no default is set, logging will go to stdout.
<code>log_format = None</code>	(StrOpt) DEPRECATED. A logging.Formatter log message format string which may use any of the available logging.LogRecord attributes. This option is deprecated. Please use <code>logging_context_format_string</code> and <code>logging_default_format_string</code> instead.
<code>logging_context_format_string = %(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [%(request_id)s %(user_identity)s] %(instance)s%(message)s</code>	(StrOpt) Format string to use for log messages with context.
<code>logging_debug_format_suffix = %(funcName)s %(pathname)s:%(lineno)d</code>	(StrOpt) Data to append to log format when level is DEBUG.
<code>logging_default_format_string = %(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [-] %(instance)s%(message)s</code>	(StrOpt) Format string to use for log messages without context.
<code>logging_exception_prefix = %(asctime)s.%(msecs)03d %(process)d TRACE %(name)s %(instance)s</code>	(StrOpt) Prefix each line of exception output with this format.
<code>publish_errors = False</code>	(BoolOpt) Enables or disables publication of error events.

Configuration option = Default value	Description
<code>syslog_log_facility = LOG_USER</code>	(StrOpt) Syslog facility to receive log lines.
<code>syslog_log_facility = LOG_USER</code>	(StrOpt) Syslog facility to receive log lines.
<code>use_syslog = False</code>	(BoolOpt) Use syslog for logging. Existing syslog format is DEPRECATED during I, and will change in J to honor RFC5424.
<code>use_syslog_rfc_format = False</code>	(BoolOpt) (Optional) Enables or disables syslog rfc5424 format for logging. If enabled, prefixes the MSG part of the syslog message with APP-NAME (RFC5424). The format without the APP-NAME is deprecated in I, and will be removed in J.
<code>use_stderr = True</code>	(BoolOpt) Log output to standard error.
<code>use_syslog = False</code>	(BoolOpt) Use syslog for logging. Existing syslog format is DEPRECATED during I, and will change in J to honor RFC5424.
<code>use_syslog_rfc_format = False</code>	(BoolOpt) (Optional) Enables or disables syslog rfc5424 format for logging. If enabled, prefixes the MSG part of the syslog message with APP-NAME (RFC5424). The format without the APP-NAME is deprecated in I, and will be removed in J.
<code>verbose = False</code>	(BoolOpt) Print more verbose output (set logging level to INFO instead of default WARNING level).

**Table 3.40. Description of metadata configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>metadata_cache_expiration = 15</code>	(IntOpt) Time in seconds to cache metadata; 0 to disable metadata caching entirely (not recommended). Increasing this 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.
<code>metadata_host = \$my_ip</code>	(StrOpt) The IP address for the metadata API server
<code>metadata_listen = 0.0.0.0</code>	(StrOpt) The IP address on which the metadata API will listen.
<code>metadata_listen_port = 8775</code>	(IntOpt) The port on which the metadata API will listen.
<code>metadata_manager = nova.api.manager.MetadataManager</code>	(StrOpt) OpenStack metadata service manager
<code>metadata_port = 8775</code>	(IntOpt) The port for the metadata API port
<code>metadata_workers = None</code>	(IntOpt) Number of workers for metadata service. The default will be the number of CPUs available.
<code>vendordata_driver = nova.api.metadata.vendordata_json.JsonFileVendorData</code>	(StrOpt) Driver to use for vendor data
<code>vendordata_jsonfile_path = None</code>	(StrOpt) File to load JSON formatted vendor data from

**Table 3.41. Description of network configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>allow_same_net_traffic = True</code>	(BoolOpt) Whether to allow network traffic from same network
<code>auto_assign_floating_ip = False</code>	(BoolOpt) Autoassigning floating IP to VM
<code>cnt_vpn_clients = 0</code>	(IntOpt) Number of addresses reserved for vpn clients

Configuration option = Default value	Description
<code>create_unique_mac_address_attempts = 5</code>	(IntOpt) Number of attempts to create unique mac address
<code>default_access_ip_network_name = None</code>	(StrOpt) Name of network to use to set access IPs for instances
<code>default_floating_pool = nova</code>	(StrOpt) Default pool for floating IPs
<code>defer iptables_apply = False</code>	(BoolOpt) Whether to batch up the application of IPTables rules during a host restart and apply all at the end of the init phase
<code>dhcp_domain = nova-local</code>	(StrOpt) Domain to use for building the hostnames
<code>dhcp_lease_time = 86400</code>	(IntOpt) Lifetime of a DHCP lease in seconds
<code>dhcpbridge = \$bindir/nova-dhcpbridge</code>	(StrOpt) Location of nova-dhcpbridge
<code>dhcpbridge_flagfile = [ '/etc/nova/nova-dhcpbridge.conf' ]</code>	(MultiStrOpt) Location of flagfiles for dhcpbridge
<code>dns_server = [ ]</code>	(MultiStrOpt) If set, uses specific DNS server for dnsmasq. Can be specified multiple times.
<code>dns_update_periodic_interval = -1</code>	(IntOpt) Number of seconds to wait between runs of updates to DNS entries.
<code>dnsmasq_config_file =</code>	(StrOpt) Override the default dnsmasq settings with this file
<code>ebtables_exec_attempts = 3</code>	(IntOpt) Number of times to retry ebtables commands on failure.
<code>ebtables_retry_interval = 1.0</code>	(FloatOpt) Number of seconds to wait between ebtables retries.
<code>firewall_driver = None</code>	(StrOpt) Firewall driver (defaults to hypervisor specific iptables driver)
<code>fixed_ip_disassociate_timeout = 600</code>	(IntOpt) Seconds after which a deallocated IP is disassociated
<code>flat_injected = False</code>	(BoolOpt) Whether to attempt to inject network setup into guest
<code>flat_interface = None</code>	(StrOpt) FlatDhcp will bridge into this interface if set
<code>flat_network_bridge = None</code>	(StrOpt) Bridge for simple network instances
<code>flat_network_dns = 8.8.4.4</code>	(StrOpt) DNS server for simple network
<code>floating_ip_dns_manager = nova.network.noop_dns_driver.NoopDNSDriver</code>	(StrOpt) Full class name for the DNS Manager for floating IPs
<code>force_dhcp_release = True</code>	(BoolOpt) If True, send a dhcp release on instance termination
<code>force_snat_range = [ ]</code>	(MultiStrOpt) Traffic to this range will always be snatted to the fallback ip, even if it would normally be bridged out of the node. Can be specified multiple times.
<code>forward_bridge_interface = [ 'all' ]</code>	(MultiStrOpt) An interface that bridges can forward to. If this is set to all then all traffic will be forwarded. Can be specified multiple times.
<code>gateway = None</code>	(StrOpt) Default IPv4 gateway
<code>injected_network_template = \$pybasedir/nova/virt/interfaces.template</code>	(StrOpt) Template file for injected network
<code>instance_dns_domain =</code>	(StrOpt) Full class name for the DNS Zone for instance IPs
<code>instance_dns_manager = nova.network.noop_dns_driver.NoopDNSDriver</code>	(StrOpt) Full class name for the DNS Manager for instance IPs
<code>iptables_bottom_regex =</code>	(StrOpt) Regular expression to match the iptables rule that should always be on the bottom.
<code>iptables_drop_action = DROP</code>	(StrOpt) The table that iptables to jump to when a packet is to be dropped.





Configuration option = Default value	Description
<code>vlan_interface = vnic0</code>	(StrOpt) Physical ethernet adapter name for vlan networking

**Table 3.42. Description of neutron configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>neutron_default_tenant_id = default</code>	(StrOpt) Default tenant id when creating neutron networks
[neutron]	
<code>admin_auth_url = http://localhost:5000/v2.0</code>	(StrOpt) Authorization URL for connecting to neutron in admin context. DEPRECATED: specify an <code>auth_plugin</code> and appropriate credentials instead.
<code>admin_password = None</code>	(StrOpt) Password for connecting to neutron in admin context DEPRECATED: specify an <code>auth_plugin</code> and appropriate credentials instead.
<code>admin_tenant_id = None</code>	(StrOpt) Tenant id for connecting to neutron in admin context DEPRECATED: specify an <code>auth_plugin</code> and appropriate credentials instead.
<code>admin_tenant_name = None</code>	(StrOpt) Tenant name for connecting to neutron in admin context. This option will be ignored if <code>neutron_admin_tenant_id</code> is set. Note that with Keystone V3 tenant names are only unique within a domain. DEPRECATED: specify an <code>auth_plugin</code> and appropriate credentials instead.
<code>admin_user_id = None</code>	(StrOpt) User id for connecting to neutron in admin context. DEPRECATED: specify an <code>auth_plugin</code> and appropriate credentials instead.
<code>admin_username = None</code>	(StrOpt) Username for connecting to neutron in admin context DEPRECATED: specify an <code>auth_plugin</code> and appropriate credentials instead.
<code>allow_duplicate_networks = False</code>	(BoolOpt) DEPRECATED: Allow an instance to have multiple vNICs attached to the same Neutron network. This option is deprecated in the 2015.1 release and will be removed in the 2015.2 release where the default behavior will be to always allow multiple ports from the same network to be attached to an instance.
<code>auth_plugin = None</code>	(StrOpt) Name of the plugin to load
<code>auth_section = None</code>	(StrOpt) Config Section from which to load plugin specific options
<code>auth_strategy = keystone</code>	(StrOpt) Authorization strategy for connecting to neutron in admin context. DEPRECATED: specify an <code>auth_plugin</code> and appropriate credentials instead. If an <code>auth_plugin</code> is specified strategy will be ignored.
<code>cafile = None</code>	(StrOpt) PEM encoded Certificate Authority to use when verifying HTTPs connections.
<code>certfile = None</code>	(StrOpt) PEM encoded client certificate cert file
<code>extension_sync_interval = 600</code>	(IntOpt) Number of seconds before querying neutron for extensions
<code>insecure = False</code>	(BoolOpt) Verify HTTPS connections.
<code>keyfile = None</code>	(StrOpt) PEM encoded client certificate key file
<code>metadata_proxy_shared_secret =</code>	(StrOpt) Shared secret to validate proxies Neutron metadata requests
<code>ovs_bridge = br-int</code>	(StrOpt) Name of Integration Bridge used by Open vSwitch



Configuration option = Default value	Description
	will update on a new reservation if max_age has passed since the last reservation
<code>max_local_block_devices = 3</code>	(IntOpt) Maximum number of devices that will result in a local image being created on the hypervisor node. Setting this to 0 means nova will allow only boot from volume. A negative number means unlimited.
<code>osapi_compute_unique_server_name_scope =</code>	(StrOpt) When set, compute API will consider duplicate hostnames invalid within the specified scope, regardless of case. Should be empty, "project" or "global".
<code>osapi_max_limit = 1000</code>	(IntOpt) The maximum number of items returned in a single response from a collection resource
<code>password_length = 12</code>	(IntOpt) Length of generated instance admin passwords
<code>policy_default_rule = default</code>	(StrOpt) Default rule. Enforced when a requested rule is not found.
<code>policy_dirs = [ 'policy.d' ]</code>	(MultiStrOpt) 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.
<code>policy_file = policy.json</code>	(StrOpt) The JSON file that defines policies.
<code>reservation_expire = 86400</code>	(IntOpt) Number of seconds until a reservation expires
<code>resize_fs_using_block_device = False</code>	(BoolOpt) 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).
<code>until_refresh = 0</code>	(IntOpt) 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.

**Table 3.47. Description of Quobyte USP volume driver configuration options**

Configuration option = Default value	Description
[libvirt]	
<code>quobyte_client_cfg = None</code>	(StrOpt) Path to a Quobyte Client configuration file.
<code>quobyte_mount_point_base = \$state_path/mnt</code>	(StrOpt) Directory where the Quobyte volume is mounted on the compute node

**Table 3.48. Description of quota configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>bandwidth_poll_interval = 600</code>	(IntOpt) Interval to pull network bandwidth usage info. Not supported on all hypervisors. Set to -1 to disable. Setting this to 0 will run at the default rate.
<code>enable_network_quota = False</code>	(BoolOpt) Enables or disables quota checking for tenant networks
<code>quota_cores = 20</code>	(IntOpt) Number of instance cores allowed per project
<code>quota_driver = nova.quota.DbQuotaDriver</code>	(StrOpt) Default driver to use for quota checks
<code>quota_fixed_ips = -1</code>	(IntOpt) Number of fixed IPs allowed per project (this should be at least the number of instances allowed)
<code>quota_floating_ips = 10</code>	(IntOpt) Number of floating IPs allowed per project

Configuration option = Default value	Description
<code>quota_injected_file_content_bytes = 10240</code>	(IntOpt) Number of bytes allowed per injected file
<code>quota_injected_file_path_length = 255</code>	(IntOpt) Length of injected file path
<code>quota_injected_files = 5</code>	(IntOpt) Number of injected files allowed
<code>quota_instances = 10</code>	(IntOpt) Number of instances allowed per project
<code>quota_key_pairs = 100</code>	(IntOpt) Number of key pairs per user
<code>quota_metadata_items = 128</code>	(IntOpt) Number of metadata items allowed per instance
<code>quota_networks = 3</code>	(IntOpt) Number of private networks allowed per project
<code>quota_ram = 51200</code>	(IntOpt) Megabytes of instance RAM allowed per project
<code>quota_security_group_rules = 20</code>	(IntOpt) Number of security rules per security group
<code>quota_security_groups = 10</code>	(IntOpt) Number of security groups per project
<code>quota_server_group_members = 10</code>	(IntOpt) Number of servers per server group
<code>quota_server_groups = 10</code>	(IntOpt) Number of server groups per project
[cells]	
<code>bandwidth_update_interval = 600</code>	(IntOpt) Seconds between bandwidth updates for cells.

**Table 3.49. Description of RDP configuration options**

Configuration option = Default value	Description
[rdp]	
<code>enabled = False</code>	(BoolOpt) Enable RDP related features
<code>html5_proxy_base_url = http://127.0.0.1:6083/</code>	(StrOpt) Location of RDP html5 console proxy, in the form "http://127.0.0.1:6083/"

**Table 3.50. Description of Redis configuration options**

Configuration option = Default value	Description
[matchmaker_redis]	
<code>host = 127.0.0.1</code>	(StrOpt) Host to locate redis.
<code>password = None</code>	(StrOpt) Password for Redis server (optional).
<code>port = 6379</code>	(IntOpt) Use this port to connect to redis host.
[matchmaker_ring]	
<code>ringfile = /etc/oslo/matchmaker_ring.json</code>	(StrOpt) Matchmaker ring file (JSON).

**Table 3.51. Description of S3 configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>buckets_path = \$state_path/buckets</code>	(StrOpt) Path to S3 buckets
<code>image_decryption_dir = /tmp</code>	(StrOpt) Parent directory for tempdir used for image decryption
<code>s3_access_key = notchecked</code>	(StrOpt) Access key to use for S3 server for images
<code>s3_affix_tenant = False</code>	(BoolOpt) Whether to affix the tenant id to the access key when downloading from S3
<code>s3_host = \$my_ip</code>	(StrOpt) Hostname or IP for OpenStack to use when accessing the S3 api
<code>s3_listen = 0.0.0.0</code>	(StrOpt) IP address for S3 API to listen
<code>s3_listen_port = 3333</code>	(IntOpt) Port for S3 API to listen
<code>s3_port = 3333</code>	(IntOpt) Port used when accessing the S3 api
<code>s3_secret_key = notchecked</code>	(StrOpt) Secret key to use for S3 server for images

Configuration option = Default value	Description
s3_use_ssl = <i>False</i>	(BoolOpt) Whether to use SSL when talking to S3

**Table 3.52. Description of scheduler configuration options**

Configuration option = Default value	Description
[DEFAULT]	
aggregate_image_properties_isolation_namespace = <i>None</i>	(StrOpt) Force the filter to consider only keys matching the given namespace.
aggregate_image_properties_isolation_separator = <i>.</i>	(StrOpt) The separator used between the namespace and keys
baremetal_scheduler_default_filters = <i>RetryFilter, AvailabilityZoneFilter, ComputeFilter, ComputeCapabilitiesFilter, ImagePropertiesFilter, ExactRamFilter, ExactDiskFilter, ExactCoreFilter</i>	(ListOpt) Which filter class names to use for filtering baremetal hosts when not specified in the request.
cpu_allocation_ratio = <i>16.0</i>	(FloatOpt) Virtual CPU to physical CPU allocation ratio which affects all CPU filters. This configuration specifies a global ratio for CoreFilter. For AggregateCoreFilter, it will fall back to this configuration value if no per-aggregate setting found.
disk_allocation_ratio = <i>1.0</i>	(FloatOpt) Virtual disk to physical disk allocation ratio
io_ops_weight_multiplier = <i>-1.0</i>	(FloatOpt) Multiplier used for weighing host io ops. Negative numbers mean a preference to choose light workload compute hosts.
isolated_hosts =	(ListOpt) Host reserved for specific images
isolated_images =	(ListOpt) Images to run on isolated host
max_instances_per_host = <i>50</i>	(IntOpt) Ignore hosts that have too many instances
max_io_ops_per_host = <i>8</i>	(IntOpt) Tells filters to ignore hosts that have this many or more instances currently in build, resize, snapshot, migrate, rescue or unshelve task states
ram_allocation_ratio = <i>1.5</i>	(FloatOpt) Virtual ram to physical ram allocation ratio which affects all ram filters. This configuration specifies a global ratio for RamFilter. For AggregateRamFilter, it will fall back to this configuration value if no per-aggregate setting found.
ram_weight_multiplier = <i>1.0</i>	(FloatOpt) Multiplier used for weighing ram. Negative numbers mean to stack vs spread.
reserved_host_disk_mb = <i>0</i>	(IntOpt) Amount of disk in MB to reserve for the host
reserved_host_memory_mb = <i>512</i>	(IntOpt) Amount of memory in MB to reserve for the host
restrict_isolated_hosts_to_isolated_images = <i>True</i>	(BoolOpt) Whether to force isolated hosts to run only isolated images
scheduler_available_filters = <i>['nova.scheduler.filters.all_filters']</i>	(MultiStrOpt) Filter classes available to the scheduler which may be specified more than once. An entry of "nova.scheduler.filters.all_filters" maps to all filters included with nova.
scheduler_default_filters = <i>RetryFilter, AvailabilityZoneFilter, RamFilter, ComputeFilter, ComputeCapabilitiesFilter, ImagePropertiesFilter, ServerGroupAntiAffinityFilter, ServerGroupAffinityFilter</i>	(ListOpt) Which filter class names to use for filtering hosts when not specified in the request.
scheduler_driver = <i>nova.scheduler.filter_scheduler.FilterScheduler</i>	(StrOpt) Default driver to use for the scheduler
scheduler_driver_task_period = <i>60</i>	(IntOpt) How often (in seconds) to run periodic tasks in the scheduler driver of your choice. Please note this is likely to interact with the value of service_down_time, but

Configuration option = Default value	Description
	exactly how they interact will depend on your choice of scheduler driver.
<code>scheduler_host_manager = nova.scheduler.host_manager.HostManager</code>	(StrOpt) The scheduler host manager class to use
<code>scheduler_host_subset_size = 1</code>	(IntOpt) New instances will be scheduled on a host chosen randomly from a subset of the N best hosts. This property defines the subset size that a host is chosen from. A value of 1 chooses the first host returned by the weighing functions. This value must be at least 1. Any value less than 1 will be ignored, and 1 will be used instead
<code>scheduler_instance_sync_interval = 120</code>	(IntOpt) Waiting time interval (seconds) between sending the scheduler a list of current instance UUIDs to verify that its view of instances is in sync with nova. If the CONF option `scheduler_tracks_instance_changes` is False, changing this option will have no effect.
<code>scheduler_json_config_location =</code>	(StrOpt) Absolute path to scheduler configuration JSON file.
<code>scheduler_manager = nova.scheduler.manager.SchedulerManager</code>	(StrOpt) Full class name for the Manager for scheduler
<code>scheduler_max_attempts = 3</code>	(IntOpt) Maximum number of attempts to schedule an instance
<code>scheduler_topic = scheduler</code>	(StrOpt) The topic scheduler nodes listen on
<code>scheduler_tracks_instance_changes = True</code>	(BoolOpt) Determines if the Scheduler tracks changes to instances to help with its filtering decisions.
<code>scheduler_use_baremetal_filters = False</code>	(BoolOpt) Flag to decide whether to use <code>baremetal_scheduler_default_filters</code> or not.
<code>scheduler_weight_classes = nova.scheduler.weights.all_weighters</code>	(ListOpt) Which weight class names to use for weighing hosts
<b>[cells]</b>	
<code>ram_weight_multiplier = 10.0</code>	(FloatOpt) Multiplier used for weighing ram. Negative numbers mean to stack vs spread.
<code>scheduler_filter_classes = nova.cells.filters.all_filters</code>	(ListOpt) Filter classes the cells scheduler should use. An entry of "nova.cells.filters.all_filters" maps to all cells filters included with nova.
<code>scheduler_retries = 10</code>	(IntOpt) How many retries when no cells are available.
<code>scheduler_retry_delay = 2</code>	(IntOpt) How often to retry in seconds when no cells are available.
<code>scheduler_weight_classes = nova.cells.weights.all_weighters</code>	(ListOpt) Weigher classes the cells scheduler should use. An entry of "nova.cells.weights.all_weighters" maps to all cell weighters included with nova.
<b>[metrics]</b>	
<code>required = True</code>	(BoolOpt) How to treat the unavailable metrics. When a metric is NOT available for a host, if it is set to be True, it would raise an exception, so it is recommended to use the scheduler filter <code>MetricFilter</code> to filter out those hosts. If it is set to be False, the unavailable metric would be treated as a negative factor in weighing process, the returned value would be set by the option <code>weight_of_unavailable</code> .
<code>weight_multiplier = 1.0</code>	(FloatOpt) Multiplier used for weighing metrics.
<code>weight_of_unavailable = -10000.0</code>	(FloatOpt) The final weight value to be returned if required is set to False and any one of the metrics set by <code>weight_setting</code> is unavailable.
<code>weight_setting =</code>	(ListOpt) How the metrics are going to be weighed. This should be in the form of "<name1>=<ratio1>, <name2>=<ratio2>, ...", where <nameX> is one of the metrics to be weighed, and <ratioX> is the corresponding









**Table 3.60. Description of volumes configuration options**

Configuration option = Default value	Description
[DEFAULT]	
block_device_allocate_retries = 60	(IntOpt) Number of times to retry block device allocation on failures
block_device_allocate_retries_interval = 3	(IntOpt) Waiting time interval (seconds) between block device allocation retries on failures
my_block_storage_ip = \$my_ip	(StrOpt) Block storage IP address of this host
volume_api_class = nova.volume.cinder.API	(StrOpt) The full class name of the volume API class to use
volume_usage_poll_interval = 0	(IntOpt) Interval in seconds for gathering volume usages
[cinder]	
cafile = None	(StrOpt) PEM encoded Certificate Authority to use when verifying HTTPs connections.
catalog_info = volumev2:cinderv2:publicURL	(StrOpt) Info to match when looking for cinder in the service catalog. Format is: separated values of the form: <service_type>:<service_name>:<endpoint_type>
certfile = None	(StrOpt) PEM encoded client certificate cert file
cross_az_attach = True	(BoolOpt) Allow attach between instance and volume in different availability zones.
endpoint_template = None	(StrOpt) Override service catalog lookup with template for cinder endpoint e.g. http://localhost:8776/v1/%(project_id)s
http_retries = 3	(IntOpt) Number of cinderclient retries on failed http calls
insecure = False	(BoolOpt) Verify HTTPS connections.
keyfile = None	(StrOpt) PEM encoded client certificate key file
os_region_name = None	(StrOpt) Region name of this node
timeout = None	(IntOpt) Timeout value for http requests
[hyperv]	
force_volumetools_v1 = False	(BoolOpt) Force V1 volume utility class
volume_attach_retry_count = 10	(IntOpt) The number of times to retry to attach a volume
volume_attach_retry_interval = 5	(IntOpt) Interval between volume attachment attempts, in seconds
[libvirt]	
glusterfs_mount_point_base = \$state_path/mnt	(StrOpt) Directory where the glusterfs volume is mounted on the compute node
nfs_mount_options = None	(StrOpt) Mount options passed to the NFS client. See section of the nfs man page for details
nfs_mount_point_base = \$state_path/mnt	(StrOpt) Directory where the NFS volume is mounted on the compute node
num_aoe_discover_tries = 3	(IntOpt) Number of times to rediscover AoE target to find volume
num_iscsi_scan_tries = 5	(IntOpt) Number of times to rescan iSCSI target to find volume
num_iscsi_scan_tries = 5	(IntOpt) Number of times to rescan iSER target to find volume
qemu_allowed_storage_drivers =	(ListOpt) Protocols listed here will be accessed directly from QEMU. Currently supported protocols: [gluster]
rbd_secret_uuid = None	(StrOpt) The libvirt UUID of the secret for the rbd_uservolumes
rbd_user = None	(StrOpt) The RADOS client name for accessing rbd volumes
scalality_sofs_config = None	(StrOpt) Path or URL to Scalality SOFS configuration file







Configuration option = Default value	Description
address = None	(StrOpt) The ZooKeeper addresses for servicegroup service in the format of host1:port,host2:port,host3:port
recv_timeout = 4000	(IntOpt) The recv_timeout parameter for the zk session
sg_prefix = /servicegroups	(StrOpt) The prefix used in ZooKeeper to store ephemeral nodes
sg_retry_interval = 5	(IntOpt) Number of seconds to wait until retrying to join the session

## Additional sample configuration files

Files in this section can be found in `/etc/nova`.

### api-paste.ini

The Compute service stores its API configuration settings in the `api-paste.ini` file.

```

#####
# Metadata #
#####
[composite:metadata]
use = egg:Paste#urlmap
/: meta

[pipeline:meta]
pipeline = ec2faultwrap logrequest metaapp

[app:metaapp]
paste.app_factory = nova.api.metadata.handler:MetadataRequestHandler.factory

#####
# EC2 #
#####

[composite:ec2]
use = egg:Paste#urlmap
/: ec2cloud

[composite:ec2cloud]
use = call:nova.api.auth:pipeline_factory
noauth = ec2faultwrap logrequest ec2noauth cloudrequest validator ec2executor
noauth2 = ec2faultwrap logrequest ec2noauth cloudrequest validator ec2executor
keystone = ec2faultwrap logrequest ec2keystoneauth cloudrequest validator
           ec2executor

[filter:ec2faultwrap]
paste.filter_factory = nova.api.ec2:FaultWrapper.factory

[filter:logrequest]
paste.filter_factory = nova.api.ec2:RequestLogging.factory

[filter:ec2lockout]
paste.filter_factory = nova.api.ec2:Lockout.factory

[filter:ec2keystoneauth]
paste.filter_factory = nova.api.ec2:EC2KeystoneAuth.factory

```

```
[filter:ec2noauth]
paste.filter_factory = nova.api.ec2:NoAuth.factory

[filter:cloudrequest]
controller = nova.api.ec2.cloud.CloudController
paste.filter_factory = nova.api.ec2:Requestify.factory

[filter:authorizer]
paste.filter_factory = nova.api.ec2:Authorizer.factory

[filter:validator]
paste.filter_factory = nova.api.ec2:Validator.factory

[app:ec2executor]
paste.app_factory = nova.api.ec2:Executor.factory

#####
# OpenStack #
#####

[composite:osapi_compute]
use = call:nova.api.openstack.urlmap:urlmap_factory
/: oscomputeversions
/v1.1: openstack_compute_api_v2
/v2: openstack_compute_api_v2
/v2.1: openstack_compute_api_v21
/v3: openstack_compute_api_v3

[composite:openstack_compute_api_v2]
use = call:nova.api.auth:pipeline_factory
noauth = compute_req_id faultwrap sizelimit noauth ratelimit
  osapi_compute_app_v2
noauth2 = compute_req_id faultwrap sizelimit noauth2 ratelimit
  osapi_compute_app_v2
keystone = compute_req_id faultwrap sizelimit authToken keystonecontext
  ratelimit osapi_compute_app_v2
keystone_nolimit = compute_req_id faultwrap sizelimit authToken
  keystonecontext osapi_compute_app_v2

[composite:openstack_compute_api_v21]
use = call:nova.api.auth:pipeline_factory_v21
noauth = compute_req_id faultwrap sizelimit noauth osapi_compute_app_v21
noauth2 = compute_req_id faultwrap sizelimit noauth2 osapi_compute_app_v21
keystone = compute_req_id faultwrap sizelimit authToken keystonecontext
  osapi_compute_app_v21

[composite:openstack_compute_api_v3]
use = call:nova.api.auth:pipeline_factory_v3
noauth = request_id faultwrap sizelimit noauth_v3 osapi_compute_app_v3
noauth2 = request_id faultwrap sizelimit noauth_v3 osapi_compute_app_v3
keystone = request_id faultwrap sizelimit authToken keystonecontext
  osapi_compute_app_v3

[filter:request_id]
paste.filter_factory = oslo.middleware:RequestId.factory

[filter:compute_req_id]
paste.filter_factory = nova.api.compute_req_id:ComputeReqIdMiddleware.factory

[filter:faultwrap]
```

























<b>Option = default value</b>	<b>(Type) Help string</b>
[cinder] cafile = None	(StrOpt) PEM encoded Certificate Authority to use when verifying HTTPS connections.
[cinder] certfile = None	(StrOpt) PEM encoded client certificate cert file
[cinder] insecure = False	(BoolOpt) Verify HTTPS connections.
[cinder] keyfile = None	(StrOpt) PEM encoded client certificate key file
[cinder] timeout = None	(IntOpt) Timeout value for http requests
[database] backend = sqlalchemy	(StrOpt) The back end to use for the database.
[database] connection = None	(StrOpt) The SQLAlchemy connection string to use to connect to the database.
[database] connection_debug = 0	(IntOpt) Verbosity of SQL debugging information: 0=None, 100=Everything.
[database] connection_trace = False	(BoolOpt) Add Python stack traces to SQL as comment strings.
[database] db_inc_retry_interval = True	(BoolOpt) If True, increases the interval between retries of a database operation up to db_max_retry_interval.
[database] db_max_retries = 20	(IntOpt) Maximum retries in case of connection error or deadlock error before error is raised. Set to -1 to specify an infinite retry count.
[database] db_max_retry_interval = 10	(IntOpt) If db_inc_retry_interval is set, the maximum seconds between retries of a database operation.
[database] db_retry_interval = 1	(IntOpt) Seconds between retries of a database transaction.
[database] idle_timeout = 3600	(IntOpt) Timeout before idle SQL connections are reaped.
[database] max_overflow = None	(IntOpt) If set, use this value for max_overflow with SQLAlchemy.
[database] max_pool_size = None	(IntOpt) Maximum number of SQL connections to keep open in a pool.
[database] max_retries = 10	(IntOpt) Maximum number of database connection retries during startup. Set to -1 to specify an infinite retry count.
[database] min_pool_size = 1	(IntOpt) Minimum number of SQL connections to keep open in a pool.
[database] mysql_sql_mode = TRADITIONAL	(StrOpt) 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=
[database] pool_timeout = None	(IntOpt) If set, use this value for pool_timeout with SQLAlchemy.
[database] retry_interval = 10	(IntOpt) Interval between retries of opening a SQL connection.
[database] slave_connection = None	(StrOpt) The SQLAlchemy connection string to use to connect to the slave database.
[database] sqlite_db = oslo.sqlite	(StrOpt) The file name to use with SQLite.
[database] sqlite_synchronous = True	(BoolOpt) If True, SQLite uses synchronous mode.
[database] use_db_reconnect = False	(BoolOpt) Enable the experimental use of database reconnect on connection lost.
[guestfs] debug = False	(BoolOpt) Enable guestfs debug
[libvirt] iscsi_iface = None	(StrOpt) The iSCSI transport iface to use to connect to target in case offload support is desired. Supported transports are be2iscsi, bnx2i, cxgb3i, cxgb4i, qla4xxx and ocs. Default format is transport_name.hwaddress and can be generated manually or via iscsiadm -m iface







**Table 3.67. Deprecated options**

Deprecated option	New Option
[DEFAULT] network_device_mtu	None
[DEFAULT] log-format	None
[DEFAULT] use-syslog	None
[cinder] http_timeout	[cinder] timeout
[DEFAULT] use_syslog	None
[ironic] client_log_level	None
[neutron] admin_username	None
[DEFAULT] osapi_max_request_body_size	[oslo_middleware] max_request_body_size
[neutron] ca_certificates_file	[neutron] cafile
[neutron] auth_strategy	None
[neutron] admin_user_id	None
[neutron] admin_tenant_id	None
[DEFAULT] log_format	None
[cinder] api_insecure	[cinder] insecure
[neutron] admin_tenant_name	None
[neutron] admin_password	None
[DEFAULT] share_dhcp_address	None
[neutron] api_insecure	[neutron] insecure
[cinder] ca_certificates_file	[cinder] cafile
[neutron] admin_auth_url	None
[neutron] url_timeout	[neutron] timeout
[neutron] allow_duplicate_networks	None







```

    },
    'horizon': {
      'handlers': ['console'],
      'propagate': False,
    },
    'novaclient': {
      'handlers': ['console'],
      'propagate': False,
    },
    'keystoneclient': {
      'handlers': ['console'],
      'propagate': False,
    },
    'nose.plugins.manager': {
      'handlers': ['console'],
      'propagate': False,
    }
  }
}

```

The service catalog configuration in the Identity Service determines whether a service appears in the dashboard. For the full listing, see [Horizon Settings and Configuration](#).

2. Restart Apache http server. For Ubuntu/Debian/SUSE:

```
# service apache2 restart
```

or for Fedora/RHEL/CentOS:

```
# service httpd restart
```

Next, restart memcached:

```
# service memcached restart
```

## Configure the dashboard for HTTPS

You can configure the dashboard for a secured HTTPS deployment. While the standard installation uses a non-encrypted HTTP channel, you can enable SSL support for the dashboard.

This example uses the `http://openstack.example.com` domain. Use a domain that fits your current setup.

1. In the `/etc/openstack-dashboard/local_settings.py` file, update the following options:

```

USE_SSL = True
CSRF_COOKIE_SECURE = True
SESSION_COOKIE_SECURE = True
SESSION_COOKIE_HTTPONLY = True

```

To enable HTTPS, the `USE_SSL = True` option is required.

The other options require that HTTPS is enabled; these options defend against cross-site scripting.

2. Edit the `/etc/apache2/conf.d/openstack-dashboard.conf` file as shown in [Example 4.2, "After" \[338\]](#):

**Example 4.1. Before**

```

WSGIScriptAlias / /usr/share/openstack-dashboard/openstack_dashboard/wsgi/django.wsgi
WSGIDaemonProcess horizon user=www-data group=www-data processes=3 threads=10
Alias /static /usr/share/openstack-dashboard/openstack_dashboard/static/
<Directory /usr/share/openstack-dashboard/openstack_dashboard/wsgi>
# For Apache http server 2.2 and earlier:
Order allow,deny
Allow from all

# For Apache http server 2.4 and later:
# Require all granted
</Directory>

```

**Example 4.2. After**

```

<VirtualHost *:80>
ServerName openstack.example.com
<IfModule mod_rewrite.c>
RewriteEngine On
RewriteCond %{HTTPS} off
RewriteRule (.*) https://%{HTTP_HOST}%{REQUEST_URI}
</IfModule>
<IfModule !mod_rewrite.c>
RedirectPermanent / https://openstack.example.com
</IfModule>
</VirtualHost>
<VirtualHost *:443>
ServerName openstack.example.com

SSLEngine On
# Remember to replace certificates and keys with valid paths in your environment
SSLCertificateFile /etc/apache2/SSL/openstack.example.com.crt
SSLCACertificateFile /etc/apache2/SSL/openstack.example.com.crt
SSLCertificateKeyFile /etc/apache2/SSL/openstack.example.com.key
SetEnvIf User-Agent ".MSIE.*" nokeepalive ssl-unclean-shutdown

# HTTP Strict Transport Security (HSTS) enforces that all communications
# with a server go over SSL. This mitigates the threat from attacks such
# as SSL-Strip which replaces links on the wire, stripping away https prefixes
# and potentially allowing an attacker to view confidential information on the
# wire
Header add Strict-Transport-Security "max-age=15768000"

WSGIScriptAlias / /usr/share/openstack-dashboard/openstack_dashboard/wsgi/django.wsgi
WSGIDaemonProcess horizon user=www-data group=www-data processes=3 threads=10
Alias /static /usr/share/openstack-dashboard/openstack_dashboard/static/
<Directory /usr/share/openstack-dashboard/openstack_dashboard/wsgi>
# For Apache http server 2.2 and earlier:
Order allow,deny
Allow from all

# For Apache http server 2.4 and later:
# Require all granted
</Directory>
</VirtualHost>

```

In this configuration, the Apache HTTP server listens on port 443 and redirects all non-secure requests to the HTTPS protocol. The secured section defines the private key, public key, and certificate to use.

- 3. Restart the Apache HTTP server.

For Debian, Ubuntu, or SUSE distributions:

```
# service apache2 restart
```

For Fedora, RHEL, or CentOS distributions:

```
# service httpd restart
```

#### 4. Restart memcached:

```
# service memcached restart
```

If you try to access the dashboard through HTTP, the browser redirects you to the HTTPS page.

## Customize the dashboard

Once you have the dashboard installed you can customize the way it looks and feels to suit your own needs.



### Note

The OpenStack dashboard by default on Ubuntu installs the `openstack-dashboard-ubuntu-theme` package.

If you do not want to use this theme you can remove it and its dependencies using the following command:

```
# apt-get remove --auto-remove openstack-dashboard-ubuntu-theme
```



### Note

This guide focuses on the `local_settings.py` file, stored in `/openstack-dashboard/openstack_dashboard/local/`.

This guide is adapted from [How To Custom Brand The OpenStack "Horizon" Dashboard](#).

The following can easily be customized:

- Site colors
- Logo
- HTML title
- Site branding link
- Help URL

### Procedure 4.1. Logo and site colors

1. Create two logo files, png format, with transparent backgrounds using the following sizes:
  - Login screen: 365 x 50
  - Logged in banner: 216 x 35
2. Upload your new images to the following location: `/usr/share/openstack-dashboard/openstack_dashboard/static/dashboard/img/`
3. Create a CSS style sheet in the following directory: `/usr/share/openstack-dashboard/openstack_dashboard/static/dashboard/css/`

- Change the colors and image file names as appropriate, though the relative directory paths should be the same. The following example file shows you how to customize your CSS file:

```

/*
 * New theme colors for dashboard that override the defaults:
 * dark blue: #355796 / rgb(53, 87, 150)
 * light blue: #BAD3E1 / rgb(186, 211, 225)
 *
 * By Preston Lee <plee@tgen.org>
 */
h1.brand {
background: #355796 repeat-x top left;
border-bottom: 2px solid #BAD3E1;
}
h1.brand a {
background: url(../img/my_cloud_logo_small.png) top left no-repeat;
}
#splash .login {
background: #355796 url(../img/my_cloud_logo_medium.png) no-repeat center 35px;
}
#splash .login .modal-header {
border-top: 1px solid #BAD3E1;
}
.btn-primary {
background-image: none !important;
background-color: #355796 !important;
border: none !important;
box-shadow: none;
}
.btn-primary:hover,
.btn-primary:active {
border: none;
box-shadow: none;
background-color: #BAD3E1 !important;
text-decoration: none;
}

```

- Open the following HTML template in an editor of your choice: `/usr/share/openstack-dashboard/openstack_dashboard/templates/_stylesheets.html`
- Add a line to include your newly created style sheet. For example `custom.css` file:

```

...
<link href='{{ STATIC_URL }}bootstrap/css/bootstrap.min.css' media='screen' rel=
'stylesheet' />
<link href='{{ STATIC_URL }}dashboard/css/{% choose_css %}' media='screen' rel='stylesheet' /
>
<link href='{{ STATIC_URL }}dashboard/css/custom.css' media='screen' rel='stylesheet' />
...

```

- Restart Apache:**

On Ubuntu:

```
# service apache2 restart
```

On Fedora, RHEL, CentOS:

```
# service httpd restart
```

On openSUSE:

```
# service apache2 restart
```

- To view your changes simply reload your dashboard. If necessary go back and modify your CSS file as appropriate.



```
  ],
  "service_or_admin": [
    [
      "rule:admin_required"
    ],
    [
      "rule:service_role"
    ]
  ],
  "owner": [
    [
      "user_id:%(user_id)s"
    ]
  ],
  "admin_or_owner": [
    [
      "rule:admin_required"
    ],
    [
      "rule:owner"
    ]
  ],
  "default": [
    [
      "rule:admin_required"
    ]
  ],
  "identity:get_service": [
    [
      "rule:admin_required"
    ]
  ],
  "identity:list_services": [
    [
      "rule:admin_required"
    ]
  ],
  "identity:create_service": [
    [
      "rule:admin_required"
    ]
  ],
  "identity:update_service": [
    [
      "rule:admin_required"
    ]
  ],
  "identity:delete_service": [
    [
      "rule:admin_required"
    ]
  ],
  "identity:get_endpoint": [
    [
      "rule:admin_required"
    ]
  ],
  "identity:list_endpoints": [
    [
      "rule:admin_required"
    ]
  ]
}
```



























Configuration option = Default value	Description
<code>trove_api_workers = None</code>	(IntOpt) Number of workers for the API service. The default will be the number of CPUs available.
<code>trove_auth_url = http://0.0.0.0:5000/v2.0</code>	(StrOpt) Trove authentication URL.
<code>trove_conductor_workers = None</code>	(IntOpt) Number of workers for the Conductor service. The default will be the number of CPUs available.
<code>trove_security_group_name_prefix = SecGroup</code>	(StrOpt) Prefix to use when creating Security Groups.
<code>trove_security_group_rule_cidr = 0.0.0.0/0</code>	(StrOpt) CIDR to use when creating Security Group Rules.
<code>trove_security_groups_support = True</code>	(BoolOpt) Whether Trove should add Security Groups on create.
<code>users_page_size = 20</code>	(IntOpt) Page size for listing users.

**Table 5.2. Description of authorization token configuration options**

Configuration option = Default value	Description
[keystone_authtoken]	
<code>admin_password = None</code>	(StrOpt) Service user password.
<code>admin_tenant_name = admin</code>	(StrOpt) Service tenant name.
<code>admin_token = None</code>	(StrOpt) This option is deprecated and may be removed in a future release. Single shared secret with the Keystone configuration used for bootstrapping a Keystone installation, or otherwise bypassing the normal authentication process. This option should not be used, use `admin_user` and `admin_password` instead.
<code>admin_user = None</code>	(StrOpt) Service username.
<code>auth_admin_prefix =</code>	(StrOpt) Prefix to prepend at the beginning of the path. Deprecated, use <code>identity_uri</code> .
<code>auth_host = 127.0.0.1</code>	(StrOpt) Host providing the admin Identity API endpoint. Deprecated, use <code>identity_uri</code> .
<code>auth_plugin = None</code>	(StrOpt) Name of the plugin to load
<code>auth_port = 35357</code>	(IntOpt) Port of the admin Identity API endpoint. Deprecated, use <code>identity_uri</code> .
<code>auth_protocol = https</code>	(StrOpt) Protocol of the admin Identity API endpoint (http or https). Deprecated, use <code>identity_uri</code> .
<code>auth_section = None</code>	(StrOpt) Config Section from which to load plugin specific options
<code>auth_uri = None</code>	(StrOpt) Complete public Identity API endpoint.
<code>auth_version = None</code>	(StrOpt) API version of the admin Identity API endpoint.
<code>cache = None</code>	(StrOpt) Env key for the swift cache.
<code>cafile = None</code>	(StrOpt) A PEM encoded Certificate Authority to use when verifying HTTPs connections. Defaults to system CAs.
<code>certfile = None</code>	(StrOpt) Required if identity server requires client certificate
<code>check_revocations_for_cached = False</code>	(BoolOpt) If true, the revocation list will be checked for cached tokens. This requires that PKI tokens are configured on the identity server.
<code>delay_auth_decision = False</code>	(BoolOpt) Do not handle authorization requests within the middleware, but delegate the authorization decision to downstream WSGI components.
<code>enforce_token_bind = permissive</code>	(StrOpt) 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 un-











**Table 5.13. Description of logging configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>debug = False</code>	(BoolOpt) Print debugging output (set logging level to DEBUG instead of default WARNING level).
<code>default_log_levels = amqp=WARN, amqpplib=WARN, boto=WARN, qpid=WARN, sqlalchemy=WARN, suds=INFO, oslo.messaging=INFO, iso8601=WARN, requests.packages.urllib3.connectionpool=WARN, urllib3.connectionpool=WARN, websocket=WARN, keystonemiddleware=WARN, routes.middleware=WARN, stevedore=WARN</code>	(ListOpt) List of logger=LEVEL pairs.
<code>fatal_deprecations = False</code>	(BoolOpt) Enables or disables fatal status of deprecations.
<code>format_options = -m 5</code>	(StrOpt) Options to use when formatting a volume.
<code>instance_format = "[instance: %(uuid)s] "</code>	(StrOpt) The format for an instance that is passed with the log message.
<code>instance_uuid_format = "[instance: %(uuid)s] "</code>	(StrOpt) The format for an instance UUID that is passed with the log message.
<code>log_config_append = None</code>	(StrOpt) 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.
<code>log_date_format = %Y-%m-%d %H:%M:%S</code>	(StrOpt) Format string for %(asctime)s in log records. Default: %(default)s .
<code>log_dir = None</code>	(StrOpt) (Optional) The base directory used for relative – log-file paths.
<code>log_file = None</code>	(StrOpt) (Optional) Name of log file to output to. If no default is set, logging will go to stdout.
<code>log_format = None</code>	(StrOpt) DEPRECATED. A logging.Formatter log message format string which may use any of the available logging.LogRecord attributes. This option is deprecated. Please use logging_context_format_string and logging_default_format_string instead.
<code>logging_context_format_string = %(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [%(request_id)s %(user_identity)s] %(instance)s%(message)s</code>	(StrOpt) Format string to use for log messages with context.
<code>logging_debug_format_suffix = %(funcName)s %(pathname)s:%(lineno)d</code>	(StrOpt) Data to append to log format when level is DEBUG.
<code>logging_default_format_string = %(asctime)s. %(msecs)03d %(process)d %(levelname)s %(name)s [-] %(instance)s%(message)s</code>	(StrOpt) Format string to use for log messages without context.
<code>logging_exception_prefix = %(asctime)s. %(msecs)03d %(process)d TRACE %(name)s %(instance)s</code>	(StrOpt) Prefix each line of exception output with this format.
<code>network_label_regex = ^private\$</code>	(StrOpt) Regular expression to match Trove network labels.
<code>publish_errors = False</code>	(BoolOpt) Enables or disables publication of error events.
<code>syslog_log_facility = LOG_USER</code>	(StrOpt) Syslog facility to receive log lines.
<code>use_stderr = True</code>	(BoolOpt) Log output to standard error.
<code>use_syslog = False</code>	(BoolOpt) Use syslog for logging. Existing syslog format is DEPRECATED during I, and will change in J to honor RFC5424.





















**Table 5.31. Description of Vertica database configuration options**

Configuration option = Default value	Description
[vertica]	
api_strategy = <i>trove.common.strategies.cluster.experimental.vertica.api.VerticaAPIStrategy</i>	(StrOpt) Class that implements datastore-specific API logic.
backup_incremental_strategy = { }	(DictOpt) Incremental Backup Runner based on the default strategy. For strategies that do not implement an incremental, the runner will use the default full backup.
backup_namespace = <i>None</i>	(StrOpt) Namespace to load backup strategies from.
backup_strategy = <i>None</i>	(StrOpt) Default strategy to perform backups.
cluster_member_count = 3	(IntOpt) Number of members in Vertica cluster.
cluster_support = <i>True</i>	(BoolOpt) Enable clusters to be created and managed.
device_path = /dev/vdb	(StrOpt) Device path for volume if volume support is enabled.
guestagent_strategy = <i>trove.common.strategies.cluster.experimental.guestagent.VerticaGuestAgentStrategy</i>	(StrOpt) Class that implements datastore-specific Guest Agent logic.
mount_point = /var/lib/vertica	(StrOpt) Filesystem path for mounting volumes if volume support is enabled.
readahead_size = 2048	(IntOpt) Size(MB) to be set as readahead_size for data volume
replication_strategy = <i>None</i>	(StrOpt) Default strategy for replication.
restore_namespace = <i>None</i>	(StrOpt) Namespace to load restore strategies from.
taskmanager_strategy = <i>trove.common.strategies.cluster.experimental.taskmanager.VerticaTaskManagerStrategy</i>	(StrOpt) Class that implements datastore-specific task manager logic.
tcp_ports = 5433, 5434, 22, 5444, 5450, 4803	(ListOpt) List of TCP ports and/or port ranges to open in the security group (only applicable if <code>trove_security_groups_support</code> is True).
udp_ports = 5433, 4803, 4804, 6453	(ListOpt) List of UDP ports and/or port ranges to open in the security group (only applicable if <code>trove_security_groups_support</code> is True).
volume_support = <i>True</i>	(BoolOpt) Whether to provision a Cinder volume for datadir.

## Configure the RPC messaging system

OpenStack projects use an open standard for messaging middleware known as AMQP. This messaging middleware enables the OpenStack services that run on multiple servers to talk to each other. OpenStack Trove RPC supports three implementations of AMQP: RabbitMQ, Qpid, and ZeroMQ.

### Configure RabbitMQ

Use these options to configure the RabbitMQ messaging system:

**Table 5.32. Description of RabbitMQ configuration options**

Configuration option = Default value	Description
[oslo_messaging_rabbit]	
amqp_auto_delete = <i>False</i>	(BoolOpt) Auto-delete queues in AMQP.
amqp_durable_queues = <i>False</i>	(BoolOpt) Use durable queues in AMQP.

Configuration option = Default value	Description
fake_rabbit = <i>False</i>	(BoolOpt) Deprecated, use rpc_backend=kombu+memory or rpc_backend=fake
heartbeat_rate = 2	(IntOpt) How often times during the heartbeat_timeout_threshold we check the heartbeat.
heartbeat_timeout_threshold = 0	(IntOpt) Number of seconds after which the Rabbit broker is considered down if heartbeat's keep-alive fails (0 disables the heartbeat, >0 enables it. Enabling heartbeats requires kombu>=3.0.7 and amqp>=1.4.0). EXPERIMENTAL
kombu_reconnect_delay = 1.0	(FloatOpt) How long to wait before reconnecting in response to an AMQP consumer cancel notification.
kombu_ssl_ca_certs =	(StrOpt) SSL certification authority file (valid only if SSL enabled).
kombu_ssl_certfile =	(StrOpt) SSL cert file (valid only if SSL enabled).
kombu_ssl_keyfile =	(StrOpt) SSL key file (valid only if SSL enabled).
kombu_ssl_version =	(StrOpt) 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.
rabbit_ha_queues = <i>False</i>	(BoolOpt) Use HA queues in RabbitMQ (x-ha-policy: all). If you change this option, you must wipe the RabbitMQ database.
rabbit_host = localhost	(StrOpt) The RabbitMQ broker address where a single node is used.
rabbit_hosts = \$rabbit_host:\$rabbit_port	(ListOpt) RabbitMQ HA cluster host:port pairs.
rabbit_login_method = AMQPPLAIN	(StrOpt) The RabbitMQ login method.
rabbit_max_retries = 0	(IntOpt) Maximum number of RabbitMQ connection retries. Default is 0 (infinite retry count).
rabbit_password = guest	(StrOpt) The RabbitMQ password.
rabbit_port = 5672	(IntOpt) The RabbitMQ broker port where a single node is used.
rabbit_retry_backoff = 2	(IntOpt) How long to backoff for between retries when connecting to RabbitMQ.
rabbit_retry_interval = 1	(IntOpt) How frequently to retry connecting with RabbitMQ.
rabbit_use_ssl = <i>False</i>	(BoolOpt) Connect over SSL for RabbitMQ.
rabbit_userid = guest	(StrOpt) The RabbitMQ userid.
rabbit_virtual_host = /	(StrOpt) The RabbitMQ virtual host.
rpc_conn_pool_size = 30	(IntOpt) Size of RPC connection pool.

## Configure Qpid

Use these options to configure the Qpid messaging system:

**Table 5.33. Description of Qpid configuration options**

Configuration option = Default value	Description
[oslo_messaging_qpid]	
amqp_auto_delete = <i>False</i>	(BoolOpt) Auto-delete queues in AMQP.
amqp_durable_queues = <i>False</i>	(BoolOpt) Use durable queues in AMQP.
qpid_heartbeat = 60	(IntOpt) Seconds between connection keepalive heartbeats.
qpid_hostname = localhost	(StrOpt) Qpid broker hostname.



Configuration option = Default value	Description
<code>control_exchange = openstack</code>	(StrOpt) The default exchange under which topics are scoped. May be overridden by an exchange name specified in the <code>transport_url</code> option.
<code>notification_driver = []</code>	(MultiStrOpt) Driver or drivers to handle sending notifications.
<code>notification_service_id = { 'vertica': 'a8d805ae-a3b2-c4fd-gb23-b62cee5201ae', 'db2': 'e040cd37-263d-4869-aaa6-c62aa97523b5', 'postgresql': 'ac277e0d-4f21-40aa-b347-1ea31e571720', 'mysql': '2f3ff068-2bfb-4f70-9a9d-a6bb65bc084b', 'couchbase': 'fa62fe68-74d9-4779-a24e-36f19602c415', 'mongodb': 'c8c907af-7375-456f-b929-b637ff9209ee', 'couchdb': 'f0a9ab7b-66f7-4352-93d7-071521d44c7c', 'redis': 'b216ffc5-1947-456c-a4cf-70f94c05f7d0', 'cassandra': '459a230d-4e97-4344-9067-2a54a310b0ed' }</code>	(DictOpt) Unique ID to tag notification events.
<code>notification_topics = notifications</code>	(ListOpt) AMQP topic used for OpenStack notifications.
<code>transport_url = None</code>	(StrOpt) A URL representing the messaging driver to use and its full configuration. If not set, we fall back to the <code>rpc_backend</code> option and driver specific configuration.

**Table 5.36. Description of RPC configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>matchmaker_heartbeat_freq = 300</code>	(IntOpt) Heartbeat frequency.
<code>matchmaker_heartbeat_ttl = 600</code>	(IntOpt) Heartbeat time-to-live.
<code>num_tries = 3</code>	(IntOpt) Number of times to check if a volume exists.
<code>report_interval = 10</code>	(IntOpt) The interval (in seconds) which periodic tasks are run.
<code>rpc_backend = rabbit</code>	(StrOpt) The messaging driver to use, defaults to rabbit. Other drivers include qpid and zmq.
<code>rpc_cast_timeout = 30</code>	(IntOpt) Seconds to wait before a cast expires (TTL). Only supported by <code>impl_zmq</code> .
<code>rpc_response_timeout = 60</code>	(IntOpt) Seconds to wait for a response from a call.
<code>rpc_thread_pool_size = 64</code>	(IntOpt) Size of RPC thread pool.
[oslo_messaging_amqp]	
<code>allow_insecure_clients = False</code>	(BoolOpt) Accept clients using either SSL or plain TCP
<code>broadcast_prefix = broadcast</code>	(StrOpt) address prefix used when broadcasting to all servers
<code>container_name = None</code>	(StrOpt) Name for the AMQP container
<code>group_request_prefix = unicast</code>	(StrOpt) address prefix when sending to any server in group
<code>idle_timeout = 0</code>	(IntOpt) Timeout for inactive connections (in seconds)
<code>server_request_prefix = exclusive</code>	(StrOpt) address prefix used when sending to a specific server
<code>ssl_ca_file =</code>	(StrOpt) CA certificate PEM file to verify server certificate
<code>ssl_cert_file =</code>	(StrOpt) Identifying certificate PEM file to present to clients
<code>ssl_key_file =</code>	(StrOpt) Private key PEM file used to sign <code>cert_file</code> certificate

Configuration option = Default value	Description
<code>ssl_key_password = None</code>	(StrOpt) Password for decrypting <code>ssl_key_file</code> (if encrypted)
<code>trace = False</code>	(BoolOpt) Debug: dump AMQP frames to stdout

## New, updated and deprecated options in Kilo for Database service

Table 5.37. New options

Option = default value	(Type) Help string
[DEFAULT] <code>agent_heartbeat_expiry = 60</code>	(IntOpt) Time (in seconds) after which a guest is considered unreachable
[DEFAULT] <code>cinder_endpoint_type = publicURL</code>	(StrOpt) Service endpoint type to use when searching catalog.
[DEFAULT] <code>guest_info = guest_info.conf</code>	(StrOpt) The guest info filename found in the injected config location. If a full path is specified then it will be used as the path to the guest info file
[DEFAULT] <code>heat_endpoint_type = publicURL</code>	(StrOpt) Service endpoint type to use when searching catalog.
[DEFAULT] <code>injected_config_location = /etc/trove/conf.d</code>	(StrOpt) Path to folder on the Guest where config files will be injected during instance creation.
[DEFAULT] <code>neutron_endpoint_type = publicURL</code>	(StrOpt) Service endpoint type to use when searching catalog.
[DEFAULT] <code>nova_compute_endpoint_type = publicURL</code>	(StrOpt) Service endpoint type to use when searching catalog.
[DEFAULT] <code>swift_endpoint_type = publicURL</code>	(StrOpt) Service endpoint type to use when searching catalog.
[DEFAULT] <code>transport_url = None</code>	(StrOpt) A URL representing the messaging driver to use and its full configuration. If not set, we fall back to the <code>rpc_backend</code> option and driver specific configuration.
[couchdb] <code>backup_incremental_strategy = { }</code>	(DictOpt) Incremental Backup Runner based on the default strategy. For strategies that do not implement an incremental, the runner will use the default full backup.
[couchdb] <code>backup_namespace = None</code>	(StrOpt) Namespace to load backup strategies from.
[couchdb] <code>backup_strategy = None</code>	(StrOpt) Default strategy to perform backups.
[couchdb] <code>device_path = /dev/vdb</code>	(StrOpt) Device path for volume if volume support is enabled.
[couchdb] <code>mount_point = /var/lib/couchdb</code>	(StrOpt) Filesystem path for mounting volumes if volume support is enabled.
[couchdb] <code>replication_strategy = None</code>	(StrOpt) Default strategy for replication.
[couchdb] <code>restore_namespace = None</code>	(StrOpt) Namespace to load restore strategies from.
[couchdb] <code>root_on_create = False</code>	(BoolOpt) Enable the automatic creation of the root user for the service during instance-create. The generated password for the root user is immediately returned in the response of instance-create as the "password" field.
[couchdb] <code>tcp_ports = 5984</code>	(ListOpt) List of TCP ports and/or port ranges to open in the security group (only applicable if <code>trove_security_groups_support</code> is True).
[couchdb] <code>udp_ports =</code>	(ListOpt) List of UDP ports and/or port ranges to open in the security group (only applicable if <code>trove_security_groups_support</code> is True).

Option = default value	(Type) Help string
[couchdb] volume_support = True	(BoolOpt) Whether to provision a Cinder volume for datadir.
[database] connection = sqlite:///trove_test.sqlite	(StrOpt) SQL Connection.
[database] idle_timeout = 3600	(IntOpt) None
[database] query_log = False	(BoolOpt) None
[db2] backup_incremental_strategy = {}	(DictOpt) Incremental Backup Runner based on the default strategy. For strategies that do not implement an incremental, the runner will use the default full backup.
[db2] backup_namespace = None	(StrOpt) Namespace to load backup strategies from.
[db2] backup_strategy = None	(StrOpt) Default strategy to perform backups.
[db2] device_path = /dev/vdb	(StrOpt) Device path for volume if volume support is enabled.
[db2] ignore_users = PUBLIC, DB2INST1	(ListOpt) None
[db2] mount_point = /home/db2inst1/db2inst1	(StrOpt) Filesystem path for mounting volumes if volume support is enabled.
[db2] replication_strategy = None	(StrOpt) Default strategy for replication.
[db2] restore_namespace = None	(StrOpt) Namespace to load restore strategies from.
[db2] root_on_create = False	(BoolOpt) Enable the automatic creation of the root user for the service during instance-create. The generated password for the root user is immediately returned in the response of instance-create as the 'password' field.
[db2] tcp_ports = 50000	(ListOpt) List of TCP ports and/or port ranges to open in the security group (only applicable if trove_security_groups_support is True).
[db2] udp_ports =	(ListOpt) List of UDP ports and/or port ranges to open in the security group (only applicable if trove_security_groups_support is True).
[db2] volume_support = True	(BoolOpt) Whether to provision a Cinder volume for datadir.
[oslo_messaging_amqp] allow_insecure_clients = False	(BoolOpt) Accept clients using either SSL or plain TCP
[oslo_messaging_amqp] broadcast_prefix = broadcast	(StrOpt) address prefix used when broadcasting to all servers
[oslo_messaging_amqp] container_name = None	(StrOpt) Name for the AMQP container
[oslo_messaging_amqp] group_request_prefix = unicast	(StrOpt) address prefix when sending to any server in group
[oslo_messaging_amqp] idle_timeout = 0	(IntOpt) Timeout for inactive connections (in seconds)
[oslo_messaging_amqp] server_request_prefix = exclusive	(StrOpt) address prefix used when sending to a specific server
[oslo_messaging_amqp] ssl_ca_file =	(StrOpt) CA certificate PEM file to verify server certificate
[oslo_messaging_amqp] ssl_cert_file =	(StrOpt) Identifying certificate PEM file to present to clients
[oslo_messaging_amqp] ssl_key_file =	(StrOpt) Private key PEM file used to sign cert_file certificate
[oslo_messaging_amqp] ssl_key_password = None	(StrOpt) Password for decrypting ssl_key_file (if encrypted)
[oslo_messaging_amqp] trace = False	(BoolOpt) Debug: dump AMQP frames to stdout
[oslo_messaging_qpid] amqp_auto_delete = False	(BoolOpt) Auto-delete queues in AMQP.
[oslo_messaging_qpid] amqp_durable_queues = False	(BoolOpt) Use durable queues in AMQP.
[oslo_messaging_qpid] qpid_heartbeat = 60	(IntOpt) Seconds between connection keepalive heartbeats.
[oslo_messaging_qpid] qpid_hostname = localhost	(StrOpt) Qpid broker hostname.



<b>Option = default value</b>	<b>(Type) Help string</b>
[oslo_messaging_rabbit] rabbit_retry_backoff = 2	(IntOpt) How long to backoff for between retries when connecting to RabbitMQ.
[oslo_messaging_rabbit] rabbit_retry_interval = 1	(IntOpt) How frequently to retry connecting with RabbitMQ.
[oslo_messaging_rabbit] rabbit_use_ssl = False	(BoolOpt) Connect over SSL for RabbitMQ.
[oslo_messaging_rabbit] rabbit_userid = guest	(StrOpt) The RabbitMQ userid.
[oslo_messaging_rabbit] rabbit_virtual_host = /	(StrOpt) The RabbitMQ virtual host.
[oslo_messaging_rabbit] rpc_conn_pool_size = 30	(IntOpt) Size of RPC connection pool.
[profiler] enabled = False	(BoolOpt) If False fully disable profiling feature.
[profiler] trace_sqlalchemy = True	(BoolOpt) If False doesn't trace SQL requests.
[upgrade_levels] conductor = icehouse	(StrOpt) Set a version cap for messages sent to conductor services
[upgrade_levels] guestagent = icehouse	(StrOpt) Set a version cap for messages sent to guestagent services
[upgrade_levels] taskmanager = icehouse	(StrOpt) Set a version cap for messages sent to taskmanager services
[vertica] api_strategy = trove.common.strategies.cluster.experimental.vertica.api.VerticaAPIStrategy	(StrOpt) Class that implements datastore-specific API logic.
[vertica] backup_incremental_strategy = { }	(DictOpt) Incremental Backup Runner based on the default strategy. For strategies that do not implement an incremental, the runner will use the default full backup.
[vertica] backup_namespace = None	(StrOpt) Namespace to load backup strategies from.
[vertica] backup_strategy = None	(StrOpt) Default strategy to perform backups.
[vertica] cluster_member_count = 3	(IntOpt) Number of members in Vertica cluster.
[vertica] cluster_support = True	(BoolOpt) Enable clusters to be created and managed.
[vertica] device_path = /dev/vdb	(StrOpt) Device path for volume if volume support is enabled.
[vertica] guestagent_strategy = trove.common.strategies.cluster.experimental.vertica.guestagent.VerticaGuestAgentStrategy	(StrOpt) Class that implements datastore-specific Guest Agent API logic.
[vertica] mount_point = /var/lib/vertica	(StrOpt) Filesystem path for mounting volumes if volume support is enabled.
[vertica] readahead_size = 2048	(IntOpt) Size(MB) to be set as readahead_size for data volume
[vertica] replication_strategy = None	(StrOpt) Default strategy for replication.
[vertica] restore_namespace = None	(StrOpt) Namespace to load restore strategies from.
[vertica] taskmanager_strategy = trove.common.strategies.cluster.experimental.vertica.taskmanager.VerticaTaskManagerStrategy	(StrOpt) Class that implements datastore-specific task manager API logic.
[vertica] tcp_ports = 5433, 5434, 22, 5444, 5450, 4803	(ListOpt) List of TCP ports and/or port ranges to open in the security group (only applicable if trove_security_groups_support is True).
[vertica] udp_ports = 5433, 4803, 4804, 6453	(ListOpt) List of UDP ports and/or port ranges to open in the security group (only applicable if trove_security_groups_support is True).
[vertica] volume_support = True	(BoolOpt) Whether to provision a Cinder volume for datadir.

**Table 5.38. New default values**

<b>Option</b>	<b>Previous default value</b>	<b>New default value</b>
[DEFAULT] default_log_levels	amqp=WARN, amqplib=WARN, boto=WARN, qpid=WARN, sqlalchemy=WARN, suds=INFO,	amqp=WARN, amqplib=WARN, boto=WARN, qpid=WARN, sqlalchemy=WARN, suds=INFO,







Configuration option = Default value	Description
<code>auth_section = None</code>	(StrOpt) Config Section from which to load plugin specific options
<code>auth_uri = None</code>	(StrOpt) Complete public Identity API endpoint.
<code>auth_version = None</code>	(StrOpt) API version of the admin Identity API endpoint.
<code>cache = None</code>	(StrOpt) Env key for the swift cache.
<code>cafile = None</code>	(StrOpt) A PEM encoded Certificate Authority to use when verifying HTTPs connections. Defaults to system CAs.
<code>certfile = None</code>	(StrOpt) Required if identity server requires client certificate
<code>check_revocations_for_cached = False</code>	(BoolOpt) If true, the revocation list will be checked for cached tokens. This requires that PKI tokens are configured on the identity server.
<code>delay_auth_decision = False</code>	(BoolOpt) Do not handle authorization requests within the middleware, but delegate the authorization decision to downstream WSGI components.
<code>enforce_token_bind = permissive</code>	(StrOpt) 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.
<code>hash_algorithms = md5</code>	(ListOpt) 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.
<code>http_connect_timeout = None</code>	(IntOpt) Request timeout value for communicating with Identity API server.
<code>http_request_max_retries = 3</code>	(IntOpt) How many times are we trying to reconnect when communicating with Identity API Server.
<code>identity_uri = None</code>	(StrOpt) Complete admin Identity API endpoint. This should specify the unversioned root endpoint e.g. https://localhost:35357/
<code>include_service_catalog = True</code>	(BoolOpt) (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.
<code>insecure = False</code>	(BoolOpt) Verify HTTPS connections.
<code>keyfile = None</code>	(StrOpt) Required if identity server requires client certificate
<code>memcache_pool_conn_get_timeout = 10</code>	(IntOpt) (Optional) Number of seconds that an operation will wait to get a memcache client connection from the pool.
<code>memcache_pool_dead_retry = 300</code>	(IntOpt) (Optional) Number of seconds memcached server is considered dead before it is tried again.
<code>memcache_pool_maxsize = 10</code>	(IntOpt) (Optional) Maximum total number of open connections to every memcached server.
<code>memcache_pool_socket_timeout = 3</code>	(IntOpt) (Optional) Socket timeout in seconds for communicating with a memcache server.

Configuration option = Default value	Description
<code>memcache_pool_unused_timeout = 60</code>	(IntOpt) (Optional) Number of seconds a connection to memcached is held unused in the pool before it is closed.
<code>memcache_secret_key = None</code>	(StrOpt) (Optional, mandatory if <code>memcache_security_strategy</code> is defined) This string is used for key derivation.
<code>memcache_security_strategy = None</code>	(StrOpt) (Optional) If defined, indicate whether token data should be authenticated or authenticated and encrypted. Acceptable values are MAC or ENCRYPT. 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.
<code>memcache_use_advanced_pool = False</code>	(BoolOpt) (Optional) Use the advanced (eventlet safe) memcache client pool. The advanced pool will only work under python 2.x.
<code>revocation_cache_time = 10</code>	(IntOpt) 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.
<code>signing_dir = None</code>	(StrOpt) Directory used to cache files related to PKI tokens.
<code>token_cache_time = 300</code>	(IntOpt) 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.

**Table 6.3. Description of CA and SSL configuration options**

Configuration option = Default value	Description
[ssl]	
<code>ca_file = None</code>	(StrOpt) CA certificate file to use to verify connecting clients.
<code>cert_file = None</code>	(StrOpt) Certificate file to use when starting the server securely.
<code>key_file = None</code>	(StrOpt) Private key file to use when starting the server securely.

**Table 6.4. Description of clients configuration options**

Configuration option = Default value	Description
[cinder]	
<code>api_insecure = False</code>	(BoolOpt) Allow to perform insecure SSL requests to cinder.
<code>api_version = 2</code>	(IntOpt) Version of the Cinder API to use.
<code>ca_file = None</code>	(StrOpt) Location of ca certificates file to use for cinder client requests.
[heat]	
<code>api_insecure = False</code>	(BoolOpt) Allow to perform insecure SSL requests to heat.
<code>ca_file = None</code>	(StrOpt) Location of ca certificates file to use for heat client requests.
[keystone]	
<code>api_insecure = False</code>	(BoolOpt) Allow to perform insecure SSL requests to keystone.

Configuration option = Default value	Description
ca_file = None	(StrOpt) Location of ca certificates file to use for keystone client requests.
[neutron]	
api_insecure = False	(BoolOpt) Allow to perform insecure SSL requests to neutron.
ca_file = None	(StrOpt) Location of ca certificates file to use for neutron client requests.
[nova]	
api_insecure = False	(BoolOpt) Allow to perform insecure SSL requests to nova.
ca_file = None	(StrOpt) Location of ca certificates file to use for nova client requests.
[swift]	
api_insecure = False	(BoolOpt) Allow to perform insecure SSL requests to swift.
ca_file = None	(StrOpt) Location of ca certificates file to use for swift client requests.

**Table 6.5. Description of common configuration options**

Configuration option = Default value	Description
[DEFAULT]	
admin_project_domain_name = default	(StrOpt) The name of the domain for the service project(ex. tenant).
admin_user_domain_name = default	(StrOpt) The name of the domain to which the admin user belongs.
api_workers = 0	(IntOpt) Number of workers for Sahara API service (0 means all-in-one-thread configuration).
cleanup_time_for_incomplete_clusters = 0	(IntOpt) 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).
cluster_remote_threshold = 70	(IntOpt) The same as global_remote_threshold, but for a single cluster.
compute_topology_file = etc/sahara/compute.topology	(StrOpt) File with nova compute topology. It should contain mapping between nova computes and racks.
disable_event_log = False	(BoolOpt) Disables event log feature.
enable_data_locality = False	(BoolOpt) Enables data locality for hadoop cluster. Also enables data locality for Swift used by hadoop. If enabled, 'compute_topology' and 'swift_topology' configuration parameters should point to OpenStack and Swift topology correspondingly.
enable_hypervisor_awareness = True	(BoolOpt) Enables four-level topology for data locality. Works only if corresponding plugin supports such mode.
enable_notifications = False	(BoolOpt) Enables sending notifications to Ceilometer
global_remote_threshold = 100	(IntOpt) Maximum number of remote operations that will be running at the same time. Note that each remote operation requires its own process to run.
infrastructure_engine = direct	(StrOpt) An engine which will be used to provision infrastructure for Hadoop cluster.
job_binary_max_KB = 5120	(IntOpt) Maximum length of job binary data in kilobytes that may be stored or retrieved in a single operation.

Configuration option = Default value	Description
<code>job_canceling_timeout = 300</code>	(IntOpt) Timeout for canceling job execution (in seconds). Sahara will try to cancel job execution during this time.
<code>job_workflow_postfix =</code>	(StrOpt) Postfix for storing jobs in hdfs. Will be added to '/user/<hdfs user>/' path.
<code>max_header_line = 16384</code>	(IntOpt) 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).
<code>memcached_servers = None</code>	(ListOpt) Memcached servers or None for in process cache.
<code>min_transient_cluster_active_time = 30</code>	(IntOpt) Minimal "lifetime" in seconds for a transient cluster. Cluster is guaranteed to be "alive" within this time period.
<code>node_domain = nova.local</code>	(StrOpt) The suffix of the node's FQDN. In nova-network that is the <code>dhcp_domain</code> config parameter.
<code>os_region_name = None</code>	(StrOpt) Region name used to get services endpoints.
<code>periodic_enable = True</code>	(BoolOpt) Enable periodic tasks.
<code>periodic_fuzzy_delay = 60</code>	(IntOpt) Range in seconds to randomly delay when starting the periodic task scheduler to reduce stampeding. (Disable by setting to 0).
<code>periodic_interval_max = 60</code>	(IntOpt) Max interval size between periodic tasks execution in seconds.
<code>plugins = vanilla, hdp, spark, cdh</code>	(ListOpt) List of plugins to be loaded. Sahara preserves the order of the list when returning it.
<code>proxy_command =</code>	(StrOpt) 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}.
<code>remote = ssh</code>	(StrOpt) A method for Sahara to execute commands on VMs.
<code>rootwrap_command = sudo sahara-rootwrap / etc/sahara/rootwrap.conf</code>	(StrOpt) Rootwrap command to leverage. Use in conjunction with <code>use_rootwrap=True</code>
<code>run_external_periodic_tasks = True</code>	(BoolOpt) Some periodic tasks can be run in a separate process. Should we run them here?
<code>swift_topology_file = etc/sahara/swift.topology</code>	(StrOpt) File with Swift topology. It should contain mapping between Swift nodes and racks.
<code>use_external_key_manager = False</code>	(BoolOpt) Enable Sahara to use an external key manager service provided by the identity service catalog. Sahara will store all keys with the manager service.
<code>use_floating_ips = True</code>	(BoolOpt) If set to True, Sahara will use floating IPs to communicate with instances. To make sure that all instances have floating IPs assigned in Nova Network set "auto_assign_floating_ip=True" in nova.conf. If Neutron is used for networking, make sure that all Node Groups have "floating_ip_pool" parameter defined.
<code>use_identity_api_v3 = True</code>	(BoolOpt) Enables Sahara to use Keystone API v3. If that flag is disabled, per-job clusters will not be terminated automatically.
<code>use_namespaces = False</code>	(BoolOpt) Use network namespaces for communication (only valid to use in conjunction with <code>use_neutron=True</code> ).
<code>use_neutron = False</code>	(BoolOpt) Use Neutron Networking (False indicates the use of Nova networking).
<code>use_rootwrap = False</code>	(BoolOpt) Use rootwrap facility to allow non-root users to run the sahara-all server instance and access pri-

Configuration option = Default value	Description
	vate network IPs (only valid to use in conjunction with <code>use_namespaces=True</code> )
[conductor]	
<code>use_local = True</code>	(BoolOpt) Perform sahara-conductor operations locally.
[keystone_auth token]	
<code>memcached_servers = None</code>	(ListOpt) Optionally specify a list of memcached server(s) to use for caching. If left undefined, tokens will instead be cached in-process.

**Table 6.6. Description of database configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>db_driver = sahara.db</code>	(StrOpt) Driver to use for database access.
[database]	
<code>backend = sqlalchemy</code>	(StrOpt) The back end to use for the database.
<code>connection = None</code>	(StrOpt) The SQLAlchemy connection string to use to connect to the database.
<code>connection_debug = 0</code>	(IntOpt) Verbosity of SQL debugging information: 0=None, 100=Everything.
<code>connection_trace = False</code>	(BoolOpt) Add Python stack traces to SQL as comment strings.
<code>db_inc_retry_interval = True</code>	(BoolOpt) If True, increases the interval between retries of a database operation up to <code>db_max_retry_interval</code> .
<code>db_max_retries = 20</code>	(IntOpt) Maximum retries in case of connection error or deadlock error before error is raised. Set to -1 to specify an infinite retry count.
<code>db_max_retry_interval = 10</code>	(IntOpt) If <code>db_inc_retry_interval</code> is set, the maximum seconds between retries of a database operation.
<code>db_retry_interval = 1</code>	(IntOpt) Seconds between retries of a database transaction.
<code>idle_timeout = 3600</code>	(IntOpt) Timeout before idle SQL connections are reaped.
<code>max_overflow = None</code>	(IntOpt) If set, use this value for <code>max_overflow</code> with SQLAlchemy.
<code>max_pool_size = None</code>	(IntOpt) Maximum number of SQL connections to keep open in a pool.
<code>max_retries = 10</code>	(IntOpt) Maximum number of database connection retries during startup. Set to -1 to specify an infinite retry count.
<code>min_pool_size = 1</code>	(IntOpt) Minimum number of SQL connections to keep open in a pool.
<code>mysql_sql_mode = TRADITIONAL</code>	(StrOpt) 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>
<code>pool_timeout = None</code>	(IntOpt) If set, use this value for <code>pool_timeout</code> with SQLAlchemy.
<code>retry_interval = 10</code>	(IntOpt) Interval between retries of opening a SQL connection.
<code>slave_connection = None</code>	(StrOpt) The SQLAlchemy connection string to use to connect to the slave database.
<code>sqlite_db = oslo.sqlite</code>	(StrOpt) The file name to use with SQLite.
<code>sqlite_synchronous = True</code>	(BoolOpt) If True, SQLite uses synchronous mode.

Configuration option = Default value	Description
<code>use_db_reconnect = False</code>	(BoolOpt) Enable the experimental use of database reconnect on connection lost.

**Table 6.7. Description of domain configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>proxy_user_domain_name = None</code>	(StrOpt) The domain Sahara will use to create new proxy users for Swift object access.
<code>proxy_user_role_names = Member</code>	(ListOpt) A list of the role names that the proxy user should assume through trust for Swift object access.
<code>use_domain_for_proxy_users = False</code>	(BoolOpt) 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.

**Table 6.8. Description of logging configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>debug = False</code>	(BoolOpt) Print debugging output (set logging level to DEBUG instead of default WARNING level).
<code>default_log_levels = amqpplib=WARN, qpid.messaging=INFO, stevedore=INFO, eventlet.wsgi.server=WARN, sqlalchemy=WARN, boto=WARN, suds=INFO, keystone=INFO, paramiko=WARN, requests=WARN, iso8601=WARN, oslo_messaging=INFO</code>	(ListOpt) List of logger=LEVEL pairs.
<code>fatal_deprecations = False</code>	(BoolOpt) Enables or disables fatal status of deprecations.
<code>instance_format = "[instance: %(uuid)s] "</code>	(StrOpt) The format for an instance that is passed with the log message.
<code>instance_uuid_format = "[instance: %(uuid)s] "</code>	(StrOpt) The format for an instance UUID that is passed with the log message.
<code>log_config_append = None</code>	(StrOpt) 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.
<code>log_date_format = %Y-%m-%d %H:%M:%S</code>	(StrOpt) Format string for %(asctime)s in log records. Default: %(default)s.
<code>log_dir = None</code>	(StrOpt) (Optional) The base directory used for relative – log-file paths.
<code>log_file = None</code>	(StrOpt) (Optional) Name of log file to output to. If no default is set, logging will go to stdout.
<code>log_format = None</code>	(StrOpt) DEPRECATED. A logging.Formatter log message format string which may use any of the available logging.LogRecord attributes. This option is deprecated. Please use <code>logging_context_format_string</code> and <code>logging_default_format_string</code> instead.
<code>log_config_append = None</code>	(StrOpt) 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.
<code>log_date_format = %Y-%m-%d %H:%M:%S</code>	(StrOpt) Format string for %(asctime)s in log records. Default: %(default)s.
<code>log_dir = None</code>	(StrOpt) (Optional) The base directory used for relative – log-file paths.



Configuration option = Default value	Description
<code>log_file = None</code>	(StrOpt) (Optional) Name of log file to output to. If no default is set, logging will go to stdout.
<code>log_format = None</code>	(StrOpt) DEPRECATED. A logging.Formatter log message format string which may use any of the available logging.LogRecord attributes. This option is deprecated. Please use <code>logging_context_format_string</code> and <code>logging_default_format_string</code> instead.
<code>logging_context_format_string = %(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [%(request_id)s %(user_identity)s] %(instance)s%(message)s</code>	(StrOpt) Format string to use for log messages with context.
<code>logging_debug_format_suffix = %(funcName)s %(pathname)s:%(lineno)d</code>	(StrOpt) Data to append to log format when level is DEBUG.
<code>logging_default_format_string = %(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [-] %(instance)s%(message)s</code>	(StrOpt) Format string to use for log messages without context.
<code>logging_exception_prefix = %(asctime)s.%(msecs)03d %(process)d TRACE %(name)s %(instance)s</code>	(StrOpt) Prefix each line of exception output with this format.
<code>publish_errors = False</code>	(BoolOpt) Enables or disables publication of error events.
<code>syslog_log_facility = LOG_USER</code>	(StrOpt) Syslog facility to receive log lines.
<code>syslog_log_facility = LOG_USER</code>	(StrOpt) Syslog facility to receive log lines.
<code>use_syslog = False</code>	(BoolOpt) Use syslog for logging. Existing syslog format is DEPRECATED during I, and will change in J to honor RFC5424.
<code>use_syslog_rfc_format = False</code>	(BoolOpt) (Optional) Enables or disables syslog rfc5424 format for logging. If enabled, prefixes the MSG part of the syslog message with APP-NAME (RFC5424). The format without the APP-NAME is deprecated in I, and will be removed in J.
<code>use_stderr = True</code>	(BoolOpt) Log output to standard error.
<code>use_syslog = False</code>	(BoolOpt) Use syslog for logging. Existing syslog format is DEPRECATED during I, and will change in J to honor RFC5424.
<code>use_syslog_rfc_format = False</code>	(BoolOpt) (Optional) Enables or disables syslog rfc5424 format for logging. If enabled, prefixes the MSG part of the syslog message with APP-NAME (RFC5424). The format without the APP-NAME is deprecated in I, and will be removed in J.
<code>verbose = False</code>	(BoolOpt) Print more verbose output (set logging level to INFO instead of default WARNING level).

**Table 6.9. Description of oslo\_middleware configuration options**

Configuration option = Default value	Description
[oslo_middleware]	
<code>max_request_body_size = 114688</code>	(IntOpt) The maximum body size for each request, in bytes.

**Table 6.10. Description of policy configuration options**

Configuration option = Default value	Description
[oslo_policy]	
<code>policy_default_rule = default</code>	(StrOpt) Default rule. Enforced when a requested rule is not found.



Configuration option = Default value	Description
kombu_ssl_certfile =	(StrOpt) SSL cert file (valid only if SSL enabled).
kombu_ssl_keyfile =	(StrOpt) SSL key file (valid only if SSL enabled).
kombu_ssl_version =	(StrOpt) 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.
rabbit_ha_queues = <i>False</i>	(BoolOpt) Use HA queues in RabbitMQ (x-ha-policy: all). If you change this option, you must wipe the RabbitMQ database.
rabbit_host = <i>localhost</i>	(StrOpt) The RabbitMQ broker address where a single node is used.
rabbit_hosts = <i>\$rabbit_host:\$rabbit_port</i>	(ListOpt) RabbitMQ HA cluster host:port pairs.
rabbit_login_method = <i>AMQPLAIN</i>	(StrOpt) The RabbitMQ login method.
rabbit_max_retries = <i>0</i>	(IntOpt) Maximum number of RabbitMQ connection retries. Default is 0 (infinite retry count).
rabbit_password = <i>guest</i>	(StrOpt) The RabbitMQ password.
rabbit_port = <i>5672</i>	(IntOpt) The RabbitMQ broker port where a single node is used.
rabbit_retry_backoff = <i>2</i>	(IntOpt) How long to backoff for between retries when connecting to RabbitMQ.
rabbit_retry_interval = <i>1</i>	(IntOpt) How frequently to retry connecting with RabbitMQ.
rabbit_use_ssl = <i>False</i>	(BoolOpt) Connect over SSL for RabbitMQ.
rabbit_userid = <i>guest</i>	(StrOpt) The RabbitMQ userid.
rabbit_virtual_host = <i>/</i>	(StrOpt) The RabbitMQ virtual host.
rpc_conn_pool_size = <i>30</i>	(IntOpt) Size of RPC connection pool.

**Table 6.13. Description of Redis configuration options**

Configuration option = Default value	Description
[matchmaker_redis]	
host = <i>127.0.0.1</i>	(StrOpt) Host to locate redis.
password = <i>None</i>	(StrOpt) Password for Redis server (optional).
port = <i>6379</i>	(IntOpt) Use this port to connect to redis host.
[matchmaker_ring]	
ringfile = <i>/etc/oslo/matchmaker_ring.json</i>	(StrOpt) Matchmaker ring file (JSON).

**Table 6.14. Description of RPC configuration options**

Configuration option = Default value	Description
[DEFAULT]	
matchmaker_heartbeat_freq = <i>300</i>	(IntOpt) Heartbeat frequency.
matchmaker_heartbeat_ttl = <i>600</i>	(IntOpt) Heartbeat time-to-live.
rpc_backend = <i>rabbit</i>	(StrOpt) The messaging driver to use, defaults to rabbit. Other drivers include qpid and zmq.
rpc_cast_timeout = <i>30</i>	(IntOpt) Seconds to wait before a cast expires (TTL). Only supported by impl_zmq.
rpc_response_timeout = <i>60</i>	(IntOpt) Seconds to wait for a response from a call.
rpc_thread_pool_size = <i>64</i>	(IntOpt) Size of RPC thread pool.
[oslo_messaging_amqp]	
allow_insecure_clients = <i>False</i>	(BoolOpt) Accept clients using either SSL or plain TCP

Configuration option = Default value	Description
broadcast_prefix = <i>broadcast</i>	(StrOpt) address prefix used when broadcasting to all servers
container_name = <i>None</i>	(StrOpt) Name for the AMQP container
group_request_prefix = <i>unicast</i>	(StrOpt) address prefix when sending to any server in group
idle_timeout = <i>0</i>	(IntOpt) Timeout for inactive connections (in seconds)
server_request_prefix = <i>exclusive</i>	(StrOpt) address prefix used when sending to a specific server
ssl_ca_file =	(StrOpt) CA certificate PEM file to verify server certificate
ssl_cert_file =	(StrOpt) Identifying certificate PEM file to present to clients
ssl_key_file =	(StrOpt) Private key PEM file used to sign cert_file certificate
ssl_key_password = <i>None</i>	(StrOpt) Password for decrypting ssl_key_file (if encrypted)
trace = <i>False</i>	(BoolOpt) Debug: dump AMQP frames to stdout

**Table 6.15. Description of timeouts configuration options**

Configuration option = Default value	Description
[timeouts]	
await_attach_volumes = <i>10</i>	(IntOpt) Wait for attaching volumes to instances, in seconds
await_for_instances_active = <i>10800</i>	(IntOpt) Wait for instances to become active, in seconds
delete_instances_timeout = <i>10800</i>	(IntOpt) Wait for instances to be deleted, in seconds
detach_volume_timeout = <i>300</i>	(IntOpt) Timeout for detaching volumes from instance, in seconds
ips_assign_timeout = <i>10800</i>	(IntOpt) Assign IPs timeout, in seconds
volume_available_timeout = <i>10800</i>	(IntOpt) Wait for volumes to become available, in seconds
wait_until_accessible = <i>10800</i>	(IntOpt) Wait for instance accessibility, in seconds

**Table 6.16. Description of ZeroMQ configuration options**

Configuration option = Default value	Description
[DEFAULT]	
rpc_zmq_bind_address = *	(StrOpt) ZeroMQ bind address. Should be a wildcard (*), an ethernet interface, or IP. The "host" option should point or resolve to this address.
rpc_zmq_contexts = <i>1</i>	(IntOpt) Number of ZeroMQ contexts, defaults to 1.
rpc_zmq_host = <i>localhost</i>	(StrOpt) Name of this node. Must be a valid hostname, FQDN, or IP address. Must match "host" option, if running Nova.
rpc_zmq_ipc_dir = <i>/var/run/openstack</i>	(StrOpt) Directory for holding IPC sockets.
rpc_zmq_matchmaker = <i>local</i>	(StrOpt) MatchMaker driver.
rpc_zmq_port = <i>9501</i>	(IntOpt) ZeroMQ receiver listening port.
rpc_zmq_topic_backlog = <i>None</i>	(IntOpt) Maximum number of ingress messages to locally buffer per topic. Default is unlimited.





Option = default value	(Type) Help string
[oslo_messaging_qpid] qpid_password =	(StrOpt) Password for Qpid connection.
[oslo_messaging_qpid] qpid_port = 5672	(IntOpt) Qpid broker port.
[oslo_messaging_qpid] qpid_protocol = tcp	(StrOpt) Transport to use, either 'tcp' or 'ssl'.
[oslo_messaging_qpid] qpid_receiver_capacity = 1	(IntOpt) The number of prefetched messages held by receiver.
[oslo_messaging_qpid] qpid_sasl_mechanisms =	(StrOpt) Space separated list of SASL mechanisms to use for auth.
[oslo_messaging_qpid] qpid_tcp_nodelay = True	(BoolOpt) Whether to disable the Nagle algorithm.
[oslo_messaging_qpid] qpid_topology_version = 1	(IntOpt) The qpid topology version to use. Version 1 is what was originally used by impl_qpid. Version 2 includes some backwards-incompatible changes that allow broker federation to work. Users should update to version 2 when they are able to take everything down, as it requires a clean break.
[oslo_messaging_qpid] qpid_username =	(StrOpt) Username for Qpid connection.
[oslo_messaging_qpid] rpc_conn_pool_size = 30	(IntOpt) Size of RPC connection pool.
[oslo_messaging_rabbit] amqp_auto_delete = False	(BoolOpt) Auto-delete queues in AMQP.
[oslo_messaging_rabbit] amqp_durable_queues = False	(BoolOpt) Use durable queues in AMQP.
[oslo_messaging_rabbit] fake_rabbit = False	(BoolOpt) Deprecated, use rpc_backend=kombu+memory or rpc_backend=fake
[oslo_messaging_rabbit] heartbeat_rate = 2	(IntOpt) How often times during the heartbeat_timeout_threshold we check the heartbeat.
[oslo_messaging_rabbit] heartbeat_timeout_threshold = 0	(IntOpt) Number of seconds after which the Rabbit broker is considered down if heartbeat's keep-alive fails (0 disables the heartbeat, >0 enables it. Enabling heartbeats requires kombu>=3.0.7 and amqp>=1.4.0). EXPERIMENTAL
[oslo_messaging_rabbit] kombu_reconnect_delay = 1.0	(FloatOpt) How long to wait before reconnecting in response to an AMQP consumer cancel notification.
[oslo_messaging_rabbit] kombu_ssl_ca_certs =	(StrOpt) SSL certification authority file (valid only if SSL enabled).
[oslo_messaging_rabbit] kombu_ssl_certfile =	(StrOpt) SSL cert file (valid only if SSL enabled).
[oslo_messaging_rabbit] kombu_ssl_keyfile =	(StrOpt) SSL key file (valid only if SSL enabled).
[oslo_messaging_rabbit] kombu_ssl_version =	(StrOpt) 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.
[oslo_messaging_rabbit] rabbit_ha_queues = False	(BoolOpt) Use HA queues in RabbitMQ (x-ha-policy: all). If you change this option, you must wipe the RabbitMQ database.
[oslo_messaging_rabbit] rabbit_host = localhost	(StrOpt) The RabbitMQ broker address where a single node is used.
[oslo_messaging_rabbit] rabbit_hosts = \$rabbit_host:\$rabbit_port	(ListOpt) RabbitMQ HA cluster host:port pairs.
[oslo_messaging_rabbit] rabbit_login_method = AMQPLAIN	(StrOpt) The RabbitMQ login method.
[oslo_messaging_rabbit] rabbit_max_retries = 0	(IntOpt) Maximum number of RabbitMQ connection retries. Default is 0 (infinite retry count).
[oslo_messaging_rabbit] rabbit_password = guest	(StrOpt) The RabbitMQ password.
[oslo_messaging_rabbit] rabbit_port = 5672	(IntOpt) The RabbitMQ broker port where a single node is used.
[oslo_messaging_rabbit] rabbit_retry_backoff = 2	(IntOpt) How long to backoff for between retries when connecting to RabbitMQ.

Option = default value	(Type) Help string
[oslo_messaging_rabbit] rabbit_retry_interval = 1	(IntOpt) How frequently to retry connecting with RabbitMQ.
[oslo_messaging_rabbit] rabbit_use_ssl = False	(BoolOpt) Connect over SSL for RabbitMQ.
[oslo_messaging_rabbit] rabbit_userid = guest	(StrOpt) The RabbitMQ userid.
[oslo_messaging_rabbit] rabbit_virtual_host = /	(StrOpt) The RabbitMQ virtual host.
[oslo_messaging_rabbit] rpc_conn_pool_size = 30	(IntOpt) Size of RPC connection pool.
[oslo_middleware] max_request_body_size = 114688	(IntOpt) The maximum body size for each request, in bytes.
[oslo_policy] policy_default_rule = default	(StrOpt) Default rule. Enforced when a requested rule is not found.
[oslo_policy] policy_dirs = ['policy.d']	(MultiStrOpt) 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.
[oslo_policy] policy_file = policy.json	(StrOpt) The JSON file that defines policies.
[ssl] ca_file = None	(StrOpt) CA certificate file to use to verify connecting clients.
[ssl] cert_file = None	(StrOpt) Certificate file to use when starting the server securely.
[ssl] key_file = None	(StrOpt) Private key file to use when starting the server securely.
[swift] api_insecure = False	(BoolOpt) Allow to perform insecure SSL requests to swift.
[swift] ca_file = None	(StrOpt) Location of ca certificates file to use for swift client requests.
[timeouts] await_attach_volumes = 10	(IntOpt) Wait for attaching volumes to instances, in seconds
[timeouts] await_for_instances_active = 10800	(IntOpt) Wait for instances to become active, in seconds
[timeouts] delete_instances_timeout = 10800	(IntOpt) Wait for instances to be deleted, in seconds
[timeouts] detach_volume_timeout = 300	(IntOpt) Timeout for detaching volumes from instance, in seconds
[timeouts] ips_assign_timeout = 10800	(IntOpt) Assign IPs timeout, in seconds
[timeouts] volume_available_timeout = 10800	(IntOpt) Wait for volumes to become available, in seconds
[timeouts] wait_until_accessible = 10800	(IntOpt) Wait for instance accessibility, in seconds

**Table 6.18. New default values**

Option	Previous default value	New default value
[DEFAULT] default_log_levels	amqpplib=WARN, qpid.messaging=INFO, stevedore=INFO, eventlet.wsgi.server=WARN, sqlalchemy=WARN, boto=WARN, suds=INFO, keystone=INFO, paramiko=WARN, requests=WARN, iso8601=WARN	amqpplib=WARN, qpid.messaging=INFO, stevedore=INFO, eventlet.wsgi.server=WARN, sqlalchemy=WARN, boto=WARN, suds=INFO, keystone=INFO, paramiko=WARN, requests=WARN, iso8601=WARN, oslo_messaging=INFO
[DEFAULT] plugins	vanilla, hdp, spark	vanilla, hdp, spark, cdh
[DEFAULT] rpc_zmq_matchmaker	oslo.messaging._drivers.matchmaker.Matchmaker	MatchmakerLocalhost





















































```
# value)
#admin_token = ADMIN

# (Deprecated) The port which the OpenStack Compute service listens on. This
# option was only used for string replacement in the templated catalog
# backend.
# Templated catalogs should replace the "${compute_port}s" substitution with
# the static port of the compute service. As of Juno, this option is
# deprecated
# and will be removed in the L release. (integer value)
#compute_port = 8774

# 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. E.g. a request to
# http://server:5000/v3/users will default to http://server:5000. You should
# only need to set this value if the base URL contains a path (e.g. /prefix/
# v3)
# or the endpoint should be found on a different server. (string value)
#public_endpoint = <None>

# 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. E.g. a request to http://server:35357/v3/users
# will default to http://server:35357. You should only need to set this value
# if the base URL contains a path (e.g. /prefix/v3) or the endpoint should be
# found on a different server. (string value)
#admin_endpoint = <None>

# Maximum depth of the project hierarchy. WARNING: setting it to a large value
# may adversely impact performance. (integer value)
#max_project_tree_depth = 5

# Limit the sizes of user & project ID/names. (integer value)
#max_param_size = 64

# Similar to max_param_size, but provides an exception for token values.
# (integer value)
#max_token_size = 8192

# Similar to the member_role_name 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.
# (string value)
#member_role_id = 9fe2ff9ee4384b1894a90878d3e92bab

# This is the role name used in combination with the member_role_id option;
# see
# that option for more detail. (string value)
#member_role_name = _member_

# The value passed as the keyword "rounds" to passlib's encrypt method.
# (integer value)
#crypt_strength = 40000

# The maximum number of entities that will be returned in a collection, with
# no
# limit set by default. This global limit may be then overridden for a
# specific
```





```
# 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. (string value)
# Deprecated group/name - [DEFAULT]/log_config
#log_config_append = <None>

# DEPRECATED. A logging.Formatter log message format string which may use any
# of the available logging.LogRecord attributes. This option is deprecated.
# Please use logging_context_format_string and logging_default_format_string
# instead. (string value)
#log_format = <None>

# Format string for %(asctime)s in log records. Default: %(default)s .
# (string
# value)
#log_date_format = %Y-%m-%d %H:%M:%S

# (Optional) Name of log file to output to. If no default is set, logging will
# go to stdout. (string value)
# Deprecated group/name - [DEFAULT]/logfile
#log_file = <None>

# (Optional) The base directory used for relative --log-file paths. (string
# value)
# Deprecated group/name - [DEFAULT]/logdir
#log_dir = <None>

# Use syslog for logging. Existing syslog format is DEPRECATED during I, and
# will change in J to honor RFC5424. (boolean value)
#use_syslog = false

# (Optional) Enables or disables syslog rfc5424 format for logging. If
# enabled,
# prefixes the MSG part of the syslog message with APP-NAME (RFC5424). The
# format without the APP-NAME is deprecated in I, and will be removed in J.
# (boolean value)
#use_syslog_rfc_format = false

# Syslog facility to receive log lines. (string value)
#syslog_log_facility = LOG_USER

# Log output to standard error. (boolean value)
#use_stderr = true

# Format string to use for log messages with context. (string value)
#logging_context_format_string = %(asctime)s.%(msecs)03d %(process)d
# %(levelname)s %(name)s [%(request_id)s %(user_identity)s] %(instance)s
# %(message)s

# Format string to use for log messages without context. (string value)
#logging_default_format_string = %(asctime)s.%(msecs)03d %(process)d
# %(levelname)s %(name)s [-] %(instance)s%(message)s

# Data to append to log format when level is DEBUG. (string value)
#logging_debug_format_suffix = %(funcName)s %(pathname)s:%(lineno)d

# Prefix each line of exception output with this format. (string value)
#logging_exception_prefix = %(asctime)s.%(msecs)03d %(process)d TRACE %(name)s
# %(instance)s
```

```

# List of logger=LEVEL pairs. (list value)
#default_log_levels = amqp=WARN,amqpplib=WARN,boto=WARN,qpidd=WARN,sqlalchemy=
WARN,suds=INFO,oslo.messaging=INFO,iso8601=WARN,requests.packages.urllib3.
connectionpool=WARN,urllib3.connectionpool=WARN,websocket=WARN,requests.
packages.urllib3.util.retry=WARN,urllib3.util.retry=WARN,keystonemiddleware=
WARN,routes.middleware=WARN,stevedore=WARN

# Enables or disables publication of error events. (boolean value)
#publish_errors = false

# Enables or disables fatal status of deprecations. (boolean value)
#fatal_deprecations = false

# The format for an instance that is passed with the log message. (string
# value)
#instance_format = "[instance: %(uuid)s] "

# The format for an instance UUID that is passed with the log message. (string
# value)
#instance_uuid_format = "[instance: %(uuid)s] "

#
# From oslo.messaging
#

# ZeroMQ bind address. Should be a wildcard (*), an ethernet interface, or IP.
# The "host" option should point or resolve to this address. (string value)
#rpc_zmq_bind_address = *

# MatchMaker driver. (string value)
#rpc_zmq_matchmaker = oslo_messaging._drivers.matchmaker.MatchMakerLocalhost

# ZeroMQ receiver listening port. (integer value)
#rpc_zmq_port = 9501

# Number of ZeroMQ contexts, defaults to 1. (integer value)
#rpc_zmq_contexts = 1

# Maximum number of ingress messages to locally buffer per topic. Default is
# unlimited. (integer value)
#rpc_zmq_topic_backlog = <None>

# Directory for holding IPC sockets. (string value)
#rpc_zmq_ipc_dir = /var/run/openstack

# Name of this node. Must be a valid hostname, FQDN, or IP address. Must match
# "host" option, if running Nova. (string value)
#rpc_zmq_host = localhost

# Seconds to wait before a cast expires (TTL). Only supported by impl_zmq.
# (integer value)
#rpc_cast_timeout = 30

# Heartbeat frequency. (integer value)
#matchmaker_heartbeat_freq = 300

# Heartbeat time-to-live. (integer value)
#matchmaker_heartbeat_ttl = 600

```

```
# Size of RPC thread pool. (integer value)
#rpc_thread_pool_size = 64

# Driver or drivers to handle sending notifications. (multi valued)
#notification_driver =

# AMQP topic used for OpenStack notifications. (list value)
# Deprecated group/name - [rpc_notifier2]/topics
#notification_topics = notifications

# Seconds to wait for a response from a call. (integer value)
#rpc_response_timeout = 60

# A URL representing the messaging driver to use and its full configuration.
  If
# not set, we fall back to the rpc_backend option and driver specific
# configuration. (string value)
#transport_url = <None>

# The messaging driver to use, defaults to rabbit. Other drivers include qpid
# and zmq. (string value)
#rpc_backend = rabbit

# The default exchange under which topics are scoped. May be overridden by an
# exchange name specified in the transport_url option. (string value)
#control_exchange = keystone

[assignment]

#
# From keystone
#

# Assignment backend driver. (string value)
#driver = <None>

[auth]

#
# From keystone
#

# Default auth methods. (list value)
#methods = external,password,token,oauth1

# The password auth plugin module. (string value)
#password = keystone.auth.plugins.password.Password

# The token auth plugin module. (string value)
#token = keystone.auth.plugins.token.Token

# The external (REMOTE_USER) auth plugin module. (string value)
#external = keystone.auth.plugins.external.DefaultDomain

# The oAuth1.0 auth plugin module. (string value)
#oauth1 = keystone.auth.plugins.oauth1.OAuth
```

```
[cache]

#
# From keystone
#

# Prefix for building the configuration dictionary for the cache region. This
# should not need to be changed unless there is another dogpile.cache region
# with the same configuration name. (string value)
#config_prefix = cache.keystone

# 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. (integer value)
#expiration_time = 600

# Dogpile.cache backend module. It is recommended that Memcache with pooling
# (keystone.cache.memcache_pool) or Redis (dogpile.cache.redis) be used in
# production deployments. Small workloads (single process) like devstack can
# use the dogpile.cache.memory backend. (string value)
#backend = keystone.common.cache.noop

# Arguments supplied to the backend module. Specify this option once per
# argument to be passed to the dogpile.cache backend. Example format:
# "<argname>:<value>". (multi valued)
#backend_argument =

# Proxy classes to import that will affect the way the dogpile.cache backend
# functions. See the dogpile.cache documentation on changing-backend-behavior.
# (list value)
#proxies =

# Global toggle for all caching using the should_cache_fn mechanism. (boolean
# value)
#enabled = false

# 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. (boolean value)
#debug_cache_backend = false

# Memcache servers in the format of "host:port". (dogpile.cache.memcache and
# keystone.cache.memcache_pool backends only). (list value)
#memcache_servers = localhost:11211

# Number of seconds memcached server is considered dead before it is tried
# again. (dogpile.cache.memcache and keystone.cache.memcache_pool backends
# only). (integer value)
#memcache_dead_retry = 300

# Timeout in seconds for every call to a server. (dogpile.cache.memcache and
# keystone.cache.memcache_pool backends only). (integer value)
#memcache_socket_timeout = 3

# Max total number of open connections to every memcached server.
# (keystone.cache.memcache_pool backend only). (integer value)
#memcache_pool_maxsize = 10
```

```
# Number of seconds a connection to memcached is held unused in the pool
before
# it is closed. (keystone.cache.memcache_pool backend only). (integer value)
#memcache_pool_unused_timeout = 60

# Number of seconds that an operation will wait to get a memcache client
# connection. (integer value)
#memcache_pool_connection_get_timeout = 10

[catalog]

#
# From keystone
#

# Catalog template file name for use with the template catalog backend.
(string
# value)
#template_file = default_catalog.templates

# Catalog backend driver. (string value)
#driver = keystone.catalog.backends.sql.Catalog

# Toggle for catalog caching. This has no effect unless global caching is
# enabled. (boolean value)
#caching = true

# Time to cache catalog data (in seconds). This has no effect unless global
and
# catalog caching are enabled. (integer value)
#cache_time = <None>

# Maximum number of entities that will be returned in a catalog collection.
# (integer value)
#list_limit = <None>

[credential]

#
# From keystone
#

# Credential backend driver. (string value)
#driver = keystone.credential.backends.sql.Credential

[database]

#
# From oslo.db
#

# The file name to use with SQLite. (string value)
# Deprecated group/name - [DEFAULT]/sqlite_db
#sqlite_db = oslo.sqlite

# If True, SQLite uses synchronous mode. (boolean value)
# Deprecated group/name - [DEFAULT]/sqlite_synchronous
```

```
#sqlite_synchronous = true

# The back end to use for the database. (string value)
# Deprecated group/name - [DEFAULT]/db_backend
#backend = sqlalchemy

# The SQLAlchemy connection string to use to connect to the database. (string
# value)
# Deprecated group/name - [DEFAULT]/sql_connection
# Deprecated group/name - [DATABASE]/sql_connection
# Deprecated group/name - [sql]/connection
#connection = <None>

# The SQLAlchemy connection string to use to connect to the slave database.
# (string value)
#slave_connection = <None>

# 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=
# (string value)
#mysql_sql_mode = TRADITIONAL

# Timeout before idle SQL connections are reaped. (integer value)
# Deprecated group/name - [DEFAULT]/sql_idle_timeout
# Deprecated group/name - [DATABASE]/sql_idle_timeout
# Deprecated group/name - [sql]/idle_timeout
#idle_timeout = 3600

# Minimum number of SQL connections to keep open in a pool. (integer value)
# Deprecated group/name - [DEFAULT]/sql_min_pool_size
# Deprecated group/name - [DATABASE]/sql_min_pool_size
#min_pool_size = 1

# Maximum number of SQL connections to keep open in a pool. (integer value)
# Deprecated group/name - [DEFAULT]/sql_max_pool_size
# Deprecated group/name - [DATABASE]/sql_max_pool_size
#max_pool_size = <None>

# Maximum number of database connection retries during startup. Set to -1 to
# specify an infinite retry count. (integer value)
# Deprecated group/name - [DEFAULT]/sql_max_retries
# Deprecated group/name - [DATABASE]/sql_max_retries
#max_retries = 10

# Interval between retries of opening a SQL connection. (integer value)
# Deprecated group/name - [DEFAULT]/sql_retry_interval
# Deprecated group/name - [DATABASE]/reconnect_interval
#retry_interval = 10

# If set, use this value for max_overflow with SQLAlchemy. (integer value)
# Deprecated group/name - [DEFAULT]/sql_max_overflow
# Deprecated group/name - [DATABASE]/sqlalchemy_max_overflow
#max_overflow = <None>

# Verbosity of SQL debugging information: 0=None, 100=Everything. (integer
# value)
# Deprecated group/name - [DEFAULT]/sql_connection_debug
#connection_debug = 0
```











```
# Identity backend driver. (string value)
#driver = keystone.identity.backends.sql.Identity

# Toggle for identity caching. This has no effect unless global caching is
# enabled. (boolean value)
#caching = true

# Time to cache identity data (in seconds). This has no effect unless global
# and identity caching are enabled. (integer value)
#cache_time = 600

# Maximum supported length for user passwords; decrease to improve
# performance.
# (integer value)
#max_password_length = 4096

# Maximum number of entities that will be returned in an identity collection.
# (integer value)
#list_limit = <None>

[identity_mapping]

#
# From keystone
#

# Keystone Identity Mapping backend driver. (string value)
#driver = keystone.identity.mapping_backends.sql.Mapping

# Public ID generator for user and group entities. The Keystone identity
# mapper
# only supports generators that produce no more than 64 characters. (string
# value)
#generator = keystone.identity.id_generators.sha256.Generator

# The format of user and group IDs changed in Juno for backends that do not
# generate UUIDs (e.g. 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, 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
# mapping for even the default LDAP driver. It is only safe to do this if you
# do not already have assignments for users and groups from the default LDAP
# domain, and it is acceptable for Keystone to provide the different IDs to
# clients than it did previously. Typically this means that the only time you
# can set this value to False is when configuring a fresh installation.
# (boolean value)
#backward_compatible_ids = true

[kvs]

#
# From keystone
#

# Extra dogpile.cache backend modules to register with the dogpile.cache
```



```
# 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. (integer value)
#debug_level = <None>

# Override the system's default referral chasing behavior for queries.
# (boolean
# value)
#chase_referrals = <None>

# Search base for users. (string value)
#user_tree_dn = <None>

# LDAP search filter for users. (string value)
#user_filter = <None>

# LDAP objectclass for users. (string value)
#user_objectclass = inetOrgPerson

# LDAP attribute mapped to user id. WARNING: must not be a multivalued
# attribute. (string value)
#user_id_attribute = cn

# LDAP attribute mapped to user name. (string value)
#user_name_attribute = sn

# LDAP attribute mapped to user email. (string value)
#user_mail_attribute = mail

# LDAP attribute mapped to password. (string value)
#user_pass_attribute = userPassword

# LDAP attribute mapped to user enabled flag. (string value)
#user_enabled_attribute = enabled

# Invert the meaning of the boolean enabled values. Some LDAP servers use a
# boolean lock attribute where "true" means an account is disabled. Setting
# "user_enabled_invert = true" will allow these lock attributes to be used.
# This setting will have no effect if "user_enabled_mask" or
# "user_enabled_emulation" settings are in use. (boolean value)
#user_enabled_invert = false

# Bitmask integer to indicate the bit that the enabled value is stored in if
# the LDAP server represents "enabled" as a bit on an integer rather than a
# boolean. A value of "0" indicates the mask is not used. If this is not set
# to
# "0" the typical value is "2". This is typically used when
# "user_enabled_attribute = userAccountControl". (integer value)
#user_enabled_mask = 0

# Default value to enable users. This should match an appropriate int 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 "True" the typical value is
# "512".
# This is typically used when "user_enabled_attribute = userAccountControl".
# (string value)
#user_enabled_default = True

# List of attributes stripped off the user on update. (list value)
```







```
#role_allow_create = true

# Allow role update in LDAP backend. (boolean value)
#role_allow_update = true

# Allow role deletion in LDAP backend. (boolean value)
#role_allow_delete = true

# Additional attribute mappings for roles. Attribute mapping format is
# <ldap_attr>:<user_attr>, where ldap_attr is the attribute in the LDAP entry
# and user_attr is the Identity API attribute. (list value)
#role_additional_attribute_mapping =

# Search base for groups. (string value)
#group_tree_dn = <None>

# LDAP search filter for groups. (string value)
#group_filter = <None>

# LDAP objectclass for groups. (string value)
#group_objectclass = groupOfNames

# LDAP attribute mapped to group id. (string value)
#group_id_attribute = cn

# LDAP attribute mapped to group name. (string value)
#group_name_attribute = ou

# LDAP attribute mapped to show group membership. (string value)
#group_member_attribute = member

# LDAP attribute mapped to group description. (string value)
#group_desc_attribute = description

# List of attributes stripped off the group on update. (list value)
#group_attribute_ignore =

# Allow group creation in LDAP backend. (boolean value)
#group_allow_create = true

# Allow group update in LDAP backend. (boolean value)
#group_allow_update = true

# Allow group deletion in LDAP backend. (boolean value)
#group_allow_delete = true

# Additional attribute mappings for groups. Attribute mapping format is
# <ldap_attr>:<user_attr>, where ldap_attr is the attribute in the LDAP entry
# and user_attr is the Identity API attribute. (list value)
#group_additional_attribute_mapping =

# CA certificate file path for communicating with LDAP servers. (string value)
#tls_cacertfile = <None>

# CA certificate directory path for communicating with LDAP servers. (string
# value)
#tls_cacertdir = <None>

# Enable TLS for communicating with LDAP servers. (boolean value)
#use_tls = false
```

```
# Valid options for tls_req_cert are demand, never, and allow. (string value)
#tls_req_cert = demand

# Enable LDAP connection pooling. (boolean value)
#use_pool = false

# Connection pool size. (integer value)
#pool_size = 10

# Maximum count of reconnect trials. (integer value)
#pool_retry_max = 3

# Time span in seconds to wait between two reconnect trials. (floating point
# value)
#pool_retry_delay = 0.1

# Connector timeout in seconds. Value -1 indicates indefinite wait for
# response. (integer value)
#pool_connection_timeout = -1

# Connection lifetime in seconds. (integer value)
#pool_connection_lifetime = 600

# Enable LDAP connection pooling for end user authentication. If use_pool is
# disabled, then this setting is meaningless and is not used at all. (boolean
# value)
#use_auth_pool = false

# End user auth connection pool size. (integer value)
#auth_pool_size = 100

# End user auth connection lifetime in seconds. (integer value)
#auth_pool_connection_lifetime = 60

[matchmaker_redis]

#
# From oslo.messaging
#

# Host to locate redis. (string value)
#host = 127.0.0.1

# Use this port to connect to redis host. (integer value)
#port = 6379

# Password for Redis server (optional). (string value)
#password = <None>

[matchmaker_ring]

#
# From oslo.messaging
#

# Matchmaker ring file (JSON). (string value)
# Deprecated group/name - [DEFAULT]/matchmaker_ringfile
```

```
#ringfile = /etc/oslo/matchmaker_ring.json

[memcache]

#
# From keystone
#

# Memcache servers in the format of "host:port". (list value)
#servers = localhost:11211

# Number of seconds memcached server is considered dead before it is tried
# again. This is used by the key value store system (e.g. token pooled
# memcached persistence backend). (integer value)
#dead_retry = 300

# Timeout in seconds for every call to a server. This is used by the key value
# store system (e.g. token pooled memcached persistence backend). (integer
# value)
#socket_timeout = 3

# Max total number of open connections to every memcached server. This is used
# by the key value store system (e.g. token pooled memcached persistence
# backend). (integer value)
#pool_maxsize = 10

# 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 (e.g. token pooled
# memcached persistence backend). (integer value)
#pool_unused_timeout = 60

# Number of seconds that an operation will wait to get a memcache client
# connection. This is used by the key value store system (e.g. token pooled
# memcached persistence backend). (integer value)
#pool_connection_get_timeout = 10

[oauth1]

#
# From keystone
#

# Credential backend driver. (string value)
#driver = keystone.contrib.oauth1.backends.sql.OAuth1

# Duration (in seconds) for the OAuth Request Token. (integer value)
#request_token_duration = 28800

# Duration (in seconds) for the OAuth Access Token. (integer value)
#access_token_duration = 86400

[os_inherit]

#
# From keystone
#
```

```
# role-assignment inheritance to projects from owning domain or from projects
# higher in the hierarchy can be optionally enabled. (boolean value)
#enabled = false

[oslo_messaging_amqp]

#
# From oslo.messaging
#

# address prefix used when sending to a specific server (string value)
# Deprecated group/name - [amqp1]/server_request_prefix
#server_request_prefix = exclusive

# address prefix used when broadcasting to all servers (string value)
# Deprecated group/name - [amqp1]/broadcast_prefix
#broadcast_prefix = broadcast

# address prefix when sending to any server in group (string value)
# Deprecated group/name - [amqp1]/group_request_prefix
#group_request_prefix = unicast

# Name for the AMQP container (string value)
# Deprecated group/name - [amqp1]/container_name
#container_name = <None>

# Timeout for inactive connections (in seconds) (integer value)
# Deprecated group/name - [amqp1]/idle_timeout
#idle_timeout = 0

# Debug: dump AMQP frames to stdout (boolean value)
# Deprecated group/name - [amqp1]/trace
#trace = false

# CA certificate PEM file for verifying server certificate (string value)
# Deprecated group/name - [amqp1]/ssl_ca_file
#ssl_ca_file =

# Identifying certificate PEM file to present to clients (string value)
# Deprecated group/name - [amqp1]/ssl_cert_file
#ssl_cert_file =

# Private key PEM file used to sign cert_file certificate (string value)
# Deprecated group/name - [amqp1]/ssl_key_file
#ssl_key_file =

# Password for decrypting ssl_key_file (if encrypted) (string value)
# Deprecated group/name - [amqp1]/ssl_key_password
#ssl_key_password = <None>

# Accept clients using either SSL or plain TCP (boolean value)
# Deprecated group/name - [amqp1]/allow_insecure_clients
#allow_insecure_clients = false

[oslo_messaging_qpid]

#
```

```
# From oslo.messaging
#

# Use durable queues in AMQP. (boolean value)
# Deprecated group/name - [DEFAULT]/rabbit_durable_queues
#amqp_durable_queues = false

# Auto-delete queues in AMQP. (boolean value)
# Deprecated group/name - [DEFAULT]/amqp_auto_delete
#amqp_auto_delete = false

# Size of RPC connection pool. (integer value)
# Deprecated group/name - [DEFAULT]/rpc_conn_pool_size
#rpc_conn_pool_size = 30

# Qpid broker hostname. (string value)
# Deprecated group/name - [DEFAULT]/qpid_hostname
#qpid_hostname = localhost

# Qpid broker port. (integer value)
# Deprecated group/name - [DEFAULT]/qpid_port
#qpid_port = 5672

# Qpid HA cluster host:port pairs. (list value)
# Deprecated group/name - [DEFAULT]/qpid_hosts
#qpid_hosts = $qpid_hostname:$qpid_port

# Username for Qpid connection. (string value)
# Deprecated group/name - [DEFAULT]/qpid_username
#qpid_username =

# Password for Qpid connection. (string value)
# Deprecated group/name - [DEFAULT]/qpid_password
#qpid_password =

# Space separated list of SASL mechanisms to use for auth. (string value)
# Deprecated group/name - [DEFAULT]/qpid_sasl_mechanisms
#qpid_sasl_mechanisms =

# Seconds between connection keepalive heartbeats. (integer value)
# Deprecated group/name - [DEFAULT]/qpid_heartbeat
#qpid_heartbeat = 60

# Transport to use, either 'tcp' or 'ssl'. (string value)
# Deprecated group/name - [DEFAULT]/qpid_protocol
#qpid_protocol = tcp

# Whether to disable the Nagle algorithm. (boolean value)
# Deprecated group/name - [DEFAULT]/qpid_tcp_nodelay
#qpid_tcp_nodelay = true

# The number of prefetched messages held by receiver. (integer value)
# Deprecated group/name - [DEFAULT]/qpid_receiver_capacity
#qpid_receiver_capacity = 1

# The qpid topology version to use. Version 1 is what was originally used by
# impl_qpid. Version 2 includes some backwards-incompatible changes that
# allow
# broker federation to work. Users should update to version 2 when they are
# able to take everything down, as it requires a clean break. (integer value)
```

```
# Deprecated group/name - [DEFAULT]/qpuid_topology_version
#qpuid_topology_version = 1

[oslo_messaging_rabbit]

#
# From oslo.messaging
#

# Use durable queues in AMQP. (boolean value)
# Deprecated group/name - [DEFAULT]/rabbit_durable_queues
#amqp_durable_queues = false

# Auto-delete queues in AMQP. (boolean value)
# Deprecated group/name - [DEFAULT]/amqp_auto_delete
#amqp_auto_delete = false

# Size of RPC connection pool. (integer value)
# Deprecated group/name - [DEFAULT]/rpc_conn_pool_size
#rpc_conn_pool_size = 30

# 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. (string value)
# Deprecated group/name - [DEFAULT]/kombu_ssl_version
#kombu_ssl_version =

# SSL key file (valid only if SSL enabled). (string value)
# Deprecated group/name - [DEFAULT]/kombu_ssl_keyfile
#kombu_ssl_keyfile =

# SSL cert file (valid only if SSL enabled). (string value)
# Deprecated group/name - [DEFAULT]/kombu_ssl_certfile
#kombu_ssl_certfile =

# SSL certification authority file (valid only if SSL enabled). (string value)
# Deprecated group/name - [DEFAULT]/kombu_ssl_ca_certs
#kombu_ssl_ca_certs =

# How long to wait before reconnecting in response to an AMQP consumer cancel
# notification. (floating point value)
# Deprecated group/name - [DEFAULT]/kombu_reconnect_delay
#kombu_reconnect_delay = 1.0

# The RabbitMQ broker address where a single node is used. (string value)
# Deprecated group/name - [DEFAULT]/rabbit_host
#rabbit_host = localhost

# The RabbitMQ broker port where a single node is used. (integer value)
# Deprecated group/name - [DEFAULT]/rabbit_port
#rabbit_port = 5672

# RabbitMQ HA cluster host:port pairs. (list value)
# Deprecated group/name - [DEFAULT]/rabbit_hosts
#rabbit_hosts = $rabbit_host:$rabbit_port

# Connect over SSL for RabbitMQ. (boolean value)
# Deprecated group/name - [DEFAULT]/rabbit_use_ssl
#rabbit_use_ssl = false
```

```
# The RabbitMQ userid. (string value)
# Deprecated group/name - [DEFAULT]/rabbit_userid
#rabbit_userid = guest

# The RabbitMQ password. (string value)
# Deprecated group/name - [DEFAULT]/rabbit_password
#rabbit_password = guest

# The RabbitMQ login method. (string value)
# Deprecated group/name - [DEFAULT]/rabbit_login_method
#rabbit_login_method = AMQPPLAIN

# The RabbitMQ virtual host. (string value)
# Deprecated group/name - [DEFAULT]/rabbit_virtual_host
#rabbit_virtual_host = /

# How frequently to retry connecting with RabbitMQ. (integer value)
#rabbit_retry_interval = 1

# How long to backoff for between retries when connecting to RabbitMQ.
(integer
# value)
# Deprecated group/name - [DEFAULT]/rabbit_retry_backoff
#rabbit_retry_backoff = 2

# Maximum number of RabbitMQ connection retries. Default is 0 (infinite retry
# count). (integer value)
# Deprecated group/name - [DEFAULT]/rabbit_max_retries
#rabbit_max_retries = 0

# Use HA queues in RabbitMQ (x-ha-policy: all). If you change this option, you
# must wipe the RabbitMQ database. (boolean value)
# Deprecated group/name - [DEFAULT]/rabbit_ha_queues
#rabbit_ha_queues = false

# Number of seconds after which the Rabbit broker is considered down if
# heartbeat's keep-alive fails (0 disable the heartbeat). (integer value)
#heartbeat_timeout_threshold = 60

# How often times during the heartbeat_timeout_threshold we check the
# heartbeat. (integer value)
#heartbeat_rate = 2

# Deprecated, use rpc_backend=kombu+memory or rpc_backend=fake (boolean value)
# Deprecated group/name - [DEFAULT]/fake_rabbit
#fake_rabbit = false

[oslo_middleware]

#
# From oslo.middleware
#

# The maximum body size for each request, in bytes. (integer value)
# Deprecated group/name - [DEFAULT]/osapi_max_request_body_size
# Deprecated group/name - [DEFAULT]/max_request_body_size
#max_request_body_size = 114688
```

```
[oslo_policy]

#
# From oslo.policy
#

# The JSON file that defines policies. (string value)
# Deprecated group/name - [DEFAULT]/policy_file
#policy_file = policy.json

# Default rule. Enforced when a requested rule is not found. (string value)
# Deprecated group/name - [DEFAULT]/policy_default_rule
#policy_default_rule = default

# 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.
# (multi
# valued)
# Deprecated group/name - [DEFAULT]/policy_dirs
#policy_dirs = policy.d

[paste_deploy]

#
# From keystone
#

# Name of the paste configuration file that defines the available pipelines.
# (string value)
#config_file = keystone-paste.ini

[policy]

#
# From keystone
#

# Policy backend driver. (string value)
#driver = keystone.policy.backends.sql.Policy

# Maximum number of entities that will be returned in a policy collection.
# (integer value)
#list_limit = <None>

[resource]

#
# From keystone
#

# Resource backend driver. If a resource driver is not specified, the
# assignment driver will choose the resource driver. (string value)
#driver = <None>
```











```
# other than KVS, which stores events in memory. (boolean value)
#revoke_by_id = true

# Allow rescoping of scoped token. Setting allow_rescoped_scoped_token to
# false
# prevents a user from exchanging a scoped token for any other token. (boolean
# value)
#allow_rescope_scoped_token = true

# The hash algorithm to use for PKI tokens. This can be set to any algorithm
# that hashlib supports. WARNING: Before changing this value, the auth_token
# middleware must be configured with the hash_algorithms, otherwise token
# revocation will not be processed correctly. (string value)
#hash_algorithm = md5

[trust]

#
# From keystone
#

# Delegation and impersonation features can be optionally disabled. (boolean
# value)
#enabled = true

# Enable redelegation feature. (boolean value)
#allow_redelegation = false

# Maximum depth of trust redelegation. (integer value)
#max_redelegation_count = 3

# Trust backend driver. (string value)
#driver = keystone.trust.backends.sql.Trust
```

## keystone-paste.ini

Use the `keystone-paste.ini` file to configure the Web Service Gateway Interface (WSGI) middleware pipeline for the Identity service.

```
# Keystone PasteDeploy configuration file.

[filter:debug]
paste.filter_factory = keystone.common.wsgi:Debug.factory

[filter:request_id]
paste.filter_factory = oslo_middleware:RequestId.factory

[filter:build_auth_context]
paste.filter_factory = keystone.middleware:AuthContextMiddleware.factory

[filter:token_auth]
paste.filter_factory = keystone.middleware:TokenAuthMiddleware.factory

[filter:admin_token_auth]
paste.filter_factory = keystone.middleware:AdminTokenAuthMiddleware.factory
```

```
[filter:json_body]
paste.filter_factory = keystone.middleware:JsonBodyMiddleware.factory

[filter:user_crud_extension]
paste.filter_factory = keystone.contrib.user_crud:CrudExtension.factory

[filter:crud_extension]
paste.filter_factory = keystone.contrib.admin_crud:CrudExtension.factory

[filter:ec2_extension]
paste.filter_factory = keystone.contrib.ec2:Ec2Extension.factory

[filter:ec2_extension_v3]
paste.filter_factory = keystone.contrib.ec2:Ec2ExtensionV3.factory

[filter:federation_extension]
paste.filter_factory = keystone.contrib.federation.
routers:FederationExtension.factory

[filter:oauth1_extension]
paste.filter_factory = keystone.contrib.oauth1.routers:OAuth1Extension.factory

[filter:s3_extension]
paste.filter_factory = keystone.contrib.s3:S3Extension.factory

[filter:endpoint_filter_extension]
paste.filter_factory = keystone.contrib.endpoint_filter.
routers:EndpointFilterExtension.factory

[filter:endpoint_policy_extension]
paste.filter_factory = keystone.contrib.endpoint_policy.
routers:EndpointPolicyExtension.factory

[filter:simple_cert_extension]
paste.filter_factory = keystone.contrib.simple_cert:SimpleCertExtension.
factory

[filter:revoke_extension]
paste.filter_factory = keystone.contrib.revoke.routers:RevokeExtension.factory

[filter:url_normalize]
paste.filter_factory = keystone.middleware:NormalizingFilter.factory

[filter:sizelimit]
paste.filter_factory = oslo_middleware.sizelimit:RequestBodySizeLimiter.
factory

[app:public_service]
paste.app_factory = keystone.service:public_app_factory

[app:service_v3]
paste.app_factory = keystone.service:v3_app_factory

[app:admin_service]
paste.app_factory = keystone.service:admin_app_factory

[pipeline:public_api]
# The last item in this pipeline must be public_service or an equivalent
# application. It cannot be a filter.
```







## policy.json

Use the `policy.json` file to define additional access controls that apply to the Identity service.

```
{
  "admin_required": "role:admin or is_admin:1",
  "service_role": "role:service",
  "service_or_admin": "rule:admin_required or rule:service_role",
  "owner" : "user_id:%(user_id)s",
  "admin_or_owner": "rule:admin_required or rule:owner",
  "token_subject": "user_id:%(target.token.user_id)s",
  "admin_or_token_subject": "rule:admin_required or rule:token_subject",

  "default": "rule:admin_required",

  "identity:get_region": "",
  "identity:list_regions": "",
  "identity:create_region": "rule:admin_required",
  "identity:update_region": "rule:admin_required",
  "identity:delete_region": "rule:admin_required",

  "identity:get_service": "rule:admin_required",
  "identity:list_services": "rule:admin_required",
  "identity:create_service": "rule:admin_required",
  "identity:update_service": "rule:admin_required",
  "identity:delete_service": "rule:admin_required",

  "identity:get_endpoint": "rule:admin_required",
  "identity:list_endpoints": "rule:admin_required",
  "identity:create_endpoint": "rule:admin_required",
  "identity:update_endpoint": "rule:admin_required",
  "identity:delete_endpoint": "rule:admin_required",

  "identity:get_domain": "rule:admin_required",
  "identity:list_domains": "rule:admin_required",
  "identity:create_domain": "rule:admin_required",
  "identity:update_domain": "rule:admin_required",
  "identity:delete_domain": "rule:admin_required",

  "identity:get_project": "rule:admin_required",
  "identity:list_projects": "rule:admin_required",
  "identity:list_user_projects": "rule:admin_or_owner",
  "identity:create_project": "rule:admin_required",
  "identity:update_project": "rule:admin_required",
  "identity:delete_project": "rule:admin_required",

  "identity:get_user": "rule:admin_required",
  "identity:list_users": "rule:admin_required",
  "identity:create_user": "rule:admin_required",
  "identity:update_user": "rule:admin_required",
  "identity:delete_user": "rule:admin_required",
  "identity:change_password": "rule:admin_or_owner",

  "identity:get_group": "rule:admin_required",
  "identity:list_groups": "rule:admin_required",
  "identity:list_groups_for_user": "rule:admin_or_owner",
  "identity:create_group": "rule:admin_required",
  "identity:update_group": "rule:admin_required",
}
```



```

"identity:list_access_tokens": "rule:admin_required",
"identity:get_access_token": "rule:admin_required",
"identity:delete_access_token": "rule:admin_required",

"identity:list_projects_for_endpoint": "rule:admin_required",
"identity:add_endpoint_to_project": "rule:admin_required",
"identity:check_endpoint_in_project": "rule:admin_required",
"identity:list_endpoints_for_project": "rule:admin_required",
"identity:remove_endpoint_from_project": "rule:admin_required",

"identity:create_endpoint_group": "rule:admin_required",
"identity:list_endpoint_groups": "rule:admin_required",
"identity:get_endpoint_group": "rule:admin_required",
"identity:update_endpoint_group": "rule:admin_required",
"identity:delete_endpoint_group": "rule:admin_required",
"identity:list_projects_associated_with_endpoint_group":
"rule:admin_required",
"identity:list_endpoints_associated_with_endpoint_group":
"rule:admin_required",
"identity:get_endpoint_group_in_project": "rule:admin_required",
"identity:add_endpoint_group_to_project": "rule:admin_required",
"identity:remove_endpoint_group_from_project": "rule:admin_required",

"identity:create_identity_provider": "rule:admin_required",
"identity:list_identity_providers": "rule:admin_required",
"identity:get_identity_providers": "rule:admin_required",
"identity:update_identity_provider": "rule:admin_required",
"identity:delete_identity_provider": "rule:admin_required",

"identity:create_protocol": "rule:admin_required",
"identity:update_protocol": "rule:admin_required",
"identity:get_protocol": "rule:admin_required",
"identity:list_protocols": "rule:admin_required",
"identity:delete_protocol": "rule:admin_required",

"identity:create_mapping": "rule:admin_required",
"identity:get_mapping": "rule:admin_required",
"identity:list_mappings": "rule:admin_required",
"identity:delete_mapping": "rule:admin_required",
"identity:update_mapping": "rule:admin_required",

"identity:create_service_provider": "rule:admin_required",
"identity:list_service_providers": "rule:admin_required",
"identity:get_service_provider": "rule:admin_required",
"identity:update_service_provider": "rule:admin_required",
"identity:delete_service_provider": "rule:admin_required",

"identity:get_auth_catalog": "",
"identity:get_auth_projects": "",
"identity:get_auth_domains": "",

"identity:list_projects_for_groups": "",
"identity:list_domains_for_groups": "",

"identity:list_revoke_events": "",

"identity:create_policy_association_for_endpoint": "rule:admin_required",
"identity:check_policy_association_for_endpoint": "rule:admin_required",
"identity:delete_policy_association_for_endpoint": "rule:admin_required",
"identity:create_policy_association_for_service": "rule:admin_required",
    
```





Option = default value	(Type) Help string
[eventlet_server_ssl] certfile = /etc/keystone/ssl/certs/keystone.pem	(StrOpt) Path of the certfile for SSL. For non-production environments, you may be interested in using `keystone-manage ssl_setup` to generate self-signed certificates.
[eventlet_server_ssl] enable = False	(BoolOpt) Toggle for SSL support on the Keystone eventlet servers.
[eventlet_server_ssl] keyfile = /etc/keystone/ssl/private/keystonekey.pem	(StrOpt) Path of the keyfile for SSL.
[federation] federated_domain_name = Federated	(StrOpt) A 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. Changing this option to empty string or None will not have any impact and default name will be used.
[federation] remote_id_attribute = None	(StrOpt) Value to be used to obtain the entity ID of the Identity Provider from the environment (e.g. if using the mod_shib plugin this value is `Shib-Identity-Provider`).
[federation] sso_callback_template = /etc/keystone/sso_callback_template.html	(StrOpt) Location of Single Sign-On callback handler, will return a token to a trusted dashboard host.
[federation] trusted_dashboard = []	(MultiStrOpt) 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 the trusted_dashboard list. This configuration option may be repeated for multiple values. For example: trusted_dashboard=http://acme.com trusted_dashboard=http://beta.com
[fernet_tokens] key_repository = /etc/keystone/fernet-keys/	(StrOpt) Directory containing Fernet token keys.
[fernet_tokens] max_active_keys = 3	(IntOpt) This controls how many keys are held in rotation by keystone-manage fernet_rotate before they are discarded. The default value of 3 means that keystone will maintain one staged key, one primary key, and one secondary key. Increasing this value means that additional secondary keys will be kept in the rotation.
[identity] cache_time = 600	(IntOpt) Time to cache identity data (in seconds). This has no effect unless global and identity caching are enabled.
[identity] caching = True	(BoolOpt) Toggle for identity caching. This has no effect unless global caching is enabled.
[identity] domain_configurations_from_database = False	(BoolOpt) Extract the domain specific configuration options from the resource backend where they have been stored with the domain data. This feature is disabled by default (in which case the domain specific options will be loaded from files in the domain configuration directory); set to true to enable.
[oslo_messaging_amqp] allow_insecure_clients = False	(BoolOpt) Accept clients using either SSL or plain TCP
[oslo_messaging_amqp] broadcast_prefix = broadcast	(StrOpt) address prefix used when broadcasting to all servers
[oslo_messaging_amqp] container_name = None	(StrOpt) Name for the AMQP container
[oslo_messaging_amqp] group_request_prefix = unicast	(StrOpt) address prefix when sending to any server in group
[oslo_messaging_amqp] idle_timeout = 0	(IntOpt) Timeout for inactive connections (in seconds)
[oslo_messaging_amqp] server_request_prefix = exclusive	(StrOpt) address prefix used when sending to a specific server
[oslo_messaging_amqp] ssl_ca_file =	(StrOpt) CA certificate PEM file to verify server certificate
[oslo_messaging_amqp] ssl_cert_file =	(StrOpt) Identifying certificate PEM file to present to clients

Option = default value	(Type) Help string
[oslo_messaging_amqp] ssl_key_file =	(StrOpt) Private key PEM file used to sign cert_file certificate
[oslo_messaging_amqp] ssl_key_password = None	(StrOpt) Password for decrypting ssl_key_file (if encrypted)
[oslo_messaging_amqp] trace = False	(BoolOpt) Debug: dump AMQP frames to stdout
[oslo_messaging_qpid] amqp_auto_delete = False	(BoolOpt) Auto-delete queues in AMQP.
[oslo_messaging_qpid] amqp_durable_queues = False	(BoolOpt) Use durable queues in AMQP.
[oslo_messaging_qpid] qpid_heartbeat = 60	(IntOpt) Seconds between connection keepalive heartbeats.
[oslo_messaging_qpid] qpid_hostname = localhost	(StrOpt) Qpid broker hostname.
[oslo_messaging_qpid] qpid_hosts = \$qpid_hostname: \$qpid_port	(ListOpt) Qpid HA cluster host:port pairs.
[oslo_messaging_qpid] qpid_password =	(StrOpt) Password for Qpid connection.
[oslo_messaging_qpid] qpid_port = 5672	(IntOpt) Qpid broker port.
[oslo_messaging_qpid] qpid_protocol = tcp	(StrOpt) Transport to use, either 'tcp' or 'ssl'.
[oslo_messaging_qpid] qpid_receiver_capacity = 1	(IntOpt) The number of prefetched messages held by receiver.
[oslo_messaging_qpid] qpid_sasl_mechanisms =	(StrOpt) Space separated list of SASL mechanisms to use for auth.
[oslo_messaging_qpid] qpid_tcp_nodelay = True	(BoolOpt) Whether to disable the Nagle algorithm.
[oslo_messaging_qpid] qpid_topology_version = 1	(IntOpt) The qpid topology version to use. Version 1 is what was originally used by impl_qpid. Version 2 includes some backwards-incompatible changes that allow broker federation to work. Users should update to version 2 when they are able to take everything down, as it requires a clean break.
[oslo_messaging_qpid] qpid_username =	(StrOpt) Username for Qpid connection.
[oslo_messaging_qpid] rpc_conn_pool_size = 30	(IntOpt) Size of RPC connection pool.
[oslo_messaging_rabbit] amqp_auto_delete = False	(BoolOpt) Auto-delete queues in AMQP.
[oslo_messaging_rabbit] amqp_durable_queues = False	(BoolOpt) Use durable queues in AMQP.
[oslo_messaging_rabbit] fake_rabbit = False	(BoolOpt) Deprecated, use rpc_backend=kombu+memory or rpc_backend=fake
[oslo_messaging_rabbit] heartbeat_rate = 2	(IntOpt) How often times during the heartbeat_timeout_threshold we check the heartbeat.
[oslo_messaging_rabbit] heartbeat_timeout_threshold = 0	(IntOpt) Number of seconds after which the Rabbit broker is considered down if heartbeat's keep-alive fails (0 disables the heartbeat, >0 enables it. Enabling heartbeats requires kombu>=3.0.7 and amqp>=1.4.0). EXPERIMENTAL
[oslo_messaging_rabbit] kombu_reconnect_delay = 1.0	(FloatOpt) How long to wait before reconnecting in response to an AMQP consumer cancel notification.
[oslo_messaging_rabbit] kombu_ssl_ca_certs =	(StrOpt) SSL certification authority file (valid only if SSL enabled).
[oslo_messaging_rabbit] kombu_ssl_certfile =	(StrOpt) SSL cert file (valid only if SSL enabled).
[oslo_messaging_rabbit] kombu_ssl_keyfile =	(StrOpt) SSL key file (valid only if SSL enabled).
[oslo_messaging_rabbit] kombu_ssl_version =	(StrOpt) 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.
[oslo_messaging_rabbit] rabbit_ha_queues = False	(BoolOpt) Use HA queues in RabbitMQ (x-ha-policy: all). If you change this option, you must wipe the RabbitMQ database.
[oslo_messaging_rabbit] rabbit_host = localhost	(StrOpt) The RabbitMQ broker address where a single node is used.

Option = default value	(Type) Help string
[oslo_messaging_rabbit] rabbit_hosts = \$rabbit_host:\$rabbit_port	(ListOpt) RabbitMQ HA cluster host:port pairs.
[oslo_messaging_rabbit] rabbit_login_method = AMQ-PLAIN	(StrOpt) The RabbitMQ login method.
[oslo_messaging_rabbit] rabbit_max_retries = 0	(IntOpt) Maximum number of RabbitMQ connection retries. Default is 0 (infinite retry count).
[oslo_messaging_rabbit] rabbit_password = guest	(StrOpt) The RabbitMQ password.
[oslo_messaging_rabbit] rabbit_port = 5672	(IntOpt) The RabbitMQ broker port where a single node is used.
[oslo_messaging_rabbit] rabbit_retry_backoff = 2	(IntOpt) How long to backoff for between retries when connecting to RabbitMQ.
[oslo_messaging_rabbit] rabbit_retry_interval = 1	(IntOpt) How frequently to retry connecting with RabbitMQ.
[oslo_messaging_rabbit] rabbit_use_ssl = False	(BoolOpt) Connect over SSL for RabbitMQ.
[oslo_messaging_rabbit] rabbit_userid = guest	(StrOpt) The RabbitMQ userid.
[oslo_messaging_rabbit] rabbit_virtual_host = /	(StrOpt) The RabbitMQ virtual host.
[oslo_messaging_rabbit] rpc_conn_pool_size = 30	(IntOpt) Size of RPC connection pool.
[oslo_middleware] max_request_body_size = 114688	(IntOpt) The maximum body size for each request, in bytes.
[oslo_policy] policy_default_rule = default	(StrOpt) Default rule. Enforced when a requested rule is not found.
[oslo_policy] policy_dirs = ['policy.d']	(MultiStrOpt) 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.
[oslo_policy] policy_file = policy.json	(StrOpt) The JSON file that defines policies.
[resource] cache_time = None	(IntOpt) TTL (in seconds) to cache resource data. This has no effect unless global caching is enabled.
[resource] caching = True	(BoolOpt) Toggle for resource caching. This has no effect unless global caching is enabled.
[resource] driver = None	(StrOpt) Resource backend driver. If a resource driver is not specified, the assignment driver will choose the resource driver.
[resource] list_limit = None	(IntOpt) Maximum number of entities that will be returned in a resource collection.
[revoke] cache_time = 3600	(IntOpt) Time to cache the revocation list and the revocation events (in seconds). This has no effect unless global and token caching are enabled.
[role] cache_time = None	(IntOpt) TTL (in seconds) to cache role data. This has no effect unless global caching is enabled.
[role] caching = True	(BoolOpt) Toggle for role caching. This has no effect unless global caching is enabled.
[role] driver = None	(StrOpt) Role backend driver.
[role] list_limit = None	(IntOpt) Maximum number of entities that will be returned in a role collection.
[saml] relay_state_prefix = ss:mem:	(StrOpt) The prefix to use for the RelayState SAML attribute, used when generating ECP wrapped assertions.
[token] allow_rescope_scoped_token = True	(BoolOpt) Allow rescopeing of scoped token. Setting allow_rescope_scoped_token to false prevents a user from exchanging a scoped token for any other token.
[trust] allow_redelegation = False	(BoolOpt) Enable redelegation feature.







































Configuration option = Default value	Description
<code>qpid_receiver_capacity = 1</code>	(IntOpt) The number of prefetched messages held by receiver.
<code>qpid_sasl_mechanisms =</code>	(StrOpt) Space separated list of SASL mechanisms to use for auth.
<code>qpid_tcp_nodelay = True</code>	(BoolOpt) Whether to disable the Nagle algorithm.
<code>qpid_topology_version = 1</code>	(IntOpt) The qpid topology version to use. Version 1 is what was originally used by <code>impl_qpid</code> . Version 2 includes some backwards-incompatible changes that allow broker federation to work. Users should update to version 2 when they are able to take everything down, as it requires a clean break.
<code>qpid_username =</code>	(StrOpt) Username for Qpid connection.
<code>rpc_conn_pool_size = 30</code>	(IntOpt) Size of RPC connection pool.

## Support for ISO images

You can load ISO images into the Image service. You can subsequently boot an ISO image using Compute.

### Procedure 8.1. To load an ISO image to an Image service data store

1. In the Image service, run the following command:

```
$ glance image-create --name "ubuntu-14.04.2-server-amd64.iso" \
--copy-from http://releases.ubuntu.com/14.04.2/ubuntu-14.04.2-server-
amd64.iso \
--is-public True --container-format bare --disk-format iso
```

In this command, `ubuntu.iso` is the name for the ISO image after it is loaded to the Image service, and `ubuntu-14.04.2-server-amd64.iso` is the name of the source ISO image.

2. Optionally, to confirm the upload in Image Service (glance), run this command:

Run this command:

```
$ glance image-list
```

## Configure back ends

The Image service supports several back ends for storing virtual machine images:

- OpenStack Block Storage (cinder)
- A directory on a local file system
- GridFS
- Ceph RBD
- Amazon S3
- Sheepdog
- OpenStack Object Storage (swift)

- VMware ESX

The following tables detail the options available for each.

**Table 8.23. Description of cinder configuration options**

Configuration option = Default value	Description
[glance_store]	
<code>cinder_api_insecure = False</code>	(BoolOpt) Allow to perform insecure SSL requests to cinder
<code>cinder_ca_certificates_file = None</code>	(StrOpt) Location of ca certificates file to use for cinder client requests.
<code>cinder_catalog_info = volume:cinder:publicURL</code>	(StrOpt) Info to match when looking for cinder in the service catalog. Format is : separated values of the form: <service_type>:<service_name>:<endpoint_type>
<code>cinder_endpoint_template = None</code>	(StrOpt) Override service catalog lookup with template for cinder endpoint e.g. http://localhost:8776/v1/%(project_id)s
<code>cinder_http_retries = 3</code>	(IntOpt) Number of cinderclient retries on failed http calls

**Table 8.24. Description of filesystem configuration options**

Configuration option = Default value	Description
[glance_store]	
<code>filesystem_store_datadir = None</code>	(StrOpt) Directory to which the Filesystem backend store writes images.
<code>filesystem_store_datadirs = None</code>	(MultiStrOpt) List of directories and its priorities to which the Filesystem backend store writes images.
<code>filesystem_store_file_perm = 0</code>	(IntOpt) The required permission for created image file. In this way the user other service used, e.g. Nova, who consumes the image could be the exclusive member of the group that owns the files created. Assigning it less than or equal to zero means don't change the default permission of the file. This value will be decoded as an octal digit.
<code>filesystem_store_metadata_file = None</code>	(StrOpt) The path to a file which contains the metadata to be returned with any location associated with this store. The file must contain a valid JSON object. The object should contain the keys 'id' and 'mountpoint'. The value for both keys should be 'string'.

**Table 8.25. Description of GridFS configuration options**

Configuration option = Default value	Description
[glance_store]	
<code>mongodb_store_db = None</code>	(StrOpt) Database to use
<code>mongodb_store_uri = None</code>	(StrOpt) Hostname or IP address of the instance to connect to, or a mongodb URI, or a list of hostnames / mongodb URIs. If host is an IPv6 literal it must be enclosed in '[' and ']' characters following the RFC2732 URL syntax (e.g. '['::1']' for localhost)

**Table 8.26. Description of RBD configuration options**

Configuration option = Default value	Description
[glance_store]	
<code>rbd_store_ceph_conf = /etc/ceph/ceph.conf</code>	(StrOpt) Ceph configuration file path. If <None>, librados will locate the default config. If using cephx authen-

Configuration option = Default value	Description
	tion, this file should include a reference to the right keyring in a client.<USER> section
<code>rbd_store_chunk_size = 8</code>	(IntOpt) RADOS images will be chunked into objects of this size (in megabytes). For best performance, this should be a power of two.
<code>rbd_store_pool = localhost</code>	(StrOpt) RADOS pool in which images are stored.
<code>rbd_store_user = None</code>	(StrOpt) RADOS user to authenticate as (only applicable if using Cephx. If <None>, a default will be chosen based on the client. section in <code>rbd_store_ceph_conf</code> )

**Table 8.27. Description of S3 configuration options**

Configuration option = Default value	Description
[glance_store]	
<code>s3_store_access_key = None</code>	(StrOpt) The S3 query token access key.
<code>s3_store_bucket = None</code>	(StrOpt) The S3 bucket to be used to store the Glance data.
<code>s3_store_bucket_url_format = subdomain</code>	(StrOpt) The S3 calling format used to determine the bucket. Either subdomain or path can be used.
<code>s3_store_create_bucket_on_put = False</code>	(BoolOpt) A boolean to determine if the S3 bucket should be created on upload if it does not exist or if an error should be returned to the user.
<code>s3_store_host = None</code>	(StrOpt) The host where the S3 server is listening.
<code>s3_store_large_object_chunk_size = 10</code>	(IntOpt) What multipart upload part size, in MB, should S3 use when uploading parts. The size must be greater than or equal to 5M.
<code>s3_store_large_object_size = 100</code>	(IntOpt) What size, in MB, should S3 start chunking image files and do a multipart upload in S3.
<code>s3_store_object_buffer_dir = None</code>	(StrOpt) The local directory where uploads will be staged before they are transferred into S3.
<code>s3_store_secret_key = None</code>	(StrOpt) The S3 query token secret key.
<code>s3_store_thread_pools = 10</code>	(IntOpt) The number of thread pools to perform a multipart upload in S3.

**Table 8.28. Description of Sheepdog configuration options**

Configuration option = Default value	Description
[glance_store]	
<code>sheepdog_store_address = localhost</code>	(StrOpt) IP address of sheep daemon.
<code>sheepdog_store_chunk_size = 64</code>	(IntOpt) Images will be chunked into objects of this size (in megabytes). For best performance, this should be a power of two.
<code>sheepdog_store_port = 7000</code>	(IntOpt) Port of sheep daemon.

**Table 8.29. Description of swift configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>default_swift_reference = ref1</code>	(StrOpt) The reference to the default swift account/backing store parameters to use for adding new images.
<code>swift_store_auth_address = None</code>	(StrOpt) The address where the Swift authentication service is listening.(deprecated)
<code>swift_store_config_file = None</code>	(StrOpt) The config file that has the swift account(s)configs.

Configuration option = Default value	Description
swift_store_key = None	(StrOpt) Auth key for the user authenticating against the Swift authentication service. (deprecated)
swift_store_user = None	(StrOpt) The user to authenticate against the Swift authentication service (deprecated)
<b>[glance_store]</b>	
default_swift_reference = <i>ref1</i>	(StrOpt) The reference to the default swift account/backup store parameters to use for adding new images.
swift_store_admin_tenants =	(ListOpt) A list of tenants that will be granted read/write access on all Swift containers created by Glance in multi-tenant mode.
swift_store_auth_address = None	(StrOpt) The address where the Swift authentication service is listening.(deprecated)
swift_store_auth_insecure = False	(BoolOpt) If True, swiftclient won't check for a valid SSL certificate when authenticating.
swift_store_auth_version = 2	(StrOpt) Version of the authentication service to use. Valid versions are 2 for keystone and 1 for swauth and rackspace. (deprecated)
swift_store_cacert = None	(StrOpt) A string giving the CA certificate file to use in SSL connections for verifying certs.
swift_store_config_file = None	(StrOpt) The config file that has the swift account(s)configs.
swift_store_container = <i>glance</i>	(StrOpt) Container within the account that the account should use for storing images in Swift when using single container mode. In multiple container mode, this will be the prefix for all containers.
swift_store_create_container_on_put = False	(BoolOpt) A boolean value that determines if we create the container if it does not exist.
swift_store_endpoint = None	(StrOpt) If set, the configured endpoint will be used. If None, the storage url from the auth response will be used.
swift_store_endpoint_type = <i>publicURL</i>	(StrOpt) A string giving the endpoint type of the swift service to use (publicURL, adminURL or internalURL). This setting is only used if swift_store_auth_version is 2.
swift_store_key = None	(StrOpt) Auth key for the user authenticating against the Swift authentication service. (deprecated)
swift_store_large_object_chunk_size = 200	(IntOpt) The amount of data written to a temporary disk buffer during the process of chunking the image file.
swift_store_large_object_size = 5120	(IntOpt) The size, in MB, that Glance will start chunking image files and do a large object manifest in Swift.
swift_store_multi_tenant = False	(BoolOpt) If set to True, enables multi-tenant storage mode which causes Glance images to be stored in tenant specific Swift accounts.
swift_store_multiple_containers_seed = 0	(IntOpt) 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.Used only when swift_store_multi_tenant is disabled. The total number of containers that will be used is equal to 16^N, so if this config option is set to 2, then 16^2=256 containers will be used to store images.
swift_store_region = None	(StrOpt) The region of the swift endpoint to be used for single tenant. This setting is only necessary if the tenant has multiple swift endpoints.
swift_store_retry_get_count = 0	(IntOpt) The number of times a Swift download will be retried before the request fails.

Configuration option = Default value	Description
<code>swift_store_service_type = object-store</code>	(StrOpt) A string giving the service type of the swift service to use. This setting is only used if <code>swift_store_auth_version</code> is 2.
<code>swift_store_ssl_compression = True</code>	(BoolOpt) If set to False, disables SSL layer compression of https swift requests. Setting to False may improve performance for images which are already in a compressed format, eg qcow2.
<code>swift_store_user = None</code>	(StrOpt) The user to authenticate against the Swift authentication service (deprecated)

## Configure vCenter data stores for the Image service back end

To use vCenter data stores for the Image service back end, you must update the `glance-api.conf` file, as follows:

- Add data store parameters to the VMware Datastore Store Options section.
- Specify vSphere as the back end.



### Note

You must configure any configured Image service data stores for the Compute service.

You can specify vCenter data stores directly by using the data store name or Storage Policy Based Management (SPBM), which requires vCenter Server 5.5 or later. For details, see [the section called “Configure vCenter data stores for the back end” \[479\]](#).



### Note

If you intend to use multiple data stores for the back end, use the SPBM feature.

In the `glance_store` section, set the `default_store` option to **vsphere**, as shown in this code sample:

```
[glance_store]
# Which back end scheme should Glance use by default is not specified
# in a request to add a new image to Glance? Known schemes are determined
# by the known_stores option below.
# Default: 'file'
default_store = vsphere
```

The following table describes the parameters in the VMware Datastore Store Options section:

**Table 8.30. Description of VMware configuration options**

Configuration option = Default value	Description
<code>[glance_store]</code>	
<code>vmware_api_insecure = False</code>	(BoolOpt) Allow to perform insecure SSL requests to ESX/VC.

Configuration option = Default value	Description
<code>vmware_api_retry_count = 10</code>	(IntOpt) Number of times VMware ESX/VC server API must be retried upon connection related issues.
<code>vmware_datacenter_path = ha-datacenter</code>	(StrOpt) DEPRECATED. Inventory path to a datacenter. If the <code>vmware_server_host</code> specified is an ESX/ESXi, the <code>vmware_datacenter_path</code> is optional. If specified, it should be "ha-datacenter". This option is deprecated in favor of <code>vmware_datastores</code> and will be removed in the Liberty release.
<code>vmware_datastore_name = None</code>	(StrOpt) DEPRECATED. Datastore associated with the datacenter. This option is deprecated in favor of <code>vmware_datastores</code> and will be removed in the Liberty release.
<code>vmware_datastores = None</code>	(MultiStrOpt) A list of datastores where the image can be stored. This option may be specified multiple times for specifying multiple datastores. Either one of <code>vmware_datastore_name</code> or <code>vmware_datastores</code> is required. 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 ":". Thus, the required format becomes <code>&lt;datacenter_path&gt;:&lt;datastore_name&gt;:&lt;optional_weight&gt;</code> . 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.
<code>vmware_server_host = None</code>	(StrOpt) ESX/ESXi or vCenter Server target system. The server value can be an IP address or a DNS name.
<code>vmware_server_password = None</code>	(StrOpt) Password for authenticating with VMware ESX/VC server.
<code>vmware_server_username = None</code>	(StrOpt) Username for authenticating with VMware ESX/VC server.
<code>vmware_store_image_dir = /openstack_glance</code>	(StrOpt) The name of the directory where the glance images will be stored in the VMware datastore.
<code>vmware_task_poll_interval = 5</code>	(IntOpt) The interval used for polling remote tasks invoked on VMware ESX/VC server.

The following block of text shows a sample configuration:

```
# ===== VMware Datastore Store Options =====
# ESX/ESXi or vCenter Server target system.
# The server value can be an IP address or a DNS name
# e.g. 127.0.0.1, 127.0.0.1:443, www.vmware-infra.com
vmware_server_host = 192.168.0.10

# Server username (string value)
vmware_server_username = ADMINISTRATOR

# Server password (string value)
vmware_server_password = password

# Inventory path to a datacenter (string value)
# Value optional when vmware_server_ip is an ESX/ESXi host: if specified
# should be `ha-datacenter`.
vmware_datacenter_path = DATACENTER

# Datastore associated with the datacenter (string value)
```

```
vmware_datastore_name = datastore1

# PBM service WSDL file location URL. e.g.
# file:///opt/SDK/spbm/wsd/pbmService.wsdl Not setting this
# will disable storage policy based placement of images.
# (string value)
#vmware_pbm_wsd_location =

# The PBM policy. If `pbm_wsd_location` is set, a PBM policy needs
# to be specified. This policy will be used to select the datastore
# in which the images will be stored.
#vmware_pbm_policy =

# The interval used for polling remote tasks
# invoked on VMware ESX/VC server in seconds (integer value)
vmware_task_poll_interval = 5

# Absolute path of the folder containing the images in the datastore
# (string value)
vmware_store_image_dir = /openstack_glance

# Allow to perform insecure SSL requests to the target system (boolean value)
vmware_api_insecure = False
```

## Configure vCenter data stores for the back end

You can specify a vCenter data store for the back end by setting the `vmware_datastore_name` parameter value to the vCenter name of the data store. This configuration limits the back end to a single data store.

### Procedure 8.2. To configure a single data store

1. If present, comment or delete the `vmware_pbm_wsd_location` and `vmware_pbm_policy` parameters.
2. Uncomment and define the `vmware_datastore_name` parameter with the name of the vCenter data store.
3. Complete the other vCenter configuration parameters as appropriate.

## Image service sample configuration files

You can find the files that are described in this section in the `/etc/glance/` directory.

### glance-api.conf

The configuration file for the Image service API is found in the `glance-api.conf` file.

This file must be modified after installation.

```
[DEFAULT]
# Show more verbose log output (sets INFO log level output)
#verbose = False

# Show debugging output in logs (sets DEBUG log level output)
#debug = False
```

```

# Maximum image size (in bytes) that may be uploaded through the
# Glance API server. Defaults to 1 TB.
# WARNING: this value should only be increased after careful consideration
# and must be set to a value under 8 EB (9223372036854775808).
#image_size_cap = 1099511627776

# Address to bind the API server
bind_host = 0.0.0.0

# Port the bind the API server to
bind_port = 9292

# Log to this file. Make sure you do not set the same log file for both the
# API
# and registry servers!
#
# If `log_file` is omitted and `use_syslog` is false, then log messages are
# sent to stdout as a fallback.
log_file = /var/log/glance/api.log

# Backlog requests when creating socket
backlog = 4096

# TCP_KEEPIDLE value in seconds when creating socket.
# Not supported on OS X.
#tcp_keepidle = 600

# Timeout (in seconds) for client connections' socket operations. If an
# incoming
# connection is idle for this period it will be closed. A value of "0"
# means wait forever.
#client_socket_timeout = 0

# API to use for accessing data. Default value points to sqlalchemy
# package, it is also possible to use: glance.db.registry.api
# data_api = glance.db.sqlalchemy.api

# The number of child process workers that will be
# created to service API requests. The default will be
# equal to the number of CPUs available. (integer value)
#workers = 4

# 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)
# max_header_line = 16384

# Role used to identify an authenticated user as administrator
#admin_role = admin

# Allow unauthenticated users to access the API with read-only
# privileges. This only applies when using ContextMiddleware.
#allow_anonymous_access = False

# Allow access to version 1 of glance api
#enable_v1_api = True

# Allow access to version 2 of glance api
#enable_v2_api = True

```













```

# Maximum number of SQL connections to keep open in a pool
# (integer value)
# Deprecated group/name - [DEFAULT]/sql_max_pool_size
# Deprecated group/name - [DATABASE]/sql_max_pool_size
#max_pool_size = <None>

# Maximum db connection retries during startup. (setting -1
# implies an infinite retry count) (integer value)
# Deprecated group/name - [DEFAULT]/sql_max_retries
# Deprecated group/name - [DATABASE]/sql_max_retries
#max_retries = 10

# Interval between retries of opening a sql connection
# (integer value)
# Deprecated group/name - [DEFAULT]/sql_retry_interval
# Deprecated group/name - [DATABASE]/reconnect_interval
#retry_interval = 10

# If set, use this value for max_overflow with sqlalchemy
# (integer value)
# Deprecated group/name - [DEFAULT]/sql_max_overflow
# Deprecated group/name - [DATABASE]/sqlalchemy_max_overflow
#max_overflow = <None>

# Verbosity of SQL debugging information. 0=None,
# 100=Everything (integer value)
# Deprecated group/name - [DEFAULT]/sql_connection_debug
#connection_debug = 0

# Add python stack traces to SQL as comment strings (boolean
# value)
# Deprecated group/name - [DEFAULT]/sql_connection_trace
#connection_trace = False

# If set, use this value for pool_timeout with sqlalchemy
# (integer value)
# Deprecated group/name - [DATABASE]/sqlalchemy_pool_timeout
#pool_timeout = <None>

# Enable the experimental use of database reconnect on
# connection lost (boolean value)
#use_db_reconnect = False

# seconds between db connection retries (integer value)
#db_retry_interval = 1

# Whether to increase interval between db connection retries,
# up to db_max_retry_interval (boolean value)
#db_inc_retry_interval = True

# max seconds between db connection retries, if
# db_inc_retry_interval is enabled (integer value)
#db_max_retry_interval = 10

# maximum db connection retries before error is raised.
# (setting -1 implies an infinite retry count) (integer value)
#db_max_retries = 20

[oslo_concurrency]
    
```







```
# When adding an image, the highest priority directory will be selected,
unless
# there is not enough space available in cases where the image size is already
# known. If no priority is given, it is assumed to be zero and the directory
# will be considered for selection last. If multiple directories have the same
# priority, then the one with the most free space available is selected.
# If same store is specified multiple times then BadStoreConfiguration
# exception will be raised.
#filesystem_store_datadirs = /var/lib/glance/images/:1

# A path to a JSON file that contains metadata describing the storage
# system. When show_multiple_locations is True the information in this
# file will be returned with any location that is contained in this
# store.
#filesystem_store_metadata_file = None

# ===== Swift Store Options =====

# Version of the authentication service to use
# Valid versions are '2' for keystone and '1' for swauth and rackspace
swift_store_auth_version = 2

# Address where the Swift authentication service lives
# Valid schemes are 'http://' and 'https://'
# If no scheme specified, default to 'https://'
# For swauth, use something like '127.0.0.1:8080/v1.0/'
swift_store_auth_address = 127.0.0.1:5000/v2.0/

# User to authenticate against the Swift authentication service
# If you use Swift authentication service, set it to 'account': 'user'
# where 'account' is a Swift storage account and 'user'
# is a user in that account
swift_store_user = jdoe:jdoe

# Auth key for the user authenticating against the
# Swift authentication service
swift_store_key = a86850deb2742ec3cb41518e26aa2d89

# Container within the account that the account should use
# for storing images in Swift
swift_store_container = glance

# Do we create the container if it does not exist?
swift_store_create_container_on_put = False

# What size, in MB, should Glance start chunking image files
# and do a large object manifest in Swift? By default, this is
# the maximum object size in Swift, which is 5GB
swift_store_large_object_size = 5120

# swift_store_config_file = glance-swift.conf
# This file contains references for each of the configured
# Swift accounts/backing stores. If used, this option can prevent
# credentials being stored in the database. Using Swift references
# is disabled if this config is left blank.

# The reference to the default Swift parameters to use for adding new images.
# default_swift_reference = 'ref1'

# When doing a large object manifest, what size, in MB, should
```

```
# Glance write chunks to Swift? This amount of data is written
# to a temporary disk buffer during the process of chunking
# the image file, and the default is 200MB
swift_store_large_object_chunk_size = 200

# If set, the configured endpoint will be used. If None, the storage URL
# from the auth response will be used. The location of an object is
# obtained by appending the container and object to the configured URL.
#
# swift_store_endpoint = https://www.example.com/v1/not_a_container
#swift_store_endpoint =

# If set to True enables multi-tenant storage mode which causes Glance images
# to be stored in tenant specific Swift accounts.
#swift_store_multi_tenant = False

# If set to an integer value between 1 and 32, a single-tenant store will
# use multiple containers to store images. If set to the default value of 0,
# only a single container will be used. Multi-tenant stores are not affected
# by this option. The max number of containers that will be used to store
# images is approximately 16^N where N is the value of this option. Discuss
# the impact of this with your swift deployment team, as this option is only
# beneficial in the largest of deployments where swift rate limiting can lead
# to unwanted throttling on a single container.
#swift_store_multiple_containers_seed = 0

# A list of swift ACL strings that will be applied as both read and
# write ACLs to the containers created by Glance in multi-tenant
# mode. This grants the specified tenants/users read and write access
# to all newly created image objects. The standard swift ACL string
# formats are allowed, including:
# <tenant_id>:<username>
# <tenant_name>:<username>
# *:<username>
# Multiple ACLs can be combined using a comma separated list, for
# example: swift_store_admin_tenants = service:glance,*:admin
#swift_store_admin_tenants =

# The region of the swift endpoint to be used for single tenant. This setting
# is only necessary if the tenant has multiple swift endpoints.
#swift_store_region =

# If set to False, disables SSL layer compression of https swift requests.
# Setting to 'False' may improve performance for images which are already
# in a compressed format, eg qcow2. If set to True, enables SSL layer
# compression (provided it is supported by the target swift proxy).
#swift_store_ssl_compression = True

# The number of times a Swift download will be retried before the
# request fails
#swift_store_retry_get_count = 0

# Bypass SSL verification for Swift
#swift_store_auth_insecure = False

# The path to a CA certificate bundle file to use for SSL verification when
# communicating with Swift.
#swift_store_cacert =

# ===== S3 Store Options =====
```

```

# Address where the S3 authentication service lives
# Valid schemes are 'http://' and 'https://'
# If no scheme specified, default to 'http://'
s3_store_host = s3.amazonaws.com

# User to authenticate against the S3 authentication service
s3_store_access_key = <20-char AWS access key>

# Auth key for the user authenticating against the
# S3 authentication service
s3_store_secret_key = <40-char AWS secret key>

# Container within the account that the account should use
# for storing images in S3. Note that S3 has a flat namespace,
# so you need a unique bucket name for your glance images. An
# easy way to do this is append your AWS access key to "glance".
# S3 buckets in AWS *must* be lowercased, so remember to lowercase
# your AWS access key if you use it in your bucket name below!
s3_store_bucket = <lowercased 20-char aws access key>glance

# Do we create the bucket if it does not exist?
s3_store_create_bucket_on_put = False

# When sending images to S3, the data will first be written to a
# temporary buffer on disk. By default the platform's temporary directory
# will be used. If required, an alternative directory can be specified here.
#s3_store_object_buffer_dir = /path/to/dir

# When forming a bucket url, boto will either set the bucket name as the
# subdomain or as the first token of the path. Amazon's S3 service will
# accept it as the subdomain, but Swift's S3 middleware requires it be
# in the path. Set this to 'path' or 'subdomain' - defaults to 'subdomain'.
#s3_store_bucket_url_format = subdomain

# Size, in MB, should S3 start chunking image files
# and do a multipart upload in S3. The default is 100MB.
#s3_store_large_object_size = 100

# Multipart upload part size, in MB, should S3 use when uploading
# parts. The size must be greater than or equal to
# 5MB. The default is 10MB.
#s3_store_large_object_chunk_size = 10

# The number of thread pools to perform a multipart upload
# in S3. The default is 10.
#s3_store_thread_pools = 10

# ===== RBD Store Options =====

# Ceph configuration file path
# If using cephx authentication, this file should
# include a reference to the right keyring
# in a client.<USER> section
#rbd_store_ceph_conf = /etc/ceph/ceph.conf

# RADOS user to authenticate as (only applicable if using cephx)
# If <None>, a default will be chosen based on the client. section
# in rbd_store_ceph_conf
#rbd_store_user = <None>
    
```

```
# RADOS pool in which images are stored
#rbd_store_pool = images

# RADOS images will be chunked into objects of this size (in megabytes).
# For best performance, this should be a power of two
#rbd_store_chunk_size = 8

# ===== Sheepdog Store Options =====

sheepdog_store_address = localhost

sheepdog_store_port = 7000

# Images will be chunked into objects of this size (in megabytes).
# For best performance, this should be a power of two
sheepdog_store_chunk_size = 64

# ===== Cinder Store Options =====

# Info to match when looking for cinder in the service catalog
# Format is : separated values of the form:
# <service_type>:<service_name>:<endpoint_type> (string value)
#cinder_catalog_info = volume:cinder:publicURL

# Override service catalog lookup with template for cinder endpoint
# e.g. http://localhost:8776/v1/%(project_id)s (string value)
#cinder_endpoint_template = <None>

# Region name of this node (string value)
#os_region_name = <None>

# Location of ca certificates file to use for cinder client requests
# (string value)
#cinder_ca_certificates_file = <None>

# Number of cinderclient retries on failed http calls (integer value)
#cinder_http_retries = 3

# Allow to perform insecure SSL requests to cinder (boolean value)
#cinder_api_insecure = False

# ===== VMware Datastore Store Options =====

# ESX/ESXi or vCenter Server target system.
# The server value can be an IP address or a DNS name
# e.g. 127.0.0.1, 127.0.0.1:443, www.vmware-infra.com
#vmware_server_host = <None>

# Server username (string value)
#vmware_server_username = <None>

# Server password (string value)
#vmware_server_password = <None>

# Inventory path to a datacenter (string value)
# Value optional when vmware_server_ip is an ESX/ESXi host: if specified
# should be `ha-datacenter`.
# Deprecated in favor of vmware_datastores.
#vmware_datacenter_path = <None>
```

```

# Datastore associated with the datacenter (string value)
# Deprecated in favor of vmware_datastores.
#vmware_datastore_name = <None>

# A list of datastores where the image can be stored.
# This option may be specified multiple times for specifying multiple
# datastores. Either one of vmware_datastore_name or vmware_datastores is
# required. 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 ":". Thus, the required format becomes
# <datacenter_path>:<datastore_name>:<optional_weight>.
# 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.
#vmware_datastores = <None>

# The number of times we retry on failures
# e.g., socket error, etc (integer value)
#vmware_api_retry_count = 10

# The interval used for polling remote tasks
# invoked on VMware ESX/VC server in seconds (integer value)
#vmware_task_poll_interval = 5

# Absolute path of the folder containing the images in the datastore
# (string value)
#vmware_store_image_dir = /openstack_glance

# Allow to perform insecure SSL requests to the target system (boolean value)
#vmware_api_insecure = False

```

## glance-registry.conf

Configuration for the Image service's registry, which stores the metadata about images, is found in the `glance-registry.conf` file.

This file must be modified after installation.

```

[DEFAULT]
# Show more verbose log output (sets INFO log level output)
#verbose = False

# Show debugging output in logs (sets DEBUG log level output)
#debug = False

# Address to bind the registry server
bind_host = 0.0.0.0

# Port the bind the registry server to
bind_port = 9191

# Log to this file. Make sure you do not set the same log file for both the
# API
# and registry servers!
#

```

```
# If `log_file` is omitted and `use_syslog` is false, then log messages are
# sent to stdout as a fallback.
log_file = /var/log/glance/registry.log

# Backlog requests when creating socket
backlog = 4096

# TCP_KEEPIDLE value in seconds when creating socket.
# Not supported on OS X.
#tcp_keepidle = 600

# Timeout (in seconds) for client connections' socket operations. If an
# incoming
# connection is idle for this period it will be closed. A value of "0"
# means wait forever.
#client_socket_timeout = 0

# API to use for accessing data. Default value points to sqlalchemy
# package.
#data_api = glance.db.sqlalchemy.api

# The number of child process workers that will be
# created to service Registry requests. The default will be
# equal to the number of CPUs available. (integer value)
#workers = None

# Enable Registry API versions individually or simultaneously
#enable_v1_registry = True
#enable_v2_registry = True

# Limit the api to return `param_limit_max` items in a call to a container. If
# a larger `limit` query param is provided, it will be reduced to this value.
api_limit_max = 1000

# If a `limit` query param is not provided in an api request, it will
# default to `limit_param_default`
limit_param_default = 25

# Role used to identify an authenticated user as administrator
#admin_role = admin

# Enable DEBUG log messages from sqlalchemy which prints every database
# query and response.
# Default: False
#sqlalchemy_debug = True

# http_keepalive option. If False, server will return the header
# "Connection: close", If True, server will return "Connection: Keep-Alive"
# in its responses. In order to close the client socket connection
# explicitly after the response is sent and read successfully by the client,
# you simply have to set this option to False when you create a wsgi server.
#http_keepalive = True

# ===== Syslog Options =====

# Send logs to syslog (/dev/log) instead of to file specified
# by `log_file`
#use_syslog = False

# Facility to use. If unset defaults to LOG_USER.
```

```
#syslog_log_facility = LOG_LOCAL1

# ===== SSL Options =====

# Certificate file to use when starting registry server securely
#cert_file = /path/to/certfile

# Private key file to use when starting registry server securely
#key_file = /path/to/keyfile

# CA certificate file to use to verify connecting clients
#ca_file = /path/to/cafile

# ===== Notification System Options =====

# Driver or drivers to handle sending notifications. Set to
# 'messaging' to send notifications to a message queue.
# notification_driver = noop

# Default publisher_id for outgoing notifications.
# default_publisher_id = image.localhost

# Messaging driver used for 'messaging' notifications driver
# rpc_backend = 'rabbit'

# Configuration options if sending notifications via rabbitmq (these are
# the defaults)
rabbit_host = localhost
rabbit_port = 5672
rabbit_use_ssl = false
rabbit_userid = guest
rabbit_password = guest
rabbit_virtual_host = /
rabbit_notification_exchange = glance
rabbit_notification_topic = notifications
rabbit_durable_queues = False

# Configuration options if sending notifications via Qpid (these are
# the defaults)
qpid_notification_exchange = glance
qpid_notification_topic = notifications
qpid_hostname = localhost
qpid_port = 5672
qpid_username =
qpid_password =
qpid_sasl_mechanisms =
qpid_reconnect_timeout = 0
qpid_reconnect_limit = 0
qpid_reconnect_interval_min = 0
qpid_reconnect_interval_max = 0
qpid_reconnect_interval = 0
qpid_heartbeat = 5
# Set to 'ssl' to enable SSL
qpid_protocol = tcp
qpid_tcp_nodelay = True

# ===== Policy Options =====

[oslo_policy]
```

```
# The JSON file that defines policies.
# Deprecated group/name - [DEFAULT]/policy_file
#policy_file = policy.json

# Default rule. Enforced when a requested rule is not found.
# Deprecated group/name - [DEFAULT]/policy_default_rule
#policy_default_rule = default

# 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.
# Deprecated group/name - [DEFAULT]/policy_dirs
#policy_dirs = policy.d

# ===== Database Options =====

[database]
# The file name to use with SQLite (string value)
#sqlite_db = glance.sqlite

# If True, SQLite uses synchronous mode (boolean value)
#sqlite_synchronous = True

# The backend to use for db (string value)
# Deprecated group/name - [DEFAULT]/db_backend
#backend = sqlalchemy

# The SQLAlchemy connection string used to connect to the
# database (string value)
# Deprecated group/name - [DEFAULT]/sql_connection
# Deprecated group/name - [DATABASE]/sql_connection
# Deprecated group/name - [sql]/connection
#connection = <None>

# 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= (string
# value)
#mysql_sql_mode = TRADITIONAL

# Timeout before idle sql connections are reaped (integer
# value)
# Deprecated group/name - [DEFAULT]/sql_idle_timeout
# Deprecated group/name - [DATABASE]/sql_idle_timeout
# Deprecated group/name - [sql]/idle_timeout
#idle_timeout = 3600

# Minimum number of SQL connections to keep open in a pool
# (integer value)
# Deprecated group/name - [DEFAULT]/sql_min_pool_size
# Deprecated group/name - [DATABASE]/sql_min_pool_size
#min_pool_size = 1

# Maximum number of SQL connections to keep open in a pool
# (integer value)
# Deprecated group/name - [DEFAULT]/sql_max_pool_size
# Deprecated group/name - [DATABASE]/sql_max_pool_size
```



```
#max_pool_size = <None>

# Maximum db connection retries during startup. (setting -1
# implies an infinite retry count) (integer value)
# Deprecated group/name - [DEFAULT]/sql_max_retries
# Deprecated group/name - [DATABASE]/sql_max_retries
#max_retries = 10

# Interval between retries of opening a sql connection
# (integer value)
# Deprecated group/name - [DEFAULT]/sql_retry_interval
# Deprecated group/name - [DATABASE]/reconnect_interval
#retry_interval = 10

# If set, use this value for max_overflow with sqlalchemy
# (integer value)
# Deprecated group/name - [DEFAULT]/sql_max_overflow
# Deprecated group/name - [DATABASE]/sqlalchemy_max_overflow
#max_overflow = <None>

# Verbosity of SQL debugging information. 0=None,
# 100=Everything (integer value)
# Deprecated group/name - [DEFAULT]/sql_connection_debug
#connection_debug = 0

# Add python stack traces to SQL as comment strings (boolean
# value)
# Deprecated group/name - [DEFAULT]/sql_connection_trace
#connection_trace = False

# If set, use this value for pool_timeout with sqlalchemy
# (integer value)
# Deprecated group/name - [DATABASE]/sqlalchemy_pool_timeout
#pool_timeout = <None>

# Enable the experimental use of database reconnect on
# connection lost (boolean value)
#use_db_reconnect = False

# seconds between db connection retries (integer value)
#db_retry_interval = 1

# Whether to increase interval between db connection retries,
# up to db_max_retry_interval (boolean value)
#db_inc_retry_interval = True

# max seconds between db connection retries, if
# db_inc_retry_interval is enabled (integer value)
#db_max_retry_interval = 10

# maximum db connection retries before error is raised.
# (setting -1 implies an infinite retry count) (integer value)
#db_max_retries = 20

[keystone_authtoken]
identity_uri = http://127.0.0.1:35357
admin_tenant_name = %SERVICE_TENANT_NAME%
admin_user = %SERVICE_USER%
admin_password = %SERVICE_PASSWORD%
```





```
[DEFAULT]

#
# From glance.manage
#

# Print debugging output (set logging level to DEBUG instead of
# default WARNING level). (boolean value)
#debug = false

# 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.
# (string value)
# Deprecated group/name - [DEFAULT]/log_config
#log_config_append = <None>

# Format string for %(asctime)s in log records. Default: %(default)s
# . (string value)
#log_date_format = %Y-%m-%d %H:%M:%S

# (Optional) The base directory used for relative --log-file paths.
# (string value)
# Deprecated group/name - [DEFAULT]/logdir
#log_dir = <None>

# (Optional) Name of log file to output to. If no default is set,
# logging will go to stdout. (string value)
# Deprecated group/name - [DEFAULT]/logfile
log_file = /var/log/glance/manage.log

# DEPRECATED. A logging.Formatter log message format string which may
# use any of the available logging.LogRecord attributes. This option
# is deprecated. Please use logging_context_format_string and
# logging_default_format_string instead. (string value)
#log_format = <None>

# Syslog facility to receive log lines. (string value)
#syslog_log_facility = LOG_USER

# Use syslog for logging. Existing syslog format is DEPRECATED during
# I, and will change in J to honor RFC5424. (boolean value)
#use_syslog = false

# (Optional) Enables or disables syslog rfc5424 format for logging. If
# enabled, prefixes the MSG part of the syslog message with APP-NAME
# (RFC5424). The format without the APP-NAME is deprecated in I, and
# will be removed in J. (boolean value)
#use_syslog_rfc_format = false

# Print more verbose output (set logging level to INFO instead of
# default WARNING level). (boolean value)
#verbose = false

[database]

#
# From oslo.db
#
```

```
# The back end to use for the database. (string value)
# Deprecated group/name - [DEFAULT]/db_backend
#backend = sqlalchemy

# The SQLAlchemy connection string to use to connect to the database.
# (string value)
# Deprecated group/name - [DEFAULT]/sql_connection
# Deprecated group/name - [DATABASE]/sql_connection
# Deprecated group/name - [sql]/connection
#connection = <None>

# Verbosity of SQL debugging information: 0=None, 100=Everything.
# (integer value)
# Deprecated group/name - [DEFAULT]/sql_connection_debug
#connection_debug = 0

# Add Python stack traces to SQL as comment strings. (boolean value)
# Deprecated group/name - [DEFAULT]/sql_connection_trace
#connection_trace = false

# If True, increases the interval between database connection retries
# up to db_max_retry_interval. (boolean value)
#db_inc_retry_interval = true

# Maximum database connection retries before error is raised. Set to
# -1 to specify an infinite retry count. (integer value)
#db_max_retries = 20

# If db_inc_retry_interval is set, the maximum seconds between
# database connection retries. (integer value)
#db_max_retry_interval = 10

# Seconds between database connection retries. (integer value)
#db_retry_interval = 1

# Timeout before idle SQL connections are reaped. (integer value)
# Deprecated group/name - [DEFAULT]/sql_idle_timeout
# Deprecated group/name - [DATABASE]/sql_idle_timeout
# Deprecated group/name - [sql]/idle_timeout
#idle_timeout = 3600

# If set, use this value for max_overflow with SQLAlchemy. (integer
# value)
# Deprecated group/name - [DEFAULT]/sql_max_overflow
# Deprecated group/name - [DATABASE]/sqlalchemy_max_overflow
#max_overflow = <None>

# Maximum number of SQL connections to keep open in a pool. (integer
# value)
# Deprecated group/name - [DEFAULT]/sql_max_pool_size
# Deprecated group/name - [DATABASE]/sql_max_pool_size
#max_pool_size = <None>

# Maximum number of database connection retries during startup. Set to
# -1 to specify an infinite retry count. (integer value)
# Deprecated group/name - [DEFAULT]/sql_max_retries
# Deprecated group/name - [DATABASE]/sql_max_retries
#max_retries = 10
```





```

# Directory that the scrubber will use to remind itself of what to delete
# Make sure this is also set in glance-api.conf
scrubber_datadir = /var/lib/glance/scrubber

# Only one server in your deployment should be designated the cleanup host
cleanup_scrubber = False

# pending_delete items older than this time are candidates for cleanup
cleanup_scrubber_time = 86400

# Address to find the registry server for cleanups
registry_host = 0.0.0.0

# Port the registry server is listening on
registry_port = 9191

# Auth settings if using Keystone
# auth_url = http://127.0.0.1:5000/v2.0/
# admin_tenant_name = %SERVICE_TENANT_NAME%
# admin_user = %SERVICE_USER%
# admin_password = %SERVICE_PASSWORD%

# API to use for accessing data. Default value points to sqlalchemy
# package, it is also possible to use: glance.db.registry.api
#data_api = glance.db.sqlalchemy.api

# ===== Security Options =====

# AES key for encrypting store 'location' metadata, including
# -- if used -- Swift or S3 credentials
# Should be set to a random string of length 16, 24 or 32 bytes
#metadata_encryption_key = <16, 24 or 32 char registry metadata key>

# ===== Policy Options =====

# The JSON file that defines policies.
#policy_file = policy.json

# Default rule. Enforced when a requested rule is not found.
#policy_default_rule = default

# 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.
#policy_dirs = policy.d

# ===== Database Options =====+=====

[database]

# The SQLAlchemy connection string used to connect to the
# database (string value)
#connection=sqlite:///glance/openstack/common/db/$sqlite_db

# The SQLAlchemy connection string used to connect to the
# slave database (string value)
#slave_connection=
    
```









































## IBM SDN-VE configuration options

**Table 9.9. Description of SDN-VE configuration options**

Configuration option = Default value	Description
[SDNVE]	
<code>base_url = /one/nb/v2/</code>	(StrOpt) Base URL for SDN-VE controller REST API.
<code>controller_ips = 127.0.0.1</code>	(ListOpt) List of IP addresses of SDN-VE controller(s).
<code>default_tenant_type = OVERLAY</code>	(StrOpt) Tenant type: OVERLAY (default) or OF.
<code>format = json</code>	(StrOpt) SDN-VE request/response format.
<code>info = sdnve_info_string</code>	(StrOpt) SDN-VE RPC subject.
<code>integration_bridge = None</code>	(StrOpt) Integration bridge to use.
<code>interface_mappings =</code>	(ListOpt) List of <physical_network_name>:<interface_name> mappings.
<code>of_signature = SDNVE-OF</code>	(StrOpt) The string in tenant description that indicates the tenant is a OF tenant.
<code>out_of_band = True</code>	(BoolOpt) Indicating if controller is out of band or not.
<code>overlay_signature = SDNVE-OVERLAY</code>	(StrOpt) The string in tenant description that indicates the tenant is a OVERLAY tenant.
<code>password = admin</code>	(StrOpt) SDN-VE administrator password.
<code>port = 8443</code>	(StrOpt) SDN-VE controller port number.
<code>reset_bridge = True</code>	(BoolOpt) Whether to reset the integration bridge before use.
<code>use_fake_controller = False</code>	(BoolOpt) Whether to use a fake controller.
<code>userid = admin</code>	(StrOpt) SDN-VE administrator user ID.
[SDNVE_AGENT]	
<code>polling_interval = 2</code>	(IntOpt) Agent polling interval if necessary.
<code>rpc = True</code>	(BoolOpt) Whether to use rpc.

## Layer 2 Gateway configuration options

**Table 9.10. Description of Layer 2 Gateway configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>default_device_name = Switch1</code>	(StrOpt) default_device_name of the l2 gateway
<code>default_l2_gw_service_uuid = None</code>	(StrOpt) Unique identifier of the NSX L2 Gateway service which will be used by default for network gateways
<code>default_l3_gw_service_uuid = None</code>	(StrOpt) Unique identifier of the NSX L3 Gateway service which will be used for implementing routers and floating IPs
<code>l2gw_callback_class = networking_l2gw.services.l2gateway.ovsdb.data_agent</code>	(StrOpt) L2 gateway plugin callback class where the RPCs from the agent are going to get invoked
<code>quota_l2_gateway = 5</code>	(IntOpt) Number of l2 gateways allowed per tenant, -1 for unlimited
[ovsdb]	
<code>l2_gw_agent_ca_cert_base_path = None</code>	(StrOpt) Trusted issuer CA cert
<code>l2_gw_agent_cert_base_path = None</code>	(StrOpt) L2 gateway agent public certificate
<code>l2_gw_agent_priv_key_base_path = None</code>	(StrOpt) L2 gateway agent private key
<code>max_connection_retries = 10</code>	(IntOpt) Maximum number of retries to open a socket with the OVSDb server

Configuration option = Default value	Description
ovsdb_hosts = <i>host1:127.0.0.1:6632</i>	(StrOpt) OVSDDB server name:host/IP:port
periodic_interval = <i>20</i>	(IntOpt) Seconds between periodic task runs

## Linux bridge Agent configuration options

**Table 9.11. Description of Linux Bridge agent configuration options**

Configuration option = Default value	Description
[LINUX_BRIDGE]	
physical_interface_mappings =	(ListOpt) List of <physical_network>:<physical_interface>
[VLANS]	
network_vlan_ranges =	(ListOpt) List of <physical_network>:<vlan_min>:<vlan_max> or <physical_network>
tenant_network_type = <i>local</i>	(StrOpt) Network type for tenant networks (local, vlan, or none)
[VXLAN]	
enable_vxlan = <i>True</i>	(BoolOpt) Enable VXLAN on the agent. Can be enabled when agent is managed by ml2 plugin using linuxbridge mechanism driver
l2_population = <i>False</i>	(BoolOpt) Extension to use alongside ml2 plugin's l2population mechanism driver. It enables the plugin to populate VXLAN forwarding table.
local_ip = <i>None</i>	(IPOpt) Local IP address of the VXLAN endpoints.
tos = <i>None</i>	(IntOpt) TOS for vxlan interface protocol packets.
ttl = <i>None</i>	(IntOpt) TTL for vxlan interface protocol packets.
vxlan_group = <i>224.0.0.1</i>	(StrOpt) Multicast group for vxlan interface.

## Meta Plug-in configuration options

The Meta Plug-in allows you to use multiple plug-ins at the same time.

**Table 9.12. Description of meta configuration options**

Configuration option = Default value	Description
[META]	
default_flavor =	(StrOpt) Default flavor to use, when flavor:network is not specified at network creation.
default_l3_flavor =	(StrOpt) Default L3 flavor to use, when flavor:router is not specified at router creation. Ignored if 'l3_plugin_list' is blank.
extension_map =	(StrOpt) Comma separated list of method:flavor to select specific plugin for a method. This has priority over method search order based on 'plugin_list'.
l3_plugin_list =	(StrOpt) Comma separated list of flavor:neutron_plugin for L3 service plugins to load. This is intended for specifying L2 plugins which support L3 functions. If you use a router service plugin, set this blank.
plugin_list =	(StrOpt) Comma separated list of flavor:neutron_plugin for plugins to load. Extension method is searched in the list order and the first one is used.

Configuration option = Default value	Description
<code>rpc_flavor =</code>	(StrOpt) Specifies flavor for plugin to handle 'q-plugin' RPC requests.
<code>supported_extension_aliases =</code>	(StrOpt) Comma separated list of supported extension aliases.

## Modular Layer 2 (ml2) configuration options

The Modular Layer 2 (ml2) plug-in has two components: network types and mechanisms. You can configure these components separately. This section describes these configuration options.



### Configure MTU for VXLAN tunnelling

Specific MTU configuration is necessary for VXLAN to function as expected:

- One option is to increase the MTU value of the physical interface and physical switch fabric by at least 50 bytes. For example, increase the MTU value to 1550. This value enables an automatic 50-byte MTU difference between the physical interface (1500) and the VXLAN interface (automatically 1500-50 = 1450). An MTU value of 1450 causes issues when virtual machine taps are configured at an MTU value of 1500.
- Another option is to decrease the virtual Ethernet devices' MTU. Set the `network_device_mtu` option to 1450 in the `neutron.conf` file, and set all guest virtual machines' MTU to the same value by using a DHCP option. For information about how to use this option, see [Configure OVS plug-in](#).

**Table 9.13. Description of ML2 configuration options**

Configuration option = Default value	Description
[ml2]	
<code>extension_drivers =</code>	(ListOpt) An ordered list of extension driver endpoints to be loaded from the <code>neutron.ml2.extension_drivers</code> namespace.
<code>mechanism_drivers =</code>	(ListOpt) An ordered list of networking mechanism driver endpoints to be loaded from the <code>neutron.ml2.mechanism_drivers</code> namespace.
<code>path_mtu = 0</code>	(IntOpt) The maximum permissible size of an unfragmented packet travelling from and to addresses where encapsulated Neutron traffic is sent. If $\leq 0$ , the path MTU is indeterminate.
<code>physical_network_mtus =</code>	(ListOpt) A list of mappings of physical networks to MTU values. The format of the mapping is <code>&lt;physnet&gt;:&lt;mtu val&gt;</code> . This mapping allows specifying a physical network MTU value that differs from the default <code>segment_mtu</code> value.
<code>segment_mtu = 0</code>	(IntOpt) The maximum permissible size of an unfragmented packet travelling a L2 network segment. If $\leq 0$ , the segment MTU is indeterminate.
<code>tenant_network_types = local</code>	(ListOpt) Ordered list of <code>network_types</code> to allocate as tenant networks.
<code>type_drivers = local, flat, vlan, gre, vxlan</code>	(ListOpt) List of network type driver endpoints to be loaded from the <code>neutron.ml2.type_drivers</code> namespace.

**Modular Layer 2 (ml2) Flat Type configuration options**

**Table 9.14. Description of ML2 Flat mechanism driver configuration options**

Configuration option = Default value	Description
[ml2_type_flat]	
flat_networks =	(ListOpt) List of physical_network names with which flat networks can be created. Use * to allow flat networks with arbitrary physical_network names.

**Modular Layer 2 (ml2) GRE Type configuration options**

**Table 9.15. Description of ML2 GRE configuration options**

Configuration option = Default value	Description
[ml2_type_gre]	
tunnel_id_ranges =	(ListOpt) Comma-separated list of <tun_min>:<tun_max> tuples enumerating ranges of GRE tunnel IDs that are available for tenant network allocation

**Modular Layer 2 (ml2) VLAN Type configuration options**

**Table 9.16. Description of ML2 VLAN configuration options**

Configuration option = Default value	Description
[ml2_type_vlan]	
network_vlan_ranges =	(ListOpt) 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.

**Modular Layer 2 (ml2) VXLAN Type configuration options**

**Table 9.17. Description of ML2 VXLAN configuration options**

Configuration option = Default value	Description
[ml2_type_nexus_vxlan]	
mcast_ranges =	(ListOpt) List of multicast groups to be used for global VNIDs in the format - a:b,c,e:f.
vni_ranges =	(ListOpt) List of global VNID ranges in the format - a:b, c:d. Multiple ranges can be separated by a comma
[ml2_type_vxlan]	
vni_ranges =	(ListOpt) Comma-separated list of <vni_min>:<vni_max> tuples enumerating ranges of VXLAN VNI IDs that are available for tenant network allocation
vxlan_group = None	(StrOpt) Multicast group for VXLAN. If unset, disables VXLAN multicast mode.

**Modular Layer 2 (ml2) Arista Mechanism configuration options**

**Table 9.18. Description of ML2 Arista mechanism driver configuration options**

Configuration option = Default value	Description
[ml2_arista]	



Configuration option = Default value	Description
node_override_vif_802.1qbg =	(ListOpt) Nova compute nodes to manually set VIF type to 802.1qbg
node_override_vif_802.1qbh =	(ListOpt) Nova compute nodes to manually set VIF type to 802.1qbh
node_override_vif_binding_failed =	(ListOpt) Nova compute nodes to manually set VIF type to binding_failed
node_override_vif_bridge =	(ListOpt) Nova compute nodes to manually set VIF type to bridge
node_override_vif_distributed =	(ListOpt) Nova compute nodes to manually set VIF type to distributed
node_override_vif_dvs =	(ListOpt) Nova compute nodes to manually set VIF type to dvs
node_override_vif_hostdev =	(ListOpt) Nova compute nodes to manually set VIF type to hostdev
node_override_vif_hw_veb =	(ListOpt) Nova compute nodes to manually set VIF type to hw_veb
node_override_vif_hyperv =	(ListOpt) Nova compute nodes to manually set VIF type to hyperv
node_override_vif_ivs =	(ListOpt) Nova compute nodes to manually set VIF type to ivs
node_override_vif_midonet =	(ListOpt) Nova compute nodes to manually set VIF type to midonet
node_override_vif_mlnx_direct =	(ListOpt) Nova compute nodes to manually set VIF type to mlnx_direct
node_override_vif_other =	(ListOpt) Nova compute nodes to manually set VIF type to other
node_override_vif_ovs =	(ListOpt) Nova compute nodes to manually set VIF type to ovs
node_override_vif_unbound =	(ListOpt) Nova compute nodes to manually set VIF type to unbound
node_override_vif_vrouter =	(ListOpt) Nova compute nodes to manually set VIF type to vrouter
vif_type = <i>ivs</i>	(StrOpt) Virtual interface type to configure on Nova compute nodes
vif_types = <i>unbound, binding_failed, ovs, ivs, bridge, 802.1qbg, 802.1qbh, hyperv, midonet, mlnx_direct, hostdev, hw_veb, dvs, other, distributed, vrouter</i>	(ListOpt) List of allowed vif_type values.
[RESTPROXY]	
add_meta_server_route = <i>True</i>	(BoolOpt) Flag to decide if a route to the metadata server should be injected into the VM
auto_sync_on_failure = <i>True</i>	(BoolOpt) If neutron fails to create a resource because the backend controller doesn't know of a dependency, the plugin automatically triggers a full data synchronization to the controller.
cache_connections = <i>True</i>	(BoolOpt) Re-use HTTP/HTTPS connections to the controller.
consistency_interval = <i>60</i>	(IntOpt) Time between verifications that the backend controller database is consistent with Neutron. (0 to disable)
neutron_id = <i>neutron-images</i>	(StrOpt) User defined identifier for this Neutron deployment
no_ssl_validation = <i>False</i>	(BoolOpt) Disables SSL certificate validation for controllers





























































### Warning

Do not run the `neutron-ns-metadata-proxy` proxy namespace as root on a node with the L3 agent running. In OpenStack Kilo and newer, you can change the permissions of `neutron-ns-metadata-proxy` after the proxy installation using the `metadata_proxy_user` and `metadata_proxy_group` options.

## Metering Agent

Use the following options in the `metering_agent.ini` file for the Metering agent.

**Table 9.59. Description of metering agent configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>driver = neutron.services.metering.drivers.noop.noop_driver</code>	(StrOpt) Metering driver <i>driver.NoopMeteringDriver</i>
<code>measure_interval = 30</code>	(IntOpt) Interval between two metering measures
[AGENT]	
<code>report_interval = 30</code>	(FloatOpt) Seconds between nodes reporting state to server; should be less than <code>agent_down_time</code> , best if it is half or less than <code>agent_down_time</code> .

## Nova

Use the following options in the `neutron.conf` file to change nova-related settings.

**Table 9.60. Description of nova configuration options**

Configuration option = Default value	Description
[nova]	
<code>auth_plugin = None</code>	(StrOpt) Name of the plugin to load
<code>auth_section = None</code>	(StrOpt) Config Section from which to load plugin specific options
<code>cafile = None</code>	(StrOpt) PEM encoded Certificate Authority to use when verifying HTTPS connections.
<code>certfile = None</code>	(StrOpt) PEM encoded client certificate cert file
<code>insecure = False</code>	(BoolOpt) Verify HTTPS connections.
<code>keyfile = None</code>	(StrOpt) PEM encoded client certificate key file
<code>region_name = None</code>	(StrOpt) Name of nova region to use. Useful if keystone manages more than one region.
<code>timeout = None</code>	(IntOpt) Timeout value for http requests

## Policy

Use the following options in the `neutron.conf` file to change oslo middleware settings.

**Table 9.61. Description of oslo\_middleware configuration options**

Configuration option = Default value	Description
[oslo_middleware]	

Configuration option = Default value	Description
<code>max_request_body_size = 114688</code>	(IntOpt) The maximum body size for each request, in bytes.

## Policy

Use the following options in the `neutron.conf` file to change policy settings.

**Table 9.62. Description of policy configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>allow_overlapping_ips = False</code>	(BoolOpt) Allow overlapping IP support in Neutron
<code>policy_default_rule = default</code>	(StrOpt) Default rule. Enforced when a requested rule is not found.
<code>policy_dirs = ['policy.d']</code>	(MultiStrOpt) 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.
<code>policy_file = policy.json</code>	(StrOpt) The JSON file that defines policies.

## Quotas

Use the following options in the `neutron.conf` file for the quota system.

**Table 9.63. Description of quotas configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>max_routes = 30</code>	(IntOpt) Maximum number of routes
[QUOTAS]	
<code>default_quota = -1</code>	(IntOpt) Default number of resource allowed per tenant. A negative value means unlimited.
<code>quota_driver = neutron.db.quota_db.DbQuotaDriver</code>	(StrOpt) Default driver to use for quota checks
<code>quota_floatingip = 50</code>	(IntOpt) Number of floating IPs allowed per tenant. A negative value means unlimited.
<code>quota_health_monitor = -1</code>	(IntOpt) Number of health monitors allowed per tenant. A negative value means unlimited.
<code>quota_items = network, subnet, port</code>	(ListOpt) Resource name(s) that are supported in quota features
<code>quota_member = -1</code>	(IntOpt) Number of pool members allowed per tenant. A negative value means unlimited.
<code>quota_network = 10</code>	(IntOpt) Number of networks allowed per tenant. A negative value means unlimited.
<code>quota_network_gateway = 5</code>	(IntOpt) Number of network gateways allowed per tenant, -1 for unlimited
<code>quota_packet_filter = 100</code>	(IntOpt) Number of packet_filters allowed per tenant, -1 for unlimited
<code>quota_pool = 10</code>	(IntOpt) Number of pools allowed per tenant. A negative value means unlimited.

Configuration option = Default value	Description
<code>quota_port = 50</code>	(IntOpt) Number of ports allowed per tenant. A negative value means unlimited.
<code>quota_router = 10</code>	(IntOpt) Number of routers allowed per tenant. A negative value means unlimited.
<code>quota_security_group = 10</code>	(IntOpt) Number of security groups allowed per tenant. A negative value means unlimited.
<code>quota_security_group_rule = 100</code>	(IntOpt) Number of security rules allowed per tenant. A negative value means unlimited.
<code>quota_subnet = 10</code>	(IntOpt) Number of subnets allowed per tenant, A negative value means unlimited.
<code>quota_vip = 10</code>	(IntOpt) Number of vips allowed per tenant. A negative value means unlimited.

## Scheduler

Use the following options in the `neutron.conf` file to change scheduler settings.

**Table 9.64. Description of scheduler configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>network_auto_schedule = True</code>	(BoolOpt) Allow auto scheduling networks to DHCP agent.
<code>network_scheduler_driver = neutron.scheduler.dhcp_agent_scheduler.ChanceAgentScheduler</code>	(StrOpt) Driver to use for scheduling network to DHCP agent
<code>router_auto_schedule = True</code>	(BoolOpt) Allow auto scheduling of routers to L3 agent.
<code>router_delete_namespaces = False</code>	(BoolOpt) Delete namespace after removing a router.
<code>router_scheduler_driver = neutron.scheduler.l3_agent_scheduler.ChanceAgentScheduler</code>	(StrOpt) Driver to use for scheduling router to a default L3 agent

## Security Groups

Use the following options in the configuration file for your driver to change security group settings.

**Table 9.65. Description of security groups configuration options**

Configuration option = Default value	Description
[SECURITYGROUP]	
<code>defer_apply = True</code>	(BoolOpt) Enable defer_apply on security bridge.
<code>enable_ipset = True</code>	(BoolOpt) Use ipset to speed-up the iptables based security groups.
<code>enable_security_group = True</code>	(BoolOpt) 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.
<code>firewall_driver = None</code>	(StrOpt) Driver for security groups firewall in the L2 agent
<code>ovsvapp_firewall_driver = networking_vsphere.drivers.ovs_firewall.OVSFirewallDriver</code>	(StrOpt) DriverManager implementation for OVS based Firewall
<code>security_bridge_mapping = br-sec</code>	(StrOpt) <security_bridge>:<phy_interface>



**Note**

Now Networking uses iptables to achieve security group functions. In L2 agent with `enable_ipset` option enabled, it makes use of IPset to improve security group's performance, as it represents a hash set which is insensitive to the number of elements.

When a port is created, L2 agent will add an additional IPset chain to its iptables chain, if the security group that this port belongs to has rules between other security group, the member of that security group will be added to the ipset chain.

If a member of a security group is changed, it used to reload iptables rules which is expensive. However, when IPset option is enabled on L2 agent, it does not need to reload iptables if only members of security group were changed, it should just update an IPset.



**Note**

A single default security group has been introduced in order to avoid race conditions when creating a tenant's default security group. The race conditions are caused by the uniqueness check of a new security group name. A table `default_security_group` implements such a group. It has `tenant_id` field as a primary key and `security_group_id`, which is an identifier of a default security group. The migration that introduces this table has a sanity check that verifies if a default security group is not duplicated in any tenant.

## SSL and Certification Authority

Use the following options in the `neutron.conf` file to enable SSL.

**Table 9.66. Description of CA and SSL configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>ssl_ca_file = None</code>	(StrOpt) CA certificate file to use to verify connecting clients
<code>ssl_cert_file = None</code>	(StrOpt) Certificate file to use when starting the server securely
<code>ssl_key_file = None</code>	(StrOpt) Private key file to use when starting the server securely

## Log files used by Networking

The corresponding log file of each Networking service is stored in the `/var/log/neutron/` directory of the host on which each service runs.

**Table 9.67. Log files used by Networking services**

Log file	Service/interface
<code>dhcp-agent.log</code>	<code>neutron-dhcp-agent</code>
<code>l3-agent.log</code>	<code>neutron-l3-agent</code>

Log file	Service/interface
lbaas-agent.log	neutron-lbaas-agent <sup>a</sup>
linuxbridge-agent.log	neutron-linuxbridge-agent
metadata-agent.log	neutron-metadata-agent
metering-agent.log	neutron-metering-agent
openvswitch-agent.log	neutron-openvswitch-agent
server.log	neutron-server

<sup>a</sup>The neutron-lbaas-agent service only runs when Load-Balancer-as-a-Service is enabled.

## Networking sample configuration files

All the files in this section can be found in `/etc/neutron/`.

### neutron.conf

Use the `neutron.conf` file to configure the majority of the OpenStack Networking options.

```
[DEFAULT]
# Print more verbose output (set logging level to INFO instead of default
# WARNING level).
# verbose = False

# =====Start Global Config Option for Distributed L3 Router=====
# Setting the "router_distributed" flag to "True" will default to the creation
# of distributed tenant routers. The admin can override this flag by
# specifying
# the type of the router on the create request (admin-only attribute). Default
# value is "False" to support legacy mode (centralized) routers.
#
# router_distributed = False
#
# =====End Global Config Option for Distributed L3 Router=====

# Print debugging output (set logging level to DEBUG instead of default
# WARNING level).
# debug = False

# Where to store Neutron state files. This directory must be writable by the
# user executing the agent.
# state_path = /var/lib/neutron

# log_format = %(asctime)s %(levelname)8s [%(name)s] %(message)s
# log_date_format = %Y-%m-%d %H:%M:%S

# use_syslog                                -> syslog
# log_file and log_dir                       -> log_dir/log_file
# (not log_file) and log_dir                 -> log_dir/{binary_name}.log
# use_stderr                                 -> stderr
# (not user_stderr) and (not log_file)       -> stdout
# publish_errors                             -> notification system

# use_syslog = False
# syslog_log_facility = LOG_USER
```

```
# use_stderr = True
# log_file =
# log_dir =

# publish_errors = False

# Address to bind the API server to
# bind_host = 0.0.0.0

# Port the bind the API server to
# bind_port = 9696

# Path to the extensions. Note that this can be a colon-separated list of
# paths. For example:
# api_extensions_path = extensions:/path/to/more/extensions:/even/more/
extensions
# The __path__ of neutron.extensions is appended to this, so if your
# extensions are in there you don't need to specify them here
# api_extensions_path =

# (StrOpt) Neutron core plugin entrypoint to be loaded from the
# neutron.core_plugins namespace. See setup.cfg for the entrypoint names of
the
# plugins included in the neutron source distribution. For compatibility with
# previous versions, the class name of a plugin can be specified instead of
its
# entrypoint name.
#
# core_plugin =
# Example: core_plugin = ml2

# (ListOpt) List of service plugin entrypoints to be loaded from the
# neutron.service_plugins namespace. See setup.cfg for the entrypoint names of
# the plugins included in the neutron source distribution. For compatibility
# with previous versions, the class name of a plugin can be specified instead
# of its entrypoint name.
#
# service_plugins =
# Example: service_plugins = router,firewall,lbaas,vpnaas,metering

# Paste configuration file
# api_paste_config = api-paste.ini

# (StrOpt) 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.
# The default value is hostname of the machine.
#
# host =

# The strategy to be used for auth.
# Supported values are 'keystone'(default), 'noauth'.
# auth_strategy = keystone

# Base MAC address. The first 3 octets will remain unchanged. If the
# 4h octet is not 00, it will also be used. The others will be
# randomly generated.
# 3 octet
# base_mac = fa:16:3e:00:00:00
# 4 octet
```

```
# base_mac = fa:16:3e:4f:00:00

# DVR Base MAC address. 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 'dvr_base_mac' *must* be different from 'base_mac' to
# avoid mixing them up with MAC's allocated for tenant ports.
# A 4 octet example would be dvr_base_mac = fa:16:3f:4f:00:00
# The default is 3 octet
# dvr_base_mac = fa:16:3f:00:00:00

# Maximum amount of retries to generate a unique MAC address
# mac_generation_retries = 16

# DHCP Lease duration (in seconds). Use -1 to
# tell dnsmasq to use infinite lease times.
# dhcp_lease_duration = 86400

# Allow sending resource operation notification to DHCP agent
# dhcp_agent_notification = True

# Enable or disable bulk create/update/delete operations
# allow_bulk = True
# Enable or disable pagination
# allow_pagination = False
# Enable or disable sorting
# allow_sorting = False
# Enable or disable overlapping IPs for subnets
# Attention: the following parameter MUST be set to False if Neutron is
# being used in conjunction with nova security groups
# allow_overlapping_ips = False
# Ensure that configured gateway is on subnet. For IPv6, validate only if
# gateway is not a link local address. Deprecated, to be removed during the
# K release, at which point the check will be mandatory.
# force_gateway_on_subnet = True

# Default maximum number of items returned in a single response,
# value == infinite and value < 0 means no max limit, and value must
# be greater than 0. If the number of items requested is greater than
# pagination_max_limit, server will just return pagination_max_limit
# of number of items.
# pagination_max_limit = -1

# Maximum number of DNS nameservers per subnet
# max_dns_nameservers = 5

# Maximum number of host routes per subnet
# max_subnet_host_routes = 20

# Maximum number of fixed ips per port
# max_fixed_ips_per_port = 5

# Maximum number of routes per router
# max_routes = 30

# Default Subnet Pool to be used for IPv4 subnet-allocation.
# Specifies by UUID the pool to be used in case of subnet-create being called
# without a subnet-pool ID. The default of None means that no pool will be
# used unless passed explicitly to subnet create. If no pool is used, then a
# CIDR must be passed to create a subnet and that subnet will not be allocated
# from any pool; it will be considered part of the tenant's private address
```

```
# space.
# default_ipv4_subnet_pool =

# Default Subnet Pool to be used for IPv6 subnet-allocation.
# Specifies by UUID the pool to be used in case of subnet-create being
# called without a subnet-pool ID. Set to "prefix_delegation"
# to enable IPv6 Prefix Delegation in a PD-capable environment.
# See the description for default_ipv4_subnet_pool for more information.
# default_ipv6_subnet_pool =

# ===== items for MTU selection and advertisement =====
# Advertise MTU. If True, effort is made to advertise MTU
# settings to VMs via network methods (ie. DHCP and RA MTU options)
# when the network's preferred MTU is known.
# advertise_mtu = False
# ===== end of items for MTU selection and advertisement =====

# ===== items for agent management extension =====
# Seconds to regard the agent as down; should be at least twice
# report_interval, to be sure the agent is down for good
# agent_down_time = 75
# ===== end of items for agent management extension =====

# ===== items for agent scheduler extension =====
# Driver to use for scheduling network to DHCP agent
# network_scheduler_driver = neutron.scheduler.dhcp_agent_scheduler.
ChanceScheduler
# Driver to use for scheduling router to a default L3 agent
# router_scheduler_driver = neutron.scheduler.l3_agent_scheduler.
ChanceScheduler
# Driver to use for scheduling a loadbalancer pool to an lbaas agent
# loadbalancer_pool_scheduler_driver = neutron.services.loadbalancer.
agent_scheduler.ChanceScheduler

# (StrOpt) 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
# Values:
# networks - number of networks hosted on the agent
# subnets - number of subnets associated with the networks hosted on the
# agent
# ports - number of ports associated with the networks hosted on the
# agent
# dhcp_load_type = networks

# Allow auto scheduling networks to DHCP agent. It will schedule non-hosted
# networks to first DHCP agent which sends get_active_networks message to
# neutron server
# network_auto_schedule = True
```













```
# Maximum number of ingress messages to locally buffer per
# topic. Default is unlimited. (integer value)
# rpc_zmq_topic_backlog=

# Directory for holding IPC sockets. (string value)
# rpc_zmq_ipc_dir=/var/run/openstack

# Name of this node. Must be a valid hostname, FQDN, or IP
# address. Must match "host" option, if running Nova. (string
# value)
# rpc_zmq_host=oslo

# Seconds to wait before a cast expires (TTL). Only supported
# by impl_zmq. (integer value)
# rpc_cast_timeout=30

# Heartbeat frequency. (integer value)
# matchmaker_heartbeat_freq=300

# Heartbeat time-to-live. (integer value)
# matchmaker_heartbeat_ttl=600

# Size of RPC greenthread pool. (integer value)
# rpc_thread_pool_size=64

# Driver or drivers to handle sending notifications. (multi
# valued)
# notification_driver=

# AMQP topic used for OpenStack notifications. (list value)
# Deprecated group/name - [rpc_notifier2]/topics
# notification_topics=notifications

# Seconds to wait for a response from a call. (integer value)
# rpc_response_timeout=60

# A URL representing the messaging driver to use and its full
# configuration. If not set, we fall back to the rpc_backend
# option and driver specific configuration. (string value)
# transport_url=

# The messaging driver to use, defaults to rabbit. Other
# drivers include qpid and zmq. (string value)
# rpc_backend=rabbit

# The default exchange under which topics are scoped. May be
# overridden by an exchange name specified in the
# transport_url option. (string value)
# control_exchange=openstack

[matchmaker_redis]

#
# Options defined in oslo.messaging
#

# Host to locate redis. (string value)
# host=127.0.0.1
```

```
# Use this port to connect to redis host. (integer value)
# port=6379

# Password for Redis server (optional). (string value)
# password=

[matchmaker_ring]

#
# Options defined in oslo.messaging
#

# Matchmaker ring file (JSON). (string value)
# Deprecated group/name - [DEFAULT]/matchmaker_ringfile
# ringfile=/etc/oslo/matchmaker_ring.json

[quotas]
# Default driver to use for quota checks
# quota_driver = neutron.db.quota_db.DbQuotaDriver

# Resource name(s) that are supported in quota features
# quota_items = network,subnet,port

# Default number of resource allowed per tenant. A negative value means
# unlimited.
# default_quota = -1

# Number of networks allowed per tenant. A negative value means unlimited.
# quota_network = 10

# Number of subnets allowed per tenant. A negative value means unlimited.
# quota_subnet = 10

# Number of ports allowed per tenant. A negative value means unlimited.
# quota_port = 50

# Number of security groups allowed per tenant. A negative value means
# unlimited.
# quota_security_group = 10

# Number of security group rules allowed per tenant. A negative value means
# unlimited.
# quota_security_group_rule = 100

# Number of vips allowed per tenant. A negative value means unlimited.
# quota_vip = 10

# Number of pools allowed per tenant. A negative value means unlimited.
# quota_pool = 10

# Number of pool members allowed per tenant. A negative value means unlimited.
# The default is unlimited because a member is not a real resource consumer
# on Openstack. However, on back-end, a member is a resource consumer
# and that is the reason why quota is possible.
# quota_member = -1

# Number of health monitors allowed per tenant. A negative value means
# unlimited.
```

```
# The default is unlimited because a health monitor is not a real resource
# consumer on Openstack. However, on back-end, a member is a resource consumer
# and that is the reason why quota is possible.
# quota_health_monitor = -1

# Number of loadbalancers allowed per tenant. A negative value means
# unlimited.
# quota_loadbalancer = 10

# Number of listeners allowed per tenant. A negative value means unlimited.
# quota_listener = -1

# Number of v2 health monitors allowed per tenant. A negative value means
# unlimited. These health monitors exist under the lbaas v2 API
# quota_healthmonitor = -1

# Number of routers allowed per tenant. A negative value means unlimited.
# quota_router = 10

# Number of floating IPs allowed per tenant. A negative value means unlimited.
# quota_floatingip = 50

# Number of firewalls allowed per tenant. A negative value means unlimited.
# quota_firewall = 1

# Number of firewall policies allowed per tenant. A negative value means
# unlimited.
# quota_firewall_policy = 1

# Number of firewall rules allowed per tenant. A negative value means
# unlimited.
# quota_firewall_rule = 100

[agent]
# Use "sudo neutron-rootwrap /etc/neutron/rootwrap.conf" to use the real
# root filter facility.
# Change to "sudo" to skip the filtering and just run the command directly
# root_helper = sudo

# Set to true to add comments to generated iptables rules that describe
# each rule's purpose. (System must support the iptables comments module.)
# comment_iptables_rules = True

# Root helper daemon application to use when possible.
# root_helper_daemon =

# 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.
# use_helper_for_ns_read = True

# The interval to check external processes for failure in seconds (0=disabled)
# check_child_processes_interval = 60

# Action to take when an external process spawned by an agent dies
# Values:
#   respawn - Respawns the external process
#   exit - Exits the agent
# check_child_processes_action = respawn
```



```
# ===== items for agent management extension =====
# 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
# report_interval = 30

# ===== end of items for agent management extension =====

[keystone_authtoken]
auth_uri = http://127.0.0.1:35357/v2.0/
identity_uri = http://127.0.0.1:5000
admin_tenant_name = %SERVICE_TENANT_NAME%
admin_user = %SERVICE_USER%
admin_password = %SERVICE_PASSWORD%

[database]
# This line MUST be changed to actually run the plugin.
# Example:
# connection = mysql://root:pass@127.0.0.1:3306/neutron
# Replace 127.0.0.1 above with the IP address of the database used by the
# main neutron server. (Leave it as is if the database runs on this host.)
# connection = sqlite://
# NOTE: In deployment the [database] section and its connection attribute may
# be set in the corresponding core plugin '.ini' file. However, it is
# suggested
# to put the [database] section and its connection attribute in this
# configuration file.

# Database engine for which script will be generated when using offline
# migration
# engine =

# The SQLAlchemy connection string used to connect to the slave database
# slave_connection =

# Database reconnection retry times - in event connectivity is lost
# set to -1 implies an infinite retry count
# max_retries = 10

# Database reconnection interval in seconds - if the initial connection to the
# database fails
# retry_interval = 10

# Minimum number of SQL connections to keep open in a pool
# min_pool_size = 1

# Maximum number of SQL connections to keep open in a pool
# max_pool_size = 10

# Timeout in seconds before idle sql connections are reaped
# idle_timeout = 3600

# If set, use this value for max_overflow with sqlalchemy
# max_overflow = 20

# Verbosity of SQL debugging information. 0=None, 100=Everything
# connection_debug = 0

# Add python stack traces to SQL as comment strings
# connection_trace = False
```





```
# Size of RPC connection pool. (integer value)
# Deprecated group/name - [DEFAULT]/rpc_conn_pool_size
# rpc_conn_pool_size = 30

# Qpid broker hostname. (string value)
# Deprecated group/name - [DEFAULT]/qpid_hostname
# qpid_hostname = localhost

# Qpid broker port. (integer value)
# Deprecated group/name - [DEFAULT]/qpid_port
# qpid_port = 5672

# Qpid HA cluster host:port pairs. (list value)
# Deprecated group/name - [DEFAULT]/qpid_hosts
# qpid_hosts = $qpid_hostname:$qpid_port

# Username for Qpid connection. (string value)
# Deprecated group/name - [DEFAULT]/qpid_username
# qpid_username =

# Password for Qpid connection. (string value)
# Deprecated group/name - [DEFAULT]/qpid_password
# qpid_password =

# Space separated list of SASL mechanisms to use for auth. (string value)
# Deprecated group/name - [DEFAULT]/qpid_sasl_mechanisms
# qpid_sasl_mechanisms =

# Seconds between connection keepalive heartbeats. (integer value)
# Deprecated group/name - [DEFAULT]/qpid_heartbeat
# qpid_heartbeat = 60

# Transport to use, either 'tcp' or 'ssl'. (string value)
# Deprecated group/name - [DEFAULT]/qpid_protocol
# qpid_protocol = tcp

# Whether to disable the Nagle algorithm. (boolean value)
# Deprecated group/name - [DEFAULT]/qpid_tcp_nodelay
# qpid_tcp_nodelay = true

# The number of prefetched messages held by receiver. (integer value)
# Deprecated group/name - [DEFAULT]/qpid_receiver_capacity
# qpid_receiver_capacity = 1

# The qpid topology version to use. Version 1 is what was originally used by
# impl_qpid. Version 2 includes some backwards-incompatible changes that
# allow
# broker federation to work. Users should update to version 2 when they are
# able to take everything down, as it requires a clean break. (integer value)
# Deprecated group/name - [DEFAULT]/qpid_topology_version
# qpid_topology_version = 1

[oslo_messaging_rabbit]

#
# From oslo.messaging
#

# Use durable queues in AMQP. (boolean value)
```

```
# Deprecated group/name - [DEFAULT]/rabbit_durable_queues
# amqp_durable_queues = false

# Auto-delete queues in AMQP. (boolean value)
# Deprecated group/name - [DEFAULT]/amqp_auto_delete
# amqp_auto_delete = false

# Size of RPC connection pool. (integer value)
# Deprecated group/name - [DEFAULT]/rpc_conn_pool_size
# rpc_conn_pool_size = 30

# 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. (string value)
# Deprecated group/name - [DEFAULT]/kombu_ssl_version
# kombu_ssl_version =

# SSL key file (valid only if SSL enabled). (string value)
# Deprecated group/name - [DEFAULT]/kombu_ssl_keyfile
# kombu_ssl_keyfile =

# SSL cert file (valid only if SSL enabled). (string value)
# Deprecated group/name - [DEFAULT]/kombu_ssl_certfile
# kombu_ssl_certfile =

# SSL certification authority file (valid only if SSL enabled). (string value)
# Deprecated group/name - [DEFAULT]/kombu_ssl_ca_certs
# kombu_ssl_ca_certs =

# How long to wait before reconnecting in response to an AMQP consumer cancel
# notification. (floating point value)
# Deprecated group/name - [DEFAULT]/kombu_reconnect_delay
# kombu_reconnect_delay = 1.0

# The RabbitMQ broker address where a single node is used. (string value)
# Deprecated group/name - [DEFAULT]/rabbit_host
# rabbit_host = localhost

# The RabbitMQ broker port where a single node is used. (integer value)
# Deprecated group/name - [DEFAULT]/rabbit_port
# rabbit_port = 5672

# RabbitMQ HA cluster host:port pairs. (list value)
# Deprecated group/name - [DEFAULT]/rabbit_hosts
# rabbit_hosts = $rabbit_host:$rabbit_port

# Connect over SSL for RabbitMQ. (boolean value)
# Deprecated group/name - [DEFAULT]/rabbit_use_ssl
# rabbit_use_ssl = false

# The RabbitMQ userid. (string value)
# Deprecated group/name - [DEFAULT]/rabbit_userid
# rabbit_userid = guest

# The RabbitMQ password. (string value)
# Deprecated group/name - [DEFAULT]/rabbit_password
# rabbit_password = guest

# The RabbitMQ login method. (string value)
# Deprecated group/name - [DEFAULT]/rabbit_login_method
```















```
# Location to store DHCP server config files
# dhcp_confs = $state_path/dhcp

# Domain to use for building the hostnames
# dhcp_domain = openstacklocal

# Override the default dnsmasq settings with this file
# dnsmasq_config_file =

# Comma-separated list of DNS servers which will be used by dnsmasq
# as forwarders.
# dnsmasq_dns_servers =

# Limit number of leases to prevent a denial-of-service.
# dnsmasq_lease_max = 16777216

# Location to DHCP lease relay UNIX domain socket
# dhcp_lease_relay_socket = $state_path/dhcp/lease_relay

# Use broadcast in DHCP replies
# dhcp_broadcast_reply = False

# dhcp_delete_namespaces, which is false by default, can be set to True if
# namespaces can be deleted cleanly on the host running the dhcp agent.
# Do not enable this until you understand the problem with the Linux iproute
# utility mentioned in https://bugs.launchpad.net/neutron/+bug/1052535 and
# you are sure that your version of iproute does not suffer from the problem.
# If True, namespaces will be deleted when a dhcp server is disabled.
# dhcp_delete_namespaces = False

# Timeout for ovs-vsctl commands.
# If the timeout expires, ovs commands will fail with ALARMCLOCK error.
# ovs_vsctl_timeout = 10
```

## **I3\_agent.ini**

```
[DEFAULT]
# Show debugging output in log (sets DEBUG log level output)
# debug = False

# L3 requires that an interface driver be set. Choose the one that best
# matches your plugin.
# interface_driver =

# Example of interface_driver option for OVS based plugins (OVS, Ryu, NEC)
# that supports L3 agent
# interface_driver = neutron.agent.linux.interface.OVSInterfaceDriver

# Use veth for an OVS interface or not.
# Support kernels with limited namespace support
# (e.g. RHEL 6.5) so long as ovs_use_veth is set to True.
# ovs_use_veth = False

# Example of interface_driver option for LinuxBridge
# interface_driver = neutron.agent.linux.interface.BridgeInterfaceDriver

# Allow overlapping IP (Must have kernel build with CONFIG_NET_NS=y and
```







# New, updated and deprecated options in Kilo for OpenStack Networking

**Table 9.68. New options**

Option = default value	(Type) Help string
[DEFAULT] advertise_mtu = False	(BoolOpt) If True, effort is made to advertise MTU settings to VMs via network methods (DHCP and RA MTU options) when the network's preferred MTU is known.
[DEFAULT] allow_automatic_dhcp_failover = True	(BoolOpt) Automatically remove networks from offline DHCP agents.
[DEFAULT] conn_idle_timeout = 900	(IntOpt) Reconnect connection to nsx if not used within this amount of time.
[DEFAULT] default_ipv4_subnet_pool = None	(StrOpt) Default IPv4 subnet-pool to be used for automatic subnet CIDR allocation
[DEFAULT] default_ipv6_subnet_pool = None	(StrOpt) Default IPv6 subnet-pool to be used for automatic subnet CIDR allocation
[DEFAULT] dhcp_broadcast_reply = False	(BoolOpt) Use broadcast in DHCP replies
[DEFAULT] dhcp_load_type = networks	(StrOpt) 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
[DEFAULT] enable_services_on_agents_with_admin_state_down = False	(BoolOpt) Enable services on an agent with admin_state_up False. If this option is False, when admin_state_up of an agent is turned False, services on it will be disabled. Agents with admin_state_up False are not selected for automatic scheduling regardless of this option. But manual scheduling to such agents is available if this option is True.
[DEFAULT] external_ingress_mark = 0x2	(StrOpt) Iptables mangle mark used to mark ingress from external network
[DEFAULT] ipv6_gateway =	(StrOpt) 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 ipv6_gateway, 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.
[DEFAULT] log-config-append = None	(StrOpt) 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.





















# 10. Object Storage

## Table of Contents

Introduction to Object Storage .....	593
Object Storage general service configuration .....	593
Object server configuration .....	595
Object expirer configuration .....	605
Container server configuration .....	609
Container sync realms configuration .....	617
Container reconciler configuration .....	618
Account server configuration .....	621
Proxy server configuration .....	627
Proxy server memcache configuration .....	647
Rsyncd configuration .....	647
Configure Object Storage features .....	648
New, updated and deprecated options in Kilo for OpenStack Object Storage .....	664

OpenStack Object Storage uses multiple configuration files for multiple services and background daemons, and **paste.deploy** to manage server configurations. Default configuration options appear in the [DEFAULT] section. You can override the default values by setting values in the other sections.

## Introduction to Object Storage

Object Storage is a robust, highly scalable and fault tolerant storage platform for unstructured data such as objects. Objects are stored bits, accessed through a RESTful, HTTP-based interface. You cannot access data at the block or file level. Object Storage is commonly used to archive and back up data, with use cases in virtual machine image, photo, video and music storage.

Object Storage provides a high degree of availability, throughput, and performance with its scale out architecture. Each object is replicated across multiple servers, residing within the same data center or across data centers, which mitigates the risk of network and hardware failure. In the event of hardware failure, Object Storage will automatically copy objects to a new location to ensure that there are always three copies available. Object Storage is an eventually consistent distributed storage platform; it sacrifices consistency for maximum availability and partition tolerance. Object Storage enables you to create a reliable platform by using commodity hardware and inexpensive storage.

For more information, review the key concepts in the developer documentation at [docs.openstack.org/developer/swift/](http://docs.openstack.org/developer/swift/).

## Object Storage general service configuration

Most Object Storage services fall into two categories, Object Storage's WSGI servers and background daemons.

Object Storage uses paste.deploy to manage server configurations. Read more at <http://pythonpaste.org/deploy/>.

Default configuration options are set in the `[DEFAULT]` section, and any options specified there can be overridden in any of the other sections when the syntax `set option_name = value` is in place.

Configuration for servers and daemons can be expressed together in the same file for each type of server, or separately. If a required section for the service trying to start is missing, there will be an error. Sections not used by the service are ignored.

Consider the example of an Object Storage node. By convention configuration for the `object-server`, `object-updater`, `object-replicator`, and `object-auditor` exist in a single file `/etc/swift/object-server.conf`:

```
[DEFAULT]

[pipeline:main]
pipeline = object-server

[app:object-server]
use = egg:swift#object

[object-replicator]
reclaim_age = 259200

[object-updater]

[object-auditor]
```

Object Storage services expect a configuration path as the first argument:

```
$ swift-object-auditor
Usage: swift-object-auditor CONFIG [options]

Error: missing config path argument
```

If you omit the `object-auditor` section, this file cannot be used as the configuration path when starting the `swift-object-auditor` daemon:

```
$ swift-object-auditor /etc/swift/object-server.conf
Unable to find object-auditor config section in /etc/swift/object-server.conf
```

If the configuration path is a directory instead of a file, all of the files in the directory with the file extension `".conf"` will be combined to generate the configuration object which is delivered to the Object Storage service. This is referred to generally as "directory-based configuration".

Directory-based configuration leverages ConfigParser's native multi-file support. Files ending in `".conf"` in the given directory are parsed in lexicographical order. File names starting with `'.'` are ignored. A mixture of file and directory configuration paths is not supported - if the configuration path is a file, only that file will be parsed.

The Object Storage service management tool `swift-init` has adopted the convention of looking for `/etc/swift/{type}-server.conf.d/` if the file `/etc/swift/{type}-server.conf` file does not exist.

When using directory-based configuration, if the same option under the same section appears more than once in different files, the last value parsed is said to override previous occurrences. You can ensure proper override precedence by prefixing the files in the configuration directory with numerical values, as in the following example file layout:

```

/etc/swift/
  default.base
  object-server.conf.d/
    000_default.conf -> ../default.base
    001_default-override.conf
    010_server.conf
    020_replicator.conf
    030_updater.conf
    040_auditor.conf
  
```

You can inspect the resulting combined configuration object using the **swift-config** command-line tool.

All the services of an Object Store deployment share a common configuration in the [swift-hash] section of the /etc/swift/swift.conf file. The swift\_hash\_path\_suffix and swift\_hash\_path\_prefix values must be identical on all the nodes.

**Table 10.1. Description of configuration options for [swift-hash] in swift.conf**

Configuration option = Default value	Description
swift_hash_path_prefix = <i>changeme</i>	A prefix used by hash_path to offer a bit more security when generating hashes for paths. It simply appends this value to all paths; if someone knows this suffix, it's easier for them to guess the hash a path will end up with. New installations are advised to set this parameter to a random secret, which would not be disclosed outside the organization. The same secret needs to be used by all swift servers of the same cluster. Existing installations should set this parameter to an empty string.
swift_hash_path_suffix = <i>changeme</i>	A suffix used by hash_path to offer a bit more security when generating hashes for paths. It simply appends this value to all paths; if someone knows this suffix, it's easier for them to guess the hash a path will end up with. New installations are advised to set this parameter to a random secret, which would not be disclosed outside the organization. The same secret needs to be used by all swift servers of the same cluster. Existing installations should set this parameter to an empty string.

## Object server configuration

Find an example object server configuration at `etc/object-server.conf-sample` in the source code repository.

The available configuration options are:

**Table 10.2. Description of configuration options for [DEFAULT] in `object-server.conf`**

Configuration option = Default value	Description
<code>backlog = 4096</code>	Maximum number of allowed pending TCP connections
<code>bind_ip = 0.0.0.0</code>	IP Address for server to bind to
<code>bind_port = 6000</code>	Port for server to bind to
<code>bind_timeout = 30</code>	Seconds to attempt bind before giving up
<code>client_timeout = 60</code>	Timeout to read one chunk from a client external services
<code>conn_timeout = 0.5</code>	Connection timeout to external services
<code>devices = /srv/node</code>	Parent directory of where devices are mounted
<code>disable_fallocate = false</code>	Disable "fast fail" fallocate checks if the underlying filesystem does not support it.
<code>disk_chunk_size = 65536</code>	Size of chunks to read/write to disk
<code>eventlet_debug = false</code>	If true, turn on debug logging for eventlet
<code>expiring_objects_account_name = expiring_objects</code>	No help text available for this option.
<code>expiring_objects_container_divisor = 86400</code>	No help text available for this option.
<code>fallocate_reserve = 0</code>	You can set <code>fallocate_reserve</code> to the number of bytes you'd like <code>fallocate</code> to reserve, whether there is space for the given file size or not. This is useful for systems that behave badly when they completely run out of space; you can make the services pretend they're out of space early. server. For most cases, this should be <code>`egg:swift#object`</code> .
<code>log_address = /dev/log</code>	Location where syslog sends the logs to
<code>log_custom_handlers =</code>	Comma-separated list of functions to call to setup custom log handlers.
<code>log_facility = LOG_LOCAL0</code>	Syslog log facility
<code>log_level = INFO</code>	Logging level
<code>log_max_line_length = 0</code>	Caps the length of log lines to the value given; no limit if set to 0, the default.
<code>log_name = swift</code>	Label used when logging
<code>log_statsd_default_sample_rate = 1.0</code>	Defines the probability of sending a sample for any given event or timing measurement.
<code>log_statsd_host = localhost</code>	If not set, the StatsD feature is disabled.
<code>log_statsd_metric_prefix =</code>	Value will be prepended to every metric sent to the StatsD server.
<code>log_statsd_port = 8125</code>	Port value for the StatsD server.
<code>log_statsd_sample_rate_factor = 1.0</code>	Not recommended to set this to a value less than 1.0, if frequency of logging is too high, tune the <code>log_statsd_default_sample_rate</code> instead.
<code>log_udp_host =</code>	If not set, the UDP receiver for syslog is disabled.
<code>log_udp_port = 514</code>	Port value for UDP receiver, if enabled.
<code>max_clients = 1024</code>	Maximum number of clients one worker can process simultaneously Lowering the number of clients handled per worker, and raising the number of workers can lessen the impact that a CPU intensive, or blocking, request can have on other requests served by the same worker. If the maximum number of clients is set to one, then a given worker will not perform another call while processing, allowing other workers a chance to process it.
<code>mount_check = true</code>	Whether or not check if the devices are mounted to prevent accidentally writing to the root device

Configuration option = Default value	Description
<code>network_chunk_size = 65536</code>	Size of chunks to read/write over the network
<code>node_timeout = 3</code>	Request timeout to external services
<code>swift_dir = /etc/swift</code>	Swift configuration directory
<code>user = swift</code>	User to run as
<code>workers = auto</code>	a much higher value, one can reduce the impact of slow file system operations in one request from negatively impacting other requests.

**Table 10.3. Description of configuration options for [ `app-object-server` ] in `object-server.conf`**

Configuration option = Default value	Description
<code>allowed_headers = Content-Disposition, Content-Encoding, X-Delete-At, X-Object-Manifest, X-Static-Large-Object</code>	Comma-separated list of headers that can be set in meta-data of an object
<code>auto_create_account_prefix = .</code>	Prefix to use when automatically creating accounts
<code>keep_cache_private = false</code>	Allow non-public objects to stay in kernel's buffer cache
<code>keep_cache_size = 5242880</code>	Largest object size to keep in buffer cache
<code>max_upload_time = 86400</code>	Maximum time allowed to upload an object
<code>mb_per_sync = 512</code>	On PUT requests, sync file every n MB
<code>replication_concurrency = 4</code>	Set to restrict the number of concurrent incoming REPLICATION requests; set to 0 for unlimited
<code>replication_failure_ratio = 1.0</code>	If the value of failures / successes of REPLICATION subrequests exceeds this ratio, the overall REPLICATION request will be aborted
<code>replication_failure_threshold = 100</code>	The number of subrequest failures before the <code>replication_failure_ratio</code> is checked
<code>replication_lock_timeout = 15</code>	Number of seconds to wait for an existing replication device lock before giving up.
<code>replication_one_per_device = True</code>	Restricts incoming REPLICATION requests to one per device, <code>replication_concurrency</code> above allowing. This can help control I/O to each device, but you may wish to set this to <code>False</code> to allow multiple REPLICATION requests (up to the above <code>replication_concurrency</code> setting) per device.
<code>replication_server = false</code>	If defined, tells server how to handle replication verbs in requests. When set to <code>True</code> (or 1), only replication verbs will be accepted. When set to <code>False</code> , replication verbs will be rejected. When undefined, server will accept any verb in the request.
<code>set log_address = /dev/log</code>	Location where syslog sends the logs to
<code>set log_facility = LOG_LOCAL0</code>	Syslog log facility
<code>set log_level = INFO</code>	Log level
<code>set log_name = object-server</code>	Label to use when logging
<code>set log_requests = true</code>	Whether or not to log requests
<code>slow = 0</code>	If > 0, Minimum time in seconds for a PUT or DELETE request to complete
<code>splice = no</code>	No help text available for this option.
<code>threads_per_disk = 0</code>	Size of the per-disk thread pool used for performing disk I/O. The default of 0 means to not use a per-disk thread pool. It is recommended to keep this value small, as large values can result in high read latencies due to large queue depths. A good starting point is 4 threads per disk.

Configuration option = Default value	Description
<code>use = egg:swift#object</code>	Entry point of paste.deploy in the server

**Table 10.4. Description of configuration options for `[pipeline-main]` in `object-server.conf`**

Configuration option = Default value	Description
<code>pipeline = healthcheck recon object-server</code>	No help text available for this option.

**Table 10.5. Description of configuration options for `[object-replicator]` in `object-server.conf`**

Configuration option = Default value	Description
<code>concurrency = 1</code>	Number of replication workers to spawn
<code>daemonize = on</code>	Whether or not to run replication as a daemon
<code>handoff_delete = auto</code>	By default handoff partitions will be removed when it has successfully replicated to all the canonical nodes. If set to an integer <code>n</code> , it will remove the partition if it is successfully replicated to <code>n</code> nodes. The default setting should not be changed, except for extreme situations. This uses what's set here, or what's set in the DEFAULT section, or 10 (though other sections use 3 as the final default).
<code>handoffs_first = False</code>	If set to True, partitions that are not supposed to be on the node will be replicated first. The default setting should not be changed, except for extreme situations.
<code>http_timeout = 60</code>	Maximum duration for an HTTP request
<code>lockup_timeout = 1800</code>	Attempts to kill all workers if nothing replications for <code>lockup_timeout</code> seconds
<code>log_address = /dev/log</code>	Location where syslog sends the logs to
<code>log_facility = LOG_LOCAL0</code>	Syslog log facility
<code>log_level = INFO</code>	Logging level
<code>log_name = object-replicator</code>	Label used when logging
<code>node_timeout = &lt;whatever's in the DEFAULT section or 10&gt;</code>	Request timeout to external services
<code>reclaim_age = 604800</code>	Time elapsed in seconds before an object can be reclaimed
<code>recon_cache_path = /var/cache/swift</code>	Directory where stats for a few items will be stored
<code>ring_check_interval = 15</code>	How often (in seconds) to check the ring
<code>rsync_bwlimit = 0</code>	No help text available for this option.
<code>rsync_error_log_line_length = 0</code>	No help text available for this option.
<code>rsync_io_timeout = 30</code>	Passed to rsync for a max duration (seconds) of an I/O op
<code>rsync_timeout = 900</code>	Max duration (seconds) of a partition rsync
<code>run_pause = 30</code>	Time in seconds to wait between replication passes
<code>stats_interval = 300</code>	Interval in seconds between logging replication statistics
<code>sync_method = rsync</code>	No help text available for this option.
<code>vm_test_mode = no</code>	Indicates that you are using a VM environment

**Table 10.6. Description of configuration options for `[object-updater]` in `object-server.conf`**

Configuration option = Default value	Description
<code>concurrency = 1</code>	Number of replication workers to spawn

Configuration option = Default value	Description
<code>interval = 300</code>	Minimum time for a pass to take
<code>log_address = /dev/log</code>	Location where syslog sends the logs to
<code>log_facility = LOG_LOCAL0</code>	Syslog log facility
<code>log_level = INFO</code>	Logging level
<code>log_name = object-updater</code>	Label used when logging
<code>node_timeout = &lt;whatever's in the DEFAULT section or 10&gt;</code>	Request timeout to external services
<code>recon_cache_path = /var/cache/swift</code>	Directory where stats for a few items will be stored
<code>slowdown = 0.01</code>	Time in seconds to wait between objects

**Table 10.7. Description of configuration options for [object-auditor] in object-server.conf**

Configuration option = Default value	Description
<code>bytes_per_second = 1000000</code>	Maximum bytes audited per second. Should be tuned according to individual system specs. 0 is unlimited. mounted to prevent accidentally writing to the root device process simultaneously (it will actually accept(2) N + 1). Setting this to one (1) will only handle one request at a time, without accepting another request concurrently. By increasing the number of workers to a much higher value, one can reduce the impact of slow file system operations in one request from negatively impacting other requests. underlying filesystem does not support it. to setup custom log handlers. bytes you'd like fallocate to reserve, whether there is space for the given file size or not. This is useful for systems that behave badly when they completely run out of space; you can make the services pretend they're out of space early. container server. For most cases, this should be `egg:swift#container`.
<code>concurrency = 1</code>	Number of replication workers to spawn
<code>disk_chunk_size = 65536</code>	Size of chunks to read/write to disk
<code>files_per_second = 20</code>	Maximum files audited per second. Should be tuned according to individual system specs. 0 is unlimited.
<code>log_address = /dev/log</code>	Location where syslog sends the logs to
<code>log_facility = LOG_LOCAL0</code>	Syslog log facility
<code>log_level = INFO</code>	Logging level
<code>log_name = object-auditor</code>	Label used when logging
<code>log_time = 3600</code>	Frequency of status logs in seconds.
<code>object_size_stats =</code>	No help text available for this option.
<code>recon_cache_path = /var/cache/swift</code>	Directory where stats for a few items will be stored
<code>zero_byte_files_per_second = 50</code>	Maximum zero byte files audited per second.

**Table 10.8. Description of configuration options for [filter-healthcheck] in object-server.conf**

Configuration option = Default value	Description
<code>disable_path =</code>	No help text available for this option.
<code>use = egg:swift#healthcheck</code>	Entry point of paste.deploy in the server









```
# You can override the default log routing for this app here (don't use set!):
# log_name = object-replicator
# log_facility = LOG_LOCAL0
# log_level = INFO
# log_address = /dev/log
#
# vm_test_mode = no
# daemonize = on
# run_pause = 30
# concurrency = 1
# stats_interval = 300
#
# The sync method to use; default is rsync but you can use ssync to try the
# EXPERIMENTAL all-swift-code-no-rsync-callouts method. Once ssync is verified
# as having performance comparable to, or better than, rsync, we plan to
# deprecate rsync so we can move on with more features for replication.
# sync_method = rsync
#
# max duration of a partition rsync
# rsync_timeout = 900
#
# bandwidth limit for rsync in kB/s. 0 means unlimited
# rsync_bwlimit = 0
#
# passed to rsync for io op timeout
# rsync_io_timeout = 30
#
# node_timeout = <whatever's in the DEFAULT section or 10>
# max duration of an http request; this is for REPLICATE finalization calls
# and
# so should be longer than node_timeout
# http_timeout = 60
#
# attempts to kill all workers if nothing replicates for lockup_timeout
# seconds
# lockup_timeout = 1800
#
# The replicator also performs reclamation
# reclaim_age = 604800
#
# ring_check_interval = 15
# recon_cache_path = /var/cache/swift
#
# limits how long rsync error log lines are
# 0 means to log the entire line
# rsync_error_log_line_length = 0
#
# handoffs_first and handoff_delete are options for a special case
# such as disk full in the cluster. These two options SHOULD NOT BE
# CHANGED, except for such an extreme situations. (e.g. disks filled up
# or are about to fill up. Anyway, DO NOT let your drives fill up)
# handoffs_first is the flag to replicate handoffs prior to canonical
# partitions. It allows to force syncing and deleting handoffs quickly.
# If set to a True value(e.g. "True" or "1"), partitions
# that are not supposed to be on the node will be replicated first.
# handoffs_first = False
#
# handoff_delete is the number of replicas which are ensured in swift.
# If the number less than the number of replicas is set, object-replicator
# could delete local handoffs even if all replicas are not ensured in the
```

















































```

# log_max_line_length = 0
#
# comma separated list of functions to call to setup custom log handlers.
# functions get passed: conf, name, log_to_console, log_route, fmt, logger,
# adapted_logger
# log_custom_handlers =
#
# If set, log_udp_host will override log_address
# log_udp_host =
# log_udp_port = 514
#
# You can enable StatsD logging here:
# log_statsd_host = localhost
# log_statsd_port = 8125
# log_statsd_default_sample_rate = 1.0
# log_statsd_sample_rate_factor = 1.0
# log_statsd_metric_prefix =
#
# If you don't mind the extra disk space usage in overhead, you can turn this
# on to preallocate disk space with SQLite databases to decrease
# fragmentation.
# db_preallocation = off
#
# eventlet_debug = false
#
# You can set fallocate_reserve to the number of bytes you'd like fallocate to
# reserve, whether there is space for the given file size or not.
# fallocate_reserve = 0

[pipeline:main]
pipeline = healthcheck recon account-server

[app:account-server]
use = egg:swift#account
# You can override the default log routing for this app here:
# set log_name = account-server
# set log_facility = LOG_LOCAL0
# set log_level = INFO
# set log_requests = true
# set log_address = /dev/log
#
# auto_create_account_prefix = .
#
# Configure parameter for creating specific server
# To handle all verbs, including replication verbs, do not specify
# "replication_server" (this is the default). To only handle replication,
# set to a True value (e.g. "True" or "1"). To handle only non-replication
# verbs, set to "False". Unless you have a separate replication network, you
# should not specify any value for "replication_server".
# replication_server = false

[filter:healthcheck]
use = egg:swift#healthcheck
# An optional filesystem path, which if present, will cause the healthcheck
# URL to return "503 Service Unavailable" with a body of "DISABLED BY FILE"
# disable_path =

[filter:recon]
use = egg:swift#recon
# recon_cache_path = /var/cache/swift

```

```
[account-replicator]
# You can override the default log routing for this app here (don't use set!):
# log_name = account-replicator
# log_facility = LOG_LOCAL0
# log_level = INFO
# log_address = /dev/log
#
# vm_test_mode = no
# per_diff = 1000
# max_diffs = 100
# concurrency = 8
# interval = 30
#
# How long without an error before a node's error count is reset. This will
# also be how long before a node is reenabled after suppression is triggered.
# error_suppression_interval = 60
#
# How many errors can accumulate before a node is temporarily ignored.
# error_suppression_limit = 10
#
# node_timeout = 10
# conn_timeout = 0.5
#
# The replicator also performs reclamation
# reclaim_age = 604800
#
# Time in seconds to wait between replication passes
# Note: if the parameter 'interval' is defined then it will be used in place
# of run_pause.
# run_pause = 30
#
# recon_cache_path = /var/cache/swift

[account-auditor]
# You can override the default log routing for this app here (don't use set!):
# log_name = account-auditor
# log_facility = LOG_LOCAL0
# log_level = INFO
# log_address = /dev/log
#
# Will audit each account at most once per interval
# interval = 1800
#
# log_facility = LOG_LOCAL0
# log_level = INFO
# accounts_per_second = 200
# recon_cache_path = /var/cache/swift

[account-reaper]
# You can override the default log routing for this app here (don't use set!):
# log_name = account-reaper
# log_facility = LOG_LOCAL0
# log_level = INFO
# log_address = /dev/log
#
# concurrency = 25
# interval = 3600
# node_timeout = 10
# conn_timeout = 0.5
```

```
#
# Normally, the reaper begins deleting account information for deleted
# accounts
# immediately; you can set this to delay its work however. The value is in
# seconds; 2592000 = 30 days for example.
# delay_reaping = 0
#
# If the account fails to be reaped due to a persistent error, the
# account reaper will log a message such as:
#     Account <name> has not been reaped since <date>
# You can search logs for this message if space is not being reclaimed
# after you delete account(s).
# Default is 2592000 seconds (30 days). This is in addition to any time
# requested by delay_reaping.
# reap_warn_after = 2592000

# Note: Put it at the beginning of the pipeline to profile all middleware. But
# it is safer to put this after healthcheck.
[filter:xprofile]
use = egg:swift#xprofile
# This option enable you to switch profilers which should inherit from python
# standard profiler. Currently the supported value can be 'cProfile',
# 'eventlet.green.profile' etc.
# profile_module = eventlet.green.profile
#
# This prefix will be used to combine process ID and timestamp to name the
# profile data file. Make sure the executing user has permission to write
# into this path (missing path segments will be created, if necessary).
# If you enable profiling in more than one type of daemon, you must override
# it with a unique value like: /var/log/swift/profile/account.profile
# log_filename_prefix = /tmp/log/swift/profile/default.profile
#
# the profile data will be dumped to local disk based on above naming rule
# in this interval.
# dump_interval = 5.0
#
# Be careful, this option will enable profiler to dump data into the file with
# time stamp which means there will be lots of files piled up in the
# directory.
# dump_timestamp = false
#
# This is the path of the URL to access the mini web UI.
# path = /__profile__
#
# Clear the data when the wsgi server shutdown.
# flush_at_shutdown = false
#
# unwind the iterator of applications
# unwind = false
```

## Proxy server configuration

Find an example proxy server configuration at `etc/proxy-server.conf-sample` in the source code repository.

The available configuration options are:

**Table 10.47. Description of configuration options for [DEFAULT] in proxy-server.conf**

Configuration option = Default value	Description
admin_key = <i>secret_admin_key</i>	to use for admin calls that are HMAC signed. Default is empty, which will disable admin calls to /info. the proxy server. For most cases, this should be `egg:swift#proxy`. request whenever it has to failover to a handoff node
backlog = 4096	Maximum number of allowed pending TCP connections
bind_ip = 0.0.0.0	IP Address for server to bind to
bind_port = 8080	Port for server to bind to
bind_timeout = 30	Seconds to attempt bind before giving up
cert_file = /etc/swift/proxy.crt	to the ssl .crt. This should be enabled for testing purposes only.
client_timeout = 60	Timeout to read one chunk from a client external services
cors_allow_origin =	is a list of hosts that are included with any CORS request by default and returned with the Access-Control-Allow-Origin header in addition to what the container has set. to call to setup custom log handlers. for eventlet the proxy server. For most cases, this should be `egg:swift#proxy`. request whenever it has to failover to a handoff node
disallowed_sections = <i>swift.valid_api_versions, container_quotas, tempurl</i>	No help text available for this option.
eventlet_debug = false	If true, turn on debug logging for eventlet
expiring_objects_account_name = <i>expiring_objects</i>	No help text available for this option.
expiring_objects_container_divisor = 86400	No help text available for this option.
expose_info = true	Enables exposing configuration settings via HTTP GET /info.
key_file = /etc/swift/proxy.key	to the ssl .key. This should be enabled for testing purposes only.
log_address = /dev/log	Location where syslog sends the logs to
log_custom_handlers =	Comma-separated list of functions to call to setup custom log handlers.
log_facility = LOG_LOCALO	Syslog log facility
log_headers = false	No help text available for this option.
log_level = INFO	Logging level
log_max_line_length = 0	Caps the length of log lines to the value given; no limit if set to 0, the default.
log_name = <i>swift</i>	Label used when logging
log_statsd_default_sample_rate = 1.0	Defines the probability of sending a sample for any given event or timing measurement.
log_statsd_host = localhost	If not set, the StatsD feature is disabled.
log_statsd_metric_prefix =	Value will be prepended to every metric sent to the StatsD server.
log_statsd_port = 8125	Port value for the StatsD server.
log_statsd_sample_rate_factor = 1.0	Not recommended to set this to a value less than 1.0, if frequency of logging is too high, tune the log_statsd_default_sample_rate instead.
log_udp_host =	If not set, the UDP receiver for syslog is disabled.
log_udp_port = 514	Port value for UDP receiver, if enabled.







**Table 10.54. Description of configuration options for [filter-container\_sync] in proxy-server.conf**

Configuration option = Default value	Description
allow_full_urls = true	No help text available for this option.
current = //REALM/CLUSTER	No help text available for this option.
use = egg:swift#container_sync	Entry point of paste.deploy in the server

**Table 10.55. Description of configuration options for [filter-dlo] in proxy-server.conf**

Configuration option = Default value	Description
max_get_time = 86400	No help text available for this option.
rate_limit_after_segment = 10	Rate limit the download of large object segments after this segment is downloaded.
rate_limit_segments_per_sec = 1	Rate limit large object downloads at this rate. contact for a normal request. You can use '*' replicas' at the end to have it use the number given times the number of replicas for the ring being used for the request. paste.deploy to use for auth. To use tempauth set to: `egg:swift#tempauth` each request
use = egg:swift#dlo	Entry point of paste.deploy in the server

**Table 10.56. Description of configuration options for [filter-gatekeeper] in proxy-server.conf**

Configuration option = Default value	Description
set log_address = /dev/log	Location where syslog sends the logs to
set log_facility = LOG_LOCAL0	Syslog log facility
set log_headers = false	If True, log headers in each request
set log_level = INFO	Log level
set log_name = gatekeeper	Label to use when logging
use = egg:swift#gatekeeper	Entry point of paste.deploy in the server

**Table 10.57. Description of configuration options for [filter-healthcheck] in proxy-server.conf**

Configuration option = Default value	Description
disable_path =	No help text available for this option.
use = egg:swift#healthcheck	Entry point of paste.deploy in the server

**Table 10.58. Description of configuration options for [filter-keystoneauth] in proxy-server.conf**

Configuration option = Default value	Description
allow_names_in_acls = true	The backwards compatible behavior can be disabled by setting this option to False.
allow_overrides = true	This option allows middleware higher in the WSGI pipeline to override auth processing, useful for middleware such as tempurl and formpost. If you know you are not going to use such middleware and you want a bit of extra security, you can set this to False.
default_domain_id = default	Name of the default domain. It is identified by its UUID, which by default has the value "default".











```
#
# How long without an error before a node's error count is reset. This will
# also be how long before a node is reenabled after suppression is triggered.
# error_suppression_interval = 60
#
# How many errors can accumulate before a node is temporarily ignored.
# error_suppression_limit = 10
#
# If set to 'true' any authorized user may create and delete accounts; if
# 'false' no one, even authorized, can.
# allow_account_management = false
#
# Set object_post_as_copy = false to turn on fast posts where only the
# metadata
# changes are stored anew and the original data file is kept in place. This
# makes for quicker posts; but since the container metadata isn't updated in
# this mode, features like container sync won't be able to sync posts.
# object_post_as_copy = true
#
# If set to 'true' authorized accounts that do not yet exist within the Swift
# cluster will be automatically created.
# account_autocreate = false
#
# If set to a positive value, trying to create a container when the account
# already has at least this maximum containers will result in a 403 Forbidden.
# Note: This is a soft limit, meaning a user might exceed the cap for
# recheck_account_existence before the 403s kick in.
# max_containers_per_account = 0
#
# This is a comma separated list of account hashes that ignore the
# max_containers_per_account cap.
# max_containers_whitelist =
#
# Comma separated list of Host headers to which the proxy will deny requests.
# deny_host_headers =
#
# Prefix used when automatically creating accounts.
# auto_create_account_prefix = .
#
# Depth of the proxy put queue.
# put_queue_depth = 10
#
# Storage nodes can be chosen at random (shuffle), by using timing
# measurements (timing), or by using an explicit match (affinity).
# Using timing measurements may allow for lower overall latency, while
# using affinity allows for finer control. In both the timing and
# affinity cases, equally-sorting nodes are still randomly chosen to
# spread load.
# The valid values for sorting_method are "affinity", "shuffle", and "timing".
# sorting_method = shuffle
#
# If the "timing" sorting_method is used, the timings will only be valid for
# the number of seconds configured by timing_expiry.
# timing_expiry = 300
#
# The maximum time (seconds) that a large object connection is allowed to
# last.
# max_large_object_get_time = 86400
#
# Set to the number of nodes to contact for a normal request. You can use
```



```
# single quote characters indicates an empty (blank) reseller prefix.
# reseller_prefix = AUTH

#
# The require_group parameter names a group that must be presented by
# either X-Auth-Token or X-Service-Token. Usually this parameter is
# used only with multiple reseller prefixes (e.g., SERVICE_require_group=
# blah).
# By default, no group is needed. Do not use .admin.
# require_group =

# The auth prefix will cause requests beginning with this prefix to be routed
# to the auth subsystem, for granting tokens, etc.
# auth_prefix = /auth/
# token_life = 86400
#
# This allows middleware higher in the WSGI pipeline to override auth
# processing, useful for middleware such as tempurl and formpost. If you know
# you're not going to use such middleware and you want a bit of extra
# security,
# you can set this to false.
# allow_overrides = true
#
# This specifies what scheme to return with storage urls:
# http, https, or default (chooses based on what the server is running as)
# This can be useful with an SSL load balancer in front of a non-SSL server.
# storage_url_scheme = default
#
# Lastly, you need to list all the accounts/users you want here. The format
# is:
# user_<account>_<user> = <key> [group] [group] [...] [storage_url]
# or if you want underscores in <account> or <user>, you can base64 encode
# them
# (with no equal signs) and use this format:
# user64_<account_b64>_<user_b64> = <key> [group] [group] [...]
# [storage_url]
# There are special groups of:
# .reseller_admin = can do anything to any account for this auth
# .admin = can do anything within the account
# If neither of these groups are specified, the user can only access
# containers
# that have been explicitly allowed for them by a .admin or .reseller_admin.
# The trailing optional storage_url allows you to specify an alternate url to
# hand back to the user upon authentication. If not specified, this defaults
# to
# $HOST/v1/<reseller_prefix>_<account> where $HOST will do its best to resolve
# to what the requester would need to use to reach this host.
# Here are example entries, required for running the tests:
user_admin_admin = admin .admin .reseller_admin
user_test_tester = testing .admin
user_test2_tester2 = testing2 .admin
user_test_tester3 = testing3
user_test5_tester5 = testing5 service

# To enable Keystone authentication you need to have the auth token
# middleware first to be configured. Here is an example below, please
# refer to the keystone's documentation for details about the
# different settings.
#
# You'll need to have as well the keystoneauth middleware enabled
```





```
# If the service_roles parameter is present, an X-Service-Token must be
# present in the request that when validated, grants at least one role listed
# in the parameter. The X-Service-Token may be scoped to any project.
# If there are several reseller prefix items, you can prefix the
# parameter so it applies only to those accounts (for example
# the parameter SERVICE_service_roles applies to the /v1/SERVICE_<project>
# path). If you omit the prefix, the option applies to all reseller
# prefix items. For the blank/empty prefix, prefix with ' ' (do not put
# underscore after the two single quote characters).
# By default, no service_roles are required.
# service_roles =
#
# For backwards compatibility, keystoneauth will match names in cross-tenant
# access control lists (ACLs) when both the requesting user and the tenant
# are in the default domain i.e the domain to which existing tenants are
# migrated. The default_domain_id value configured here should be the same as
# the value used during migration of tenants to keystone domains.
# default_domain_id = default
#
# For a new installation, or an installation in which keystone projects may
# move between domains, you should disable backwards compatible name matching
# in ACLs by setting allow_names_in_acls to false:
# allow_names_in_acls = true

[filter:healthcheck]
use = egg:swift#healthcheck
# An optional filesystem path, which if present, will cause the healthcheck
# URL to return "503 Service Unavailable" with a body of "DISABLED BY FILE".
# This facility may be used to temporarily remove a Swift node from a load
# balancer pool during maintenance or upgrade (remove the file to allow the
# node back into the load balancer pool).
# disable_path =

[filter:cache]
use = egg:swift#memcache
# You can override the default log routing for this filter here:
# set log_name = cache
# set log_facility = LOG_LOCAL0
# set log_level = INFO
# set log_headers = false
# set log_address = /dev/log
#
# If not set here, the value for memcache_servers will be read from
# memcache.conf (see memcache.conf-sample) or lacking that file, it will
# default to the value below. You can specify multiple servers separated with
# commas, as in: 10.1.2.3:11211,10.1.2.4:11211
# memcache_servers = 127.0.0.1:11211
#
# Sets how memcache values are serialized and deserialized:
# 0 = older, insecure pickle serialization
# 1 = json serialization but pickles can still be read (still insecure)
# 2 = json serialization only (secure and the default)
# If not set here, the value for memcache_serialization_support will be read
# from /etc/swift/memcache.conf (see memcache.conf-sample).
# To avoid an instant full cache flush, existing installations should
# upgrade with 0, then set to 1 and reload, then after some time (24 hours)
# set to 2 and reload.
# In the future, the ability to use pickle serialization will be removed.
# memcache_serialization_support = 2
#
```

```
# Sets the maximum number of connections to each memcached server per worker
# memcache_max_connections = 2
#
# More options documented in memcache.conf-sample

[filter:ratelimit]
use = egg:swift#ratelimit
# You can override the default log routing for this filter here:
# set log_name = ratelimit
# set log_facility = LOG_LOCAL0
# set log_level = INFO
# set log_headers = false
# set log_address = /dev/log
#
# clock_accuracy should represent how accurate the proxy servers' system
# clocks
# are with each other. 1000 means that all the proxies' clock are accurate to
# each other within 1 millisecond. No ratelimit should be higher than the
# clock accuracy.
# clock_accuracy = 1000
#
# max_sleep_time_seconds = 60
#
# log_sleep_time_seconds of 0 means disabled
# log_sleep_time_seconds = 0
#
# allows for slow rates (e.g. running up to 5 sec's behind) to catch up.
# rate_buffer_seconds = 5
#
# account_ratelimit of 0 means disabled
# account_ratelimit = 0

# DEPRECATED- these will continue to work but will be replaced
# by the X-Account-Sysmeta-Global-Write-Ratelimit flag.
# Please see ratelimiting docs for details.
# these are comma separated lists of account names
# account_whitelist = a,b
# account_blacklist = c,d

# with container_limit_x = r
# for containers of size x limit write requests per second to r. The
# container
# rate will be linearly interpolated from the values given. With the values
# below, a container of size 5 will get a rate of 75.
# container_ratelimit_0 = 100
# container_ratelimit_10 = 50
# container_ratelimit_50 = 20

# Similarly to the above container-level write limits, the following will
# limit
# container GET (listing) requests.
# container_listing_ratelimit_0 = 100
# container_listing_ratelimit_10 = 50
# container_listing_ratelimit_50 = 20

[filter:domain_remap]
use = egg:swift#domain_remap
# You can override the default log routing for this filter here:
# set log_name = domain_remap
# set log_facility = LOG_LOCAL0
```

```
# set log_level = INFO
# set log_headers = false
# set log_address = /dev/log
#
# storage_domain = example.com
# path_root = v1
# reseller_prefixes = AUTH

[filter:catch_errors]
use = egg:swift#catch_errors
# You can override the default log routing for this filter here:
# set log_name = catch_errors
# set log_facility = LOG_LOCAL0
# set log_level = INFO
# set log_headers = false
# set log_address = /dev/log

[filter:cname_lookup]
# Note: this middleware requires python-dnspython
use = egg:swift#cname_lookup
# You can override the default log routing for this filter here:
# set log_name = cname_lookup
# set log_facility = LOG_LOCAL0
# set log_level = INFO
# set log_headers = false
# set log_address = /dev/log
#
# Specify the storage_domain that match your cloud, multiple domains
# can be specified separated by a comma
# storage_domain = example.com
#
# lookup_depth = 1

# Note: Put staticweb just after your auth filter(s) in the pipeline
[filter:staticweb]
use = egg:swift#staticweb

# Note: Put tempurl before dlo, slo and your auth filter(s) in the pipeline
[filter:tempurl]
use = egg:swift#tempurl
# The methods allowed with Temp URLs.
# methods = GET HEAD PUT POST DELETE
#
# The headers to remove from incoming requests. Simply a whitespace delimited
# list of header names and names can optionally end with '*' to indicate a
# prefix match. incoming_allow_headers is a list of exceptions to these
# removals.
# incoming_remove_headers = x-timestamp
#
# The headers allowed as exceptions to incoming_remove_headers. Simply a
# whitespace delimited list of header names and names can optionally end with
# '*' to indicate a prefix match.
# incoming_allow_headers =
#
# The headers to remove from outgoing responses. Simply a whitespace delimited
# list of header names and names can optionally end with '*' to indicate a
# prefix match. outgoing_allow_headers is a list of exceptions to these
# removals.
# outgoing_remove_headers = x-object-meta-*
#
```





```
#
# Time limit on GET requests (seconds)
# max_get_time = 86400

# Note: Put after auth in the pipeline.
[filter:container-quotas]
use = egg:swift#container_quotas

# Note: Put after auth in the pipeline.
[filter:account-quotas]
use = egg:swift#account_quotas

[filter:gatekeeper]
use = egg:swift#gatekeeper
# You can override the default log routing for this filter here:
# set log_name = gatekeeper
# set log_facility = LOG_LOCAL0
# set log_level = INFO
# set log_headers = false
# set log_address = /dev/log

[filter:container_sync]
use = egg:swift#container_sync
# Set this to false if you want to disallow any full url values to be set for
# any new X-Container-Sync-To headers. This will keep any new full urls from
# coming in, but won't change any existing values already in the cluster.
# Updating those will have to be done manually, as knowing what the true realm
# endpoint should be cannot always be guessed.
# allow_full_urls = true
# Set this to specify this clusters //realm/cluster as "current" in /info
# current = //REALM/CLUSTER

# Note: Put it at the beginning of the pipeline to profile all middleware. But
# it is safer to put this after catch_errors, gatekeeper and healthcheck.
[filter:xprofile]
use = egg:swift#xprofile
# This option enable you to switch profilers which should inherit from python
# standard profiler. Currently the supported value can be 'cProfile',
# 'eventlet.green.profile' etc.
# profile_module = eventlet.green.profile
#
# This prefix will be used to combine process ID and timestamp to name the
# profile data file. Make sure the executing user has permission to write
# into this path (missing path segments will be created, if necessary).
# If you enable profiling in more than one type of daemon, you must override
# it with an unique value like: /var/log/swift/profile/proxy.profile
# log_filename_prefix = /tmp/log/swift/profile/default.profile
#
# the profile data will be dumped to local disk based on above naming rule
# in this interval.
# dump_interval = 5.0
#
# Be careful, this option will enable profiler to dump data into the file with
# time stamp which means there will be lots of files piled up in the
# directory.
# dump_timestamp = false
#
# This is the path of the URL to access the mini web UI.
# path = /__profile__
#
```































- DELETE Bucket
- GET Bucket (List Objects)
- PUT Bucket
- DELETE Object
- GET Object
- HEAD Object
- PUT Object
- PUT Object (Copy)

To use this middleware, first download the latest version from its repository to your proxy server(s).

```
$ git clone https://git.openstack.org/cgit/stackforge/swift3/
```

Then, install it using standard python mechanisms, such as:

```
# python setup.py install
```

Alternatively, if you have configured the Ubuntu Cloud Archive, you may use:

```
# apt-get install swift-python-s3
```

To add this middleware to your configuration, add the `swift3` middleware in front of the `swauth` middleware, and before any other middleware that looks at Object Storage requests (like rate limiting).

Ensure that your `proxy-server.conf` file contains `swift3` in the pipeline and the `[filter:swift3]` section, as shown below:

```
[pipeline:main]
pipeline = healthcheck cache swift3 swauth proxy-server

[filter:swift3]
use = egg:swift3#swift3
```

Next, configure the tool that you use to connect to the S3 API. For `S3curl`, for example, you must add your host IP information by adding your host IP to the `@endpoints` array (line 33 in `s3curl.pl`):

```
my @endpoints = ( '1.2.3.4');
```

Now you can send commands to the endpoint, such as:

```
$ ./s3curl.pl - 'a7811544507ebaf6c9a7a8804f47ea1c' -key 'a7d8e981-e296-d2ba-cb3b-dbd23159bd' -get - -s -v http://1.2.3.4:8080
```

To set up your client, ensure you are using the `ec2` credentials, which can be downloaded from the **API Endpoints** tab of the dashboard. The host should also point to the Object Storage node's hostname. It also will have to use the old-style calling format, and not the hostname-based container format. Here is an example client setup using the Python `boto` library on a locally installed all-in-one Object Storage installation.







<http://10.1.1.1:6000/sda1/2/a>

## New, updated and deprecated options in Kilo for OpenStack Object Storage

**Table 10.82. New options**

Option = default value	(Type) Help string
container-server.conf: [container-sync] conn_timeout = 5	(StrOpt) Connection timeout to external services
container-server.conf: [container-sync] internal_client_conf_path = /etc/swift/internal-client.conf	(StrOpt) No help text available for this option.
container-server.conf: [container-sync] request_tries = 3	(StrOpt) No help text available for this option.
drive-audit.conf: [drive-audit] log_to_console = False	(StrOpt) No help text available for this option.
drive-audit.conf: [drive-audit] recon_cache_path = /var/cache/swift	(StrOpt) Directory where stats for a few items will be stored
drive-audit.conf: [drive-audit] unmount_failed_device = True	(StrOpt) No help text available for this option.
internal-client.conf: [DEFAULT] log_address = /dev/log	(StrOpt) Location where syslog sends the logs to
internal-client.conf: [DEFAULT] log_custom_handlers =	(StrOpt) Comma-separated list of functions to call to setup custom log handlers.
internal-client.conf: [DEFAULT] log_facility = LOG_LOCAL0	(StrOpt) Syslog log facility
internal-client.conf: [DEFAULT] log_level = INFO	(StrOpt) Logging level
internal-client.conf: [DEFAULT] log_name = swift	(StrOpt) Label used when logging
internal-client.conf: [DEFAULT] log_statsd_default_sample_rate = 1.0	(StrOpt) Defines the probability of sending a sample for any given event or timing measurement.
internal-client.conf: [DEFAULT] log_statsd_host = localhost	(StrOpt) If not set, the StatsD feature is disabled.
internal-client.conf: [DEFAULT] log_statsd_metric_prefix =	(StrOpt) Value will be prepended to every metric sent to the StatsD server.
internal-client.conf: [DEFAULT] log_statsd_port = 8125	(StrOpt) Port value for the StatsD server.
internal-client.conf: [DEFAULT] log_statsd_sample_rate_factor = 1.0	(StrOpt) Not recommended to set this to a value less than 1.0, if frequency of logging is too high, tune the log_statsd_default_sample_rate instead.
internal-client.conf: [DEFAULT] log_udp_host =	(StrOpt) If not set, the UDP receiver for syslog is disabled.
internal-client.conf: [DEFAULT] log_udp_port = 514	(StrOpt) Port value for UDP receiver, if enabled.
internal-client.conf: [DEFAULT] swift_dir = /etc/swift	(StrOpt) Swift configuration directory
internal-client.conf: [DEFAULT] user = swift	(StrOpt) User to run as
internal-client.conf: [app-proxy-server] use = egg:swift#proxy	(StrOpt) Entry point of paste.deploy in the server
internal-client.conf: [filter-cache] use = egg:swift#memcache	(StrOpt) Entry point of paste.deploy in the server
internal-client.conf: [filter-catch_errors] use = egg:swift#catch_errors	(StrOpt) Entry point of paste.deploy in the server
internal-client.conf: [filter-proxy-logging] use = egg:swift#proxy_logging	(StrOpt) Entry point of paste.deploy in the server
internal-client.conf: [pipeline-main] pipeline = catch_errors proxy-logging cache proxy-server	(StrOpt) No help text available for this option.
memcache.conf: [memcache] connect_timeout = 0.3	(StrOpt) Timeout in seconds (float) for connection
memcache.conf: [memcache] io_timeout = 2.0	(StrOpt) Timeout in seconds (float) for read and write
memcache.conf: [memcache] pool_timeout = 1.0	(StrOpt) Timeout in seconds (float) for pooled connection



Option = default value	(Type) Help string
memcache.conf: [memcache] tries = 3	(StrOpt) Number of servers to retry on failures getting a pooled connection
object-server.conf: [object-reconstructor] concurrency = 1	(StrOpt) Number of replication workers to spawn
object-server.conf: [object-reconstructor] daemonize = on	(StrOpt) Whether or not to run replication as a daemon
object-server.conf: [object-reconstructor] handoffs_first = False	(StrOpt) If set to True, partitions that are not supposed to be on the node will be replicated first. The default setting should not be changed, except for extreme situations.
object-server.conf: [object-reconstructor] http_timeout = 60	(StrOpt) Maximum duration for an HTTP request
object-server.conf: [object-reconstructor] lockup_timeout = 1800	(StrOpt) Attempts to kill all workers if nothing replications for lockup_timeout seconds
object-server.conf: [object-reconstructor] log_address = /dev/log	(StrOpt) Location where syslog sends the logs to
object-server.conf: [object-reconstructor] log_facility = LOG_LOCAL0	(StrOpt) Syslog log facility
object-server.conf: [object-reconstructor] log_level = INFO	(StrOpt) Logging level
object-server.conf: [object-reconstructor] log_name = object-reconstructor	(StrOpt) Label used when logging
object-server.conf: [object-reconstructor] node_timeout = 10	(StrOpt) Request timeout to external services
object-server.conf: [object-reconstructor] reclaim_age = 604800	(StrOpt) Time elapsed in seconds before an object can be reclaimed
object-server.conf: [object-reconstructor] recon_cache_path = /var/cache/swift	(StrOpt) Directory where stats for a few items will be stored
object-server.conf: [object-reconstructor] ring_check_interval = 15	(StrOpt) How often (in seconds) to check the ring
object-server.conf: [object-reconstructor] run_pause = 30	(StrOpt) Time in seconds to wait between replication passes
object-server.conf: [object-reconstructor] stats_interval = 300	(StrOpt) Interval in seconds between logging replication statistics
proxy-server.conf: [filter-authtoken] identity_uri = http://keystonehost:35357/	(StrOpt) No help text available for this option.
proxy-server.conf: [filter-keystoneauth] allow_overrides = true	(StrOpt) This option allows middleware higher in the WSGI pipeline to override auth processing, useful for middleware such as tempurl and formpost. If you know you are not going to use such middleware and you want a bit of extra security, you can set this to False.
proxy-server.conf: [filter-keystoneauth] is_admin = false	(StrOpt) If this option is set to True, it allows to give a user whose username is the same as the project name and who has any role in the project access rights elevated to be the same as if the user had one of the operator_roles. Note that the condition compares names rather than UUIDs. This option is deprecated. It is False by default.
proxy-server.conf: [filter-keystoneauth] reseller_prefix = AUTH	(StrOpt) The naming scope for the auth service. Swift
proxy-server.conf: [filter-keystoneauth] service_roles =	(StrOpt) When present, this option requires that the X-Service-Token header supplies a token from a user who has a role listed in service_roles. This parameter may be prefixed with an appropriate prefix.
proxy-server.conf: [filter-tempauth] require_group =	(StrOpt) No help text available for this option.
proxy-server.conf: [filter-tempauth] user_test5_tester5 = testing5 service	(StrOpt) No help text available for this option.
swift.conf: [storage-policy-0] policy_type = replication	(StrOpt) No help text available for this option.
swift.conf: [storage-policy-1] policy_type = replication	(StrOpt) No help text available for this option.

Option = default value	(Type) Help string
swift.conf: [storage-policy-2] ec_num_data_fragments = 10	(StrOpt) No help text available for this option.
swift.conf: [storage-policy-2] ec_num_parity_fragments = 4	(StrOpt) No help text available for this option.
swift.conf: [storage-policy-2] ec_object_segment_size = 1048576	(StrOpt) No help text available for this option.
swift.conf: [storage-policy-2] ec_type = jerasure_rs_vand	(StrOpt) No help text available for this option.
swift.conf: [storage-policy-2] name = deepfreeze10-4	(StrOpt) No help text available for this option.
swift.conf: [storage-policy-2] policy_type = erasure_coding	(StrOpt) No help text available for this option.
swift.conf: [swift-constraints] valid_api_versions = v0,v1,v2	(StrOpt) No help text available for this option.

**Table 10.83. New default values**

Option	Previous default value	New default value
object-server.conf: [app-object-server] keep_cache_size	5424880	5242880
proxy-server.conf: [DEFAULT] disallowed_sections	container_quotas, tempurl, bulk_delete.max_failed_deletes	swift.valid_api_versions, container_quotas, tempurl
proxy-server.conf: [app-proxy-server] swift_owner_headers	x-container-read, x-container-write, x-container-sync-key, x-container-sync-to, x-account-meta-temp-url-key, x-account-meta-temp-url-key-2, x-account-access-control	x-container-read, x-container-write, x-container-sync-key, x-container-sync-to, x-account-meta-temp-url-key, x-account-meta-temp-url-key-2, x-container-meta-temp-url-key, x-container-meta-temp-url-key-2, x-account-access-control
proxy-server.conf: [filter-authtoken] delay_auth_decision	1	False



Configuration option = Default value	Description
<code>check_revocations_for_cached = False</code>	(BoolOpt) If true, the revocation list will be checked for cached tokens. This requires that PKI tokens are configured on the identity server.
<code>delay_auth_decision = False</code>	(BoolOpt) Do not handle authorization requests within the middleware, but delegate the authorization decision to downstream WSGI components.
<code>enforce_token_bind = permissive</code>	(StrOpt) 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.
<code>hash_algorithms = md5</code>	(ListOpt) 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.
<code>http_connect_timeout = None</code>	(IntOpt) Request timeout value for communicating with Identity API server.
<code>http_request_max_retries = 3</code>	(IntOpt) How many times are we trying to reconnect when communicating with Identity API Server.
<code>identity_uri = None</code>	(StrOpt) Complete admin Identity API endpoint. This should specify the unversioned root endpoint e.g. <code>https://localhost:35357/</code>
<code>include_service_catalog = True</code>	(BoolOpt) (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.
<code>insecure = False</code>	(BoolOpt) Verify HTTPS connections.
<code>keyfile = None</code>	(StrOpt) Required if identity server requires client certificate
<code>memcache_pool_conn_get_timeout = 10</code>	(IntOpt) (Optional) Number of seconds that an operation will wait to get a memcache client connection from the pool.
<code>memcache_pool_dead_retry = 300</code>	(IntOpt) (Optional) Number of seconds memcached server is considered dead before it is tried again.
<code>memcache_pool_maxsize = 10</code>	(IntOpt) (Optional) Maximum total number of open connections to every memcached server.
<code>memcache_pool_socket_timeout = 3</code>	(IntOpt) (Optional) Socket timeout in seconds for communicating with a memcache server.
<code>memcache_pool_unused_timeout = 60</code>	(IntOpt) (Optional) Number of seconds a connection to memcached is held unused in the pool before it is closed.
<code>memcache_secret_key = None</code>	(StrOpt) (Optional, mandatory if <code>memcache_security_strategy</code> is defined) This string is used for key derivation.
<code>memcache_security_strategy = None</code>	(StrOpt) (Optional) If defined, indicate whether token data should be authenticated or authenticated and encrypted. Acceptable values are MAC or ENCRYPT. If MAC, token data is authenticated (with HMAC) in the cache. If ENCRYPT, token data is encrypted and authenticated in the







Configuration option = Default value	Description
<code>fatal_deprecations = False</code>	(BoolOpt) Enables or disables fatal status of deprecations.
<code>fatal_exception_format_errors = False</code>	(BoolOpt) Make exception message format errors fatal
<code>instance_format = "[instance: %(uuid)s] "</code>	(StrOpt) The format for an instance that is passed with the log message.
<code>instance_uuid_format = "[instance: %(uuid)s]"</code>	(StrOpt) The format for an instance UUID that is passed with the log message.
<code>log_config_append = None</code>	(StrOpt) 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.
<code>log_date_format = "%Y-%m-%d %H:%M:%S"</code>	(StrOpt) Format string for <code>%(asctime)s</code> in log records. Default: <code>%(default)s</code> .
<code>log_dir = None</code>	(StrOpt) (Optional) The base directory used for relative log-file paths.
<code>log_file = None</code>	(StrOpt) (Optional) Name of log file to output to. If no default is set, logging will go to stdout.
<code>log_format = None</code>	(StrOpt) DEPRECATED. A logging.Formatter log message format string which may use any of the available logging.LogRecord attributes. This option is deprecated. Please use <code>logging_context_format_string</code> and <code>logging_default_format_string</code> instead.
<code>logging_context_format_string = %(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [%(request_id)s %(user_identity)s] %(instance)s%(message)s</code>	(StrOpt) Format string to use for log messages with context.
<code>logging_debug_format_suffix = %(funcName)s %(pathname)s: %(lineno)d</code>	(StrOpt) Data to append to log format when level is DEBUG.
<code>logging_default_format_string = %(asctime)s.%(msecs)03d %(process)d %(levelname)s %(name)s [-] %(instance)s%(message)s</code>	(StrOpt) Format string to use for log messages without context.
<code>logging_exception_prefix = %(asctime)s.%(msecs)03d %(process)d TRACE %(name)s %(instance)s</code>	(StrOpt) Prefix each line of exception output with this format.
<code>publish_errors = False</code>	(BoolOpt) Enables or disables publication of error events.
<code>syslog_log_facility = LOG_USER</code>	(StrOpt) Syslog facility to receive log lines.
<code>use_syslog = False</code>	(BoolOpt) Use syslog for logging. Existing syslog format is DEPRECATED during I, and will change in J to honor RFC5424.
<code>use_syslog_rfc_format = False</code>	(BoolOpt) (Optional) Enables or disables syslog rfc5424 format for logging. If enabled, prefixes the MSG part of the syslog message with APP-NAME (RFC5424). The format without the APP-NAME is deprecated in I, and will be removed in J.
<code>use_stderr = True</code>	(BoolOpt) Log output to standard error.
<code>verbose = False</code>	(BoolOpt) Print more verbose output (set logging level to INFO instead of default WARNING level).

**Table 11.8. Description of oslo\_middleware configuration options**

Configuration option = Default value	Description
<code>[oslo_middleware]</code>	
<code>max_request_body_size = 114688</code>	(IntOpt) The maximum body size for each request, in bytes.





















**Configure Qpid**

Use these options to configure the Qpid messaging system for OpenStack Oslo RPC. Qpid is not the default messaging system, so you must enable it by setting the `rpc_backend` option in the `heat.conf` file:

```
rpc_backend=heat.openstack.common.rpc.impl_qpid
```

This critical option points the compute nodes to the Qpid broker (server). Set the `qpid_hostname` option to the host name where the broker runs in the `heat.conf` file.



**Note**

The `qpid_hostname` option accepts a host name or IP address value.

```
qpid_hostname = hostname.example.com
```

If the Qpid broker listens on a port other than the AMQP default of 5672, you must set the `qpid_port` option to that value:

```
qpid_port = 12345
```

If you configure the Qpid broker to require authentication, you must add a user name and password to the configuration:

```
qpid_username = username
qpid_password = password
```

By default, TCP is used as the transport. To enable SSL, set the `qpid_protocol` option:

```
qpid_protocol = ssl
```

Use these additional options to configure the Qpid messaging driver for OpenStack Oslo RPC. These options are used infrequently.

**Table 11.30. Description of Qpid configuration options**

Configuration option = Default value	Description
[oslo_messaging_qpid]	
<code>amqp_auto_delete = False</code>	(BoolOpt) Auto-delete queues in AMQP.
<code>amqp_durable_queues = False</code>	(BoolOpt) Use durable queues in AMQP.
<code>qpid_heartbeat = 60</code>	(IntOpt) Seconds between connection keepalive heartbeats.
<code>qpid_hostname = localhost</code>	(StrOpt) Qpid broker hostname.
<code>qpid_hosts = \$qpid_hostname:\$qpid_port</code>	(ListOpt) Qpid HA cluster host:port pairs.
<code>qpid_password =</code>	(StrOpt) Password for Qpid connection.
<code>qpid_port = 5672</code>	(IntOpt) Qpid broker port.
<code>qpid_protocol = tcp</code>	(StrOpt) Transport to use, either 'tcp' or 'ssl'.
<code>qpid_receiver_capacity = 1</code>	(IntOpt) The number of prefetched messages held by receiver.
<code>qpid_sasl_mechanisms =</code>	(StrOpt) Space separated list of SASL mechanisms to use for auth.
<code>qpid_tcp_nodelay = True</code>	(BoolOpt) Whether to disable the Nagle algorithm.
<code>qpid_topology_version = 1</code>	(IntOpt) The qpid topology version to use. Version 1 is what was originally used by impl_qpid. Version 2 includes some backwards-incompatible changes that allow bro-

Configuration option = Default value	Description
	ker federation to work. Users should update to version 2 when they are able to take everything down, as it requires a clean break.
qpid_username =	(StrOpt) Username for Qpid connection.
rpc_conn_pool_size = 30	(IntOpt) Size of RPC connection pool.

## Configure ZeroMQ

Use these options to configure the ZeroMQ messaging system for OpenStack Oslo RPC. ZeroMQ is not the default messaging system, so you must enable it by setting the `rpc_backend` option in the `heat.conf` file:

**Table 11.31. Description of ZeroMQ configuration options**

Configuration option = Default value	Description
[DEFAULT]	
rpc_zmq_bind_address = *	(StrOpt) ZeroMQ bind address. Should be a wildcard (*), an ethernet interface, or IP. The "host" option should point or resolve to this address.
rpc_zmq_contexts = 1	(IntOpt) Number of ZeroMQ contexts, defaults to 1.
rpc_zmq_host = localhost	(StrOpt) Name of this node. Must be a valid hostname, FQDN, or IP address. Must match "host" option, if running Nova.
rpc_zmq_ipc_dir = /var/run/openstack	(StrOpt) Directory for holding IPC sockets.
rpc_zmq_matchmaker = local	(StrOpt) MatchMaker driver.
rpc_zmq_port = 9501	(IntOpt) ZeroMQ receiver listening port.
rpc_zmq_topic_backlog = None	(IntOpt) Maximum number of ingress messages to locally buffer per topic. Default is unlimited.

## Configure messaging

Use these common options to configure the RabbitMQ, Qpid, and ZeroMq messaging drivers:

**Table 11.32. Description of AMQP configuration options**

Configuration option = Default value	Description
[DEFAULT]	
control_exchange = openstack	(StrOpt) The default exchange under which topics are scoped. May be overridden by an exchange name specified in the <code>transport_url</code> option.
default_notification_level = INFO	(StrOpt) Default notification level for outgoing notifications.
default_publisher_id = None	(StrOpt) Default publisher_id for outgoing notifications.
list_notifier_drivers = None	(MultiStrOpt) List of drivers to send notifications (DEPRECATED).
notification_driver = []	(MultiStrOpt) Driver or drivers to handle sending notifications.
notification_topics = notifications	(ListOpt) AMQP topic used for OpenStack notifications.
transport_url = None	(StrOpt) A URL representing the messaging driver to use and its full configuration. If not set, we fall back to the <code>rpc_backend</code> option and driver specific configuration.



<b>Option = default value</b>	<b>(Type) Help string</b>
[DEFAULT] default_software_config_transport = POLL_SERVER_CFN	(StrOpt) 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).
[DEFAULT] error_wait_time = 240	(IntOpt) Error wait time in seconds for stack action (ie. create or update).
[DEFAULT] fatal_exception_format_errors = False	(BoolOpt) Make exception message format errors fatal
[DEFAULT] log-config-append = None	(StrOpt) 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.
[DEFAULT] log-date-format = %Y-%m-%d %H:%M:%S	(StrOpt) Format string for %(asctime)s in log records. Default: %(default)s .
[DEFAULT] log-dir = None	(StrOpt) (Optional) The base directory used for relative – log-file paths.
[DEFAULT] log-file = None	(StrOpt) (Optional) Name of log file to output to. If no default is set, logging will go to stdout.
[DEFAULT] log-format = None	(StrOpt) DEPRECATED. A logging.Formatter log message format string which may use any of the available logging.LogRecord attributes. This option is deprecated. Please use logging_context_format_string and logging_default_format_string instead.
[DEFAULT] policy_dirs = ['policy.d']	(MultiStrOpt) 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.
[DEFAULT] stack_scheduler_hints = False	(BoolOpt) When this feature is enabled, scheduler hints identifying the heat stack context of a server resource are passed to the configured schedulers in nova, for server creates done using heat resource types 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 tuples, (stackresource name, stackname) with list[0] being (None, rootstackname), and heat_resource_name will be set to the resource's name.
[DEFAULT] syslog-log-facility = LOG_USER	(StrOpt) Syslog facility to receive log lines.
[DEFAULT] use-syslog = False	(BoolOpt) Use syslog for logging. Existing syslog format is DEPRECATED during I, and will change in J to honor RFC5424.
[DEFAULT] use-syslog-rfc-format = False	(BoolOpt) (Optional) Enables or disables syslog rfc5424 format for logging. If enabled, prefixes the MSG part of the syslog message with APP-NAME (RFC5424). The format without the APP-NAME is deprecated in I, and will be removed in J.

Option = default value	(Type) Help string
[clients_sahara] ca_file = None	(StrOpt) Optional CA cert file to use in SSL connections.
[clients_sahara] cert_file = None	(StrOpt) Optional PEM-formatted certificate chain file.
[clients_sahara] endpoint_type = None	(StrOpt) Type of endpoint in Identity service catalog to use for communication with the OpenStack service.
[clients_sahara] insecure = None	(BoolOpt) If set, then the server's certificate will not be verified.
[clients_sahara] key_file = None	(StrOpt) Optional PEM-formatted file that contains the private key.
[ec2authtoken] ca_file = None	(StrOpt) Optional CA cert file to use in SSL connections.
[ec2authtoken] cert_file = None	(StrOpt) Optional PEM-formatted certificate chain file.
[ec2authtoken] insecure = False	(BoolOpt) If set, then the server's certificate will not be verified.
[ec2authtoken] key_file = None	(StrOpt) Optional PEM-formatted file that contains the private key.
[oslo_messaging_amqp] allow_insecure_clients = False	(BoolOpt) Accept clients using either SSL or plain TCP
[oslo_messaging_amqp] broadcast_prefix = broadcast	(StrOpt) address prefix used when broadcasting to all servers
[oslo_messaging_amqp] container_name = None	(StrOpt) Name for the AMQP container
[oslo_messaging_amqp] group_request_prefix = unicast	(StrOpt) address prefix when sending to any server in group
[oslo_messaging_amqp] idle_timeout = 0	(IntOpt) Timeout for inactive connections (in seconds)
[oslo_messaging_amqp] server_request_prefix = exclusive	(StrOpt) address prefix used when sending to a specific server
[oslo_messaging_amqp] ssl_ca_file =	(StrOpt) CA certificate PEM file to verify server certificate
[oslo_messaging_amqp] ssl_cert_file =	(StrOpt) Identifying certificate PEM file to present to clients
[oslo_messaging_amqp] ssl_key_file =	(StrOpt) Private key PEM file used to sign cert_file certificate
[oslo_messaging_amqp] ssl_key_password = None	(StrOpt) Password for decrypting ssl_key_file (if encrypted)
[oslo_messaging_amqp] trace = False	(BoolOpt) Debug: dump AMQP frames to stdout
[oslo_messaging_qpid] amqp_auto_delete = False	(BoolOpt) Auto-delete queues in AMQP.
[oslo_messaging_qpid] amqp_durable_queues = False	(BoolOpt) Use durable queues in AMQP.
[oslo_messaging_qpid] qpid_heartbeat = 60	(IntOpt) Seconds between connection keepalive heartbeats.
[oslo_messaging_qpid] qpid_hostname = localhost	(StrOpt) Qpid broker hostname.
[oslo_messaging_qpid] qpid_hosts = \$qpid_hostname:\$qpid_port	(ListOpt) Qpid HA cluster host:port pairs.
[oslo_messaging_qpid] qpid_password =	(StrOpt) Password for Qpid connection.
[oslo_messaging_qpid] qpid_port = 5672	(IntOpt) Qpid broker port.
[oslo_messaging_qpid] qpid_protocol = tcp	(StrOpt) Transport to use, either 'tcp' or 'ssl'.
[oslo_messaging_qpid] qpid_receiver_capacity = 1	(IntOpt) The number of prefetched messages held by receiver.
[oslo_messaging_qpid] qpid_sasl_mechanisms =	(StrOpt) Space separated list of SASL mechanisms to use for auth.
[oslo_messaging_qpid] qpid_tcp_nodelay = True	(BoolOpt) Whether to disable the Nagle algorithm.
[oslo_messaging_qpid] qpid_topology_version = 1	(IntOpt) The qpid topology version to use. Version 1 is what was originally used by impl_qpid. Version 2 includes some backwards-incompatible changes that allow broker federation to work. Users should update to version 2

Option = default value	(Type) Help string
	when they are able to take everything down, as it requires a clean break.
[oslo_messaging_qpid] qpid_username =	(StrOpt) Username for Qpid connection.
[oslo_messaging_qpid] rpc_conn_pool_size = 30	(IntOpt) Size of RPC connection pool.
[oslo_messaging_rabbit] amqp_auto_delete = False	(BoolOpt) Auto-delete queues in AMQP.
[oslo_messaging_rabbit] amqp_durable_queues = False	(BoolOpt) Use durable queues in AMQP.
[oslo_messaging_rabbit] fake_rabbit = False	(BoolOpt) Deprecated, use rpc_backend=kombu+memory or rpc_backend=fake
[oslo_messaging_rabbit] heartbeat_rate = 2	(IntOpt) How often times during the heartbeat_timeout_threshold we check the heartbeat.
[oslo_messaging_rabbit] heartbeat_timeout_threshold = 0	(IntOpt) Number of seconds after which the Rabbit broker is considered down if heartbeat's keep-alive fails (0 disables the heartbeat, >0 enables it. Enabling heartbeats requires kombu>=3.0.7 and amqp>=1.4.0). EXPERIMENTAL
[oslo_messaging_rabbit] kombu_reconnect_delay = 1.0	(FloatOpt) How long to wait before reconnecting in response to an AMQP consumer cancel notification.
[oslo_messaging_rabbit] kombu_ssl_ca_certs =	(StrOpt) SSL certification authority file (valid only if SSL enabled).
[oslo_messaging_rabbit] kombu_ssl_certfile =	(StrOpt) SSL cert file (valid only if SSL enabled).
[oslo_messaging_rabbit] kombu_ssl_keyfile =	(StrOpt) SSL key file (valid only if SSL enabled).
[oslo_messaging_rabbit] kombu_ssl_version =	(StrOpt) 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.
[oslo_messaging_rabbit] rabbit_ha_queues = False	(BoolOpt) Use HA queues in RabbitMQ (x-ha-policy: all). If you change this option, you must wipe the RabbitMQ database.
[oslo_messaging_rabbit] rabbit_host = localhost	(StrOpt) The RabbitMQ broker address where a single node is used.
[oslo_messaging_rabbit] rabbit_hosts = \$rabbit_host:\$rabbit_port	(ListOpt) RabbitMQ HA cluster host:port pairs.
[oslo_messaging_rabbit] rabbit_login_method = AMQPLAIN	(StrOpt) The RabbitMQ login method.
[oslo_messaging_rabbit] rabbit_max_retries = 0	(IntOpt) Maximum number of RabbitMQ connection retries. Default is 0 (infinite retry count).
[oslo_messaging_rabbit] rabbit_password = guest	(StrOpt) The RabbitMQ password.
[oslo_messaging_rabbit] rabbit_port = 5672	(IntOpt) The RabbitMQ broker port where a single node is used.
[oslo_messaging_rabbit] rabbit_retry_backoff = 2	(IntOpt) How long to backoff for between retries when connecting to RabbitMQ.
[oslo_messaging_rabbit] rabbit_retry_interval = 1	(IntOpt) How frequently to retry connecting with RabbitMQ.
[oslo_messaging_rabbit] rabbit_use_ssl = False	(BoolOpt) Connect over SSL for RabbitMQ.
[oslo_messaging_rabbit] rabbit_userid = guest	(StrOpt) The RabbitMQ userid.
[oslo_messaging_rabbit] rabbit_virtual_host = /	(StrOpt) The RabbitMQ virtual host.
[oslo_messaging_rabbit] rpc_conn_pool_size = 30	(IntOpt) Size of RPC connection pool.
[oslo_middleware] max_request_body_size = 114688	(IntOpt) The maximum body size for each request, in bytes.
[profiler] profiler_enabled = False	(BoolOpt) If False fully disable profiling feature.
[profiler] trace_sqlalchemy = False	(BoolOpt) If False do not trace SQL requests.

**Table 11.36. New default values**

Option	Previous default value	New default value
[DEFAULT] auth_encryption_key	notgood but just long enough i think	notgood but just long enough i t
[DEFAULT] default_log_levels	amqp=WARN, amqplib=WARN, boto=WARN, qpid=WARN, sqlalchemy=WARN, suds=INFO, oslo.messaging=INFO, iso8601=WARN, requests.packages.urllib3.connectionpool=	amqp=WARN, amqplib=WARN, boto=WARN, qpid=WARN, sqlalchemy=WARN, suds=INFO, oslo.messaging=INFO, iso8601=WARN, requests.packages.urllib3.connectionpool=
[DEFAULT] deferred_auth_method	password	trusts
[DEFAULT] max_nested_stack_depth	3	5
[DEFAULT] num_engine_workers	1	4
[DEFAULT] plugin_dirs	/usr/lib64/heat, /usr/lib/heat	/usr/lib64/heat, /usr/lib/heat, /usr/local/lib/heat, /usr/local/lib64/heat
[DEFAULT] rpc_zmq_matchmaker	oslo.messaging_drivers.matchmaker.MatchMakerLocalhost	MatchMakerLocalhost
[DEFAULT] trusts_delegated_roles	heat_stack_owner	

**Table 11.37. Deprecated options**

Deprecated option	New Option
[DEFAULT] log-format	None
[DEFAULT] use-syslog	None
[DEFAULT] list_notifier_drivers	None

# 12. Telemetry

## Table of Contents

Telemetry sample configuration files .....	702
New, updated and deprecated options in Kilo for Telemetry .....	733

The Telemetry service collects measurements within OpenStack. Its various agents and services are configured in the `/etc/ceilometer/ceilometer.conf` file.

To install Telemetry, see the *OpenStack Installation Guide* for your distribution ([docs.openstack.org](http://docs.openstack.org)).

The following tables provide a comprehensive list of the Telemetry configuration options.

**Table 12.1. Description of alarm configuration options**

Configuration option = Default value	Description
[alarm]	
<code>evaluation_interval = 60</code>	(IntOpt) Period of evaluation cycle, should be >= than configured pipeline interval for collection of underlying metrics.
<code>evaluation_service = default</code>	(StrOpt) Driver to use for alarm evaluation service. DEPRECATED: "singleton" and "partitioned" alarm evaluator services will be removed in Kilo in favour of the default alarm evaluation service using tooz for partitioning.
<code>notifier_rpc_topic = alarm_notifier</code>	(StrOpt) The topic that ceilometer uses for alarm notifier messages.
<code>partition_rpc_topic = alarm_partition_coordination</code>	(StrOpt) The topic that ceilometer uses for alarm partition coordination messages. DEPRECATED: RPC-based partitioned alarm evaluation service will be removed in Kilo in favour of the default alarm evaluation service using tooz for partitioning.
<code>project_alarm_quota = None</code>	(IntOpt) Maximum number of alarms defined for a project.
<code>record_history = True</code>	(BoolOpt) Record alarm change events.
<code>rest_notifier_certificate_file =</code>	(StrOpt) SSL Client certificate for REST notifier.
<code>rest_notifier_certificate_key =</code>	(StrOpt) SSL Client private key for REST notifier.
<code>rest_notifier_max_retries = 0</code>	(IntOpt) Number of retries for REST notifier
<code>rest_notifier_ssl_verify = True</code>	(BoolOpt) Whether to verify the SSL Server certificate when calling alarm action.
<code>user_alarm_quota = None</code>	(IntOpt) Maximum number of alarms defined for a user.

**Table 12.2. Description of alarms configuration options**

Configuration option = Default value	Description
[alarms]	
<code>gnocchi_url = http://localhost:8041</code>	(StrOpt) URL to Gnocchi.



**Table 12.3. Description of AMQP configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>control_exchange = openstack</code>	(StrOpt) The default exchange under which topics are scoped. May be overridden by an exchange name specified in the <code>transport_url</code> option.
<code>notification_driver = []</code>	(MultiStrOpt) Driver or drivers to handle sending notifications.
<code>notification_topics = notifications</code>	(ListOpt) AMQP topic used for OpenStack notifications.
<code>transport_url = None</code>	(StrOpt) A URL representing the messaging driver to use and its full configuration. If not set, we fall back to the <code>rpc_backend</code> option and driver specific configuration.

**Table 12.4. Description of API configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>api_paste_config = api_paste.ini</code>	(StrOpt) Configuration file for WSGI definition of API.
<code>api_workers = 1</code>	(IntOpt) Number of workers for Ceilometer API server.
<code>event_pipeline_cfg_file = event_pipeline.yaml</code>	(StrOpt) Configuration file for event pipeline definition.
<code>pipeline_cfg_file = pipeline.yaml</code>	(StrOpt) Configuration file for pipeline definition.
<code>reserved_metadata_keys =</code>	(ListOpt) List of metadata keys reserved for metering use. And these keys are additional to the ones included in the namespace.
<code>reserved_metadata_length = 256</code>	(IntOpt) Limit on length of reserved metadata values.
<code>reserved_metadata_namespace = metering.</code>	(ListOpt) List of metadata prefixes reserved for metering use.
[api]	
<code>host = 0.0.0.0</code>	(StrOpt) The listen IP for the ceilometer API server.
<code>pecan_debug = False</code>	(BoolOpt) Toggle Pecan Debug Middleware.
<code>port = 8777</code>	(IntOpt) The port for the ceilometer API server.

**Table 12.5. Description of authorization configuration options**

Configuration option = Default value	Description
[service_credentials]	
<code>insecure = False</code>	(BoolOpt) Disables X.509 certificate validation when an SSL connection to Identity Service is established.
<code>os_auth_url = http://localhost:5000/v2.0</code>	(StrOpt) Auth URL to use for OpenStack service access.
<code>os_cacert = None</code>	(StrOpt) Certificate chain for SSL validation.
<code>os_endpoint_type = publicURL</code>	(StrOpt) Type of endpoint in Identity service catalog to use for communication with OpenStack services.
<code>os_password = admin</code>	(StrOpt) Password to use for OpenStack service access.
<code>os_region_name = None</code>	(StrOpt) Region name to use for OpenStack service endpoints.
<code>os_tenant_id =</code>	(StrOpt) Tenant ID to use for OpenStack service access.
<code>os_tenant_name = admin</code>	(StrOpt) Tenant name to use for OpenStack service access.
<code>os_username = ceilometer</code>	(StrOpt) User name to use for OpenStack service access.

**Table 12.6. Description of authorization token configuration options**

Configuration option = Default value	Description
[keystone_authtoken]	
admin_password = <i>None</i>	(StrOpt) Service user password.
admin_tenant_name = <i>admin</i>	(StrOpt) Service tenant name.
admin_token = <i>None</i>	(StrOpt) This option is deprecated and may be removed in a future release. Single shared secret with the Keystone configuration used for bootstrapping a Keystone installation, or otherwise bypassing the normal authentication process. This option should not be used, use `admin_user` and `admin_password` instead.
admin_user = <i>None</i>	(StrOpt) Service username.
auth_admin_prefix =	(StrOpt) Prefix to prepend at the beginning of the path. Deprecated, use <code>identity_uri</code> .
auth_host = <i>127.0.0.1</i>	(StrOpt) Host providing the admin Identity API endpoint. Deprecated, use <code>identity_uri</code> .
auth_plugin = <i>None</i>	(StrOpt) Name of the plugin to load
auth_port = <i>35357</i>	(IntOpt) Port of the admin Identity API endpoint. Deprecated, use <code>identity_uri</code> .
auth_protocol = <i>https</i>	(StrOpt) Protocol of the admin Identity API endpoint (http or https). Deprecated, use <code>identity_uri</code> .
auth_section = <i>None</i>	(StrOpt) Config Section from which to load plugin specific options
auth_uri = <i>None</i>	(StrOpt) Complete public Identity API endpoint.
auth_version = <i>None</i>	(StrOpt) API version of the admin Identity API endpoint.
cache = <i>None</i>	(StrOpt) Env key for the swift cache.
cafile = <i>None</i>	(StrOpt) A PEM encoded Certificate Authority to use when verifying HTTPS connections. Defaults to system CAs.
certfile = <i>None</i>	(StrOpt) Required if identity server requires client certificate
check_revocations_for_cached = <i>False</i>	(BoolOpt) If true, the revocation list will be checked for cached tokens. This requires that PKI tokens are configured on the identity server.
delay_auth_decision = <i>False</i>	(BoolOpt) Do not handle authorization requests within the middleware, but delegate the authorization decision to downstream WSGI components.
enforce_token_bind = <i>permissive</i>	(StrOpt) 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.
hash_algorithms = <i>md5</i>	(ListOpt) 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.
http_connect_timeout = <i>None</i>	(IntOpt) Request timeout value for communicating with Identity API server.

Configuration option = Default value	Description
<code>http_request_max_retries = 3</code>	(IntOpt) How many times are we trying to reconnect when communicating with Identity API Server.
<code>identity_uri = None</code>	(StrOpt) Complete admin Identity API endpoint. This should specify the unversioned root endpoint e.g. <code>https://localhost:35357/</code>
<code>include_service_catalog = True</code>	(BoolOpt) (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.
<code>insecure = False</code>	(BoolOpt) Verify HTTPS connections.
<code>keyfile = None</code>	(StrOpt) Required if identity server requires client certificate
<code>memcache_pool_conn_get_timeout = 10</code>	(IntOpt) (Optional) Number of seconds that an operation will wait to get a memcache client connection from the pool.
<code>memcache_pool_dead_retry = 300</code>	(IntOpt) (Optional) Number of seconds memcached server is considered dead before it is tried again.
<code>memcache_pool_maxsize = 10</code>	(IntOpt) (Optional) Maximum total number of open connections to every memcached server.
<code>memcache_pool_socket_timeout = 3</code>	(IntOpt) (Optional) Socket timeout in seconds for communicating with a memcache server.
<code>memcache_pool_unused_timeout = 60</code>	(IntOpt) (Optional) Number of seconds a connection to memcached is held unused in the pool before it is closed.
<code>memcache_secret_key = None</code>	(StrOpt) (Optional, mandatory if <code>memcache_security_strategy</code> is defined) This string is used for key derivation.
<code>memcache_security_strategy = None</code>	(StrOpt) (Optional) If defined, indicate whether token data should be authenticated or authenticated and encrypted. Acceptable values are MAC or ENCRYPT. 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.
<code>memcache_use_advanced_pool = False</code>	(BoolOpt) (Optional) Use the advanced (eventlet safe) memcache client pool. The advanced pool will only work under python 2.x.
<code>revocation_cache_time = 10</code>	(IntOpt) 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.
<code>signing_dir = None</code>	(StrOpt) Directory used to cache files related to PKI tokens.
<code>token_cache_time = 300</code>	(IntOpt) 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.

**Table 12.7. Description of collector configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>collector_workers = 1</code>	(IntOpt) Number of workers for collector service. A single collector is enabled by default.
[collector]	



Configuration option = Default value	Description
<code>memcached_servers = None</code>	(ListOpt) Optionally specify a list of memcached server(s) to use for caching. If left undefined, tokens will instead be cached in-process.
[polling]	
<code>partitioning_group_prefix = None</code>	(StrOpt) 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.

**Table 12.9. Description of concurrency configuration options**

Configuration option = Default value	Description
[oslo_concurrency]	
<code>disable_process_locking = False</code>	(BoolOpt) Enables or disables inter-process locks.
<code>lock_path = None</code>	(StrOpt) 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 <code>OSLO_LOCK_PATH</code> . If external locks are used, a lock path must be set.

**Table 12.10. Description of database configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>database_connection = None</code>	(StrOpt) DEPRECATED - Database connection string.
[database]	
<code>alarm_connection = None</code>	(StrOpt) The connection string used to connect to the alarm database. (if unset, connection is used)
<code>backend = sqlalchemy</code>	(StrOpt) The back end to use for the database.
<code>connection = None</code>	(StrOpt) The SQLAlchemy connection string to use to connect to the database.
<code>connection_debug = 0</code>	(IntOpt) Verbosity of SQL debugging information: 0=None, 100=Everything.
<code>connection_trace = False</code>	(BoolOpt) Add Python stack traces to SQL as comment strings.
<code>db2nosql_resource_id_maxlen = 512</code>	(IntOpt) The max length of resources id in DB2 nosql, the value should be larger than <code>len(hostname) * 2</code> as compute node's resource id is <code>&lt;hostname&gt;_&lt;nodename&gt;</code> .
<code>db_inc_retry_interval = True</code>	(BoolOpt) If True, increases the interval between retries of a database operation up to <code>db_max_retry_interval</code> .
<code>db_max_retries = 20</code>	(IntOpt) Maximum retries in case of connection error or deadlock error before error is raised. Set to -1 to specify an infinite retry count.
<code>db_max_retry_interval = 10</code>	(IntOpt) If <code>db_inc_retry_interval</code> is set, the maximum seconds between retries of a database operation.
<code>db_retry_interval = 1</code>	(IntOpt) Seconds between retries of a database transaction.
<code>event_connection = None</code>	(StrOpt) The connection string used to connect to the event database. (if unset, connection is used)
<code>event_time_to_live = -1</code>	(IntOpt) Number of seconds that events are kept in the database for (<= 0 means forever).
<code>idle_timeout = 3600</code>	(IntOpt) Timeout before idle SQL connections are reaped.

Configuration option = Default value	Description
<code>max_overflow = None</code>	(IntOpt) If set, use this value for <code>max_overflow</code> with SQLAlchemy.
<code>max_pool_size = None</code>	(IntOpt) Maximum number of SQL connections to keep open in a pool.
<code>max_retries = 10</code>	(IntOpt) Maximum number of database connection retries during startup. Set to -1 to specify an infinite retry count.
<code>metering_connection = None</code>	(StrOpt) The connection string used to connect to the metering database. (if unset, connection is used)
<code>metering_time_to_live = -1</code>	(IntOpt) Number of seconds that samples are kept in the database for (<= 0 means forever).
<code>min_pool_size = 1</code>	(IntOpt) Minimum number of SQL connections to keep open in a pool.
<code>mongodb_replica_set =</code>	(StrOpt) The name of the replica set which is used to connect to MongoDB database. If it is set, <code>MongoReplicaSetClient</code> will be used instead of <code>MongoClient</code> .
<code>mysql_sql_mode = TRADITIONAL</code>	(StrOpt) 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>
<code>pool_timeout = None</code>	(IntOpt) If set, use this value for <code>pool_timeout</code> with SQLAlchemy.
<code>retry_interval = 10</code>	(IntOpt) Interval between retries of opening a SQL connection.
<code>slave_connection = None</code>	(StrOpt) The SQLAlchemy connection string to use to connect to the slave database.
<code>sqlite_db = oslo.sqlite</code>	(StrOpt) The file name to use with SQLite.
<code>sqlite_synchronous = True</code>	(BoolOpt) If True, SQLite uses synchronous mode.
<code>use_db_reconnect = False</code>	(BoolOpt) Enable the experimental use of database reconnect on connection lost.

**Table 12.11. Description of logging configuration options**

Configuration option = Default value	Description
[DEFAULT]	
<code>backdoor_port = None</code>	(StrOpt) 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.
<code>nova_http_log_debug = False</code>	(BoolOpt) Allow novaclient's debug log output.

**Table 12.12. Description of HTTP dispatcher configuration options**

Configuration option = Default value	Description
[dispatcher_http]	
<code>cadf_only = False</code>	(BoolOpt) The flag that indicates if only cadf message should be posted. If false, all meters will be posted.
<code>event_target = None</code>	(StrOpt) The target for event data where the http request will be sent to. If this is not set, it will default to same as Sample target.



**Table 12.16. Description of inspector configuration options**

Configuration option = Default value	Description
[DEFAULT]	
hypervisor_inspector = <i>libvirt</i>	(StrOpt) Inspector to use for inspecting the hypervisor layer.
libvirt_type = <i>kvm</i>	(StrOpt) Libvirt domain type.
libvirt_uri =	(StrOpt) Override the default libvirt URI (which is dependent on <i>libvirt_type</i> ).

**Table 12.17. Description of IPMI configuration options**

Configuration option = Default value	Description
[ipmi]	
node_manager_init_retry = 3	(IntOpt) Number of retries upon Intel Node Manager initialization failure
polling_retry = 3	(IntOpt) Tolerance of IPMI/NM polling failures before disable this pollster. Negative indicates retrying forever.

**Table 12.18. Description of oslo\_middleware configuration options**

Configuration option = Default value	Description
[oslo_middleware]	
max_request_body_size = 114688	(IntOpt) The maximum body size for each request, in bytes.

**Table 12.19. Description of logging configuration options**

Configuration option = Default value	Description
[DEFAULT]	
debug = <i>False</i>	(BoolOpt) Print debugging output (set logging level to DEBUG instead of default WARNING level).
default_log_levels = <i>amqp=WARN, amqpplib=WARN, boto=WARN, qpid=WARN, sqlalchemy=WARN, suds=INFO, oslo.messaging=INFO, iso8601=WARN, requests.packages.urllib3.connectionpool=WARN, urllib3.connectionpool=WARN, websocket=WARN, keystonemiddleware=WARN, routes.middleware=WARN, stevedore=WARN</i>	(ListOpt) List of logger=LEVEL pairs.
fatal_deprecations = <i>False</i>	(BoolOpt) Enables or disables fatal status of deprecations.
instance_format = "[instance: %(uid)s] "	(StrOpt) The format for an instance that is passed with the log message.
instance_uuid_format = "[instance: %(uid)s] "	(StrOpt) The format for an instance UUID that is passed with the log message.
log_config_append = <i>None</i>	(StrOpt) 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.
log_date_format = <i>%Y-%m-%d %H:%M:%S</i>	(StrOpt) Format string for %(asctime)s in log records. Default: %(default)s .
log_dir = <i>None</i>	(StrOpt) (Optional) The base directory used for relative – log-file paths.
log_file = <i>None</i>	(StrOpt) (Optional) Name of log file to output to. If no default is set, logging will go to stdout.















```
# (unknown type)
#pollster_list = []

# Exchange name for Nova notifications. (string value)
#nova_control_exchange = nova

# List of metadata prefixes reserved for metering use. (list value)
#reserved_metadata_namespace = metering.

# Limit on length of reserved metadata values. (integer value)
#reserved_metadata_length = 256

# List of metadata keys reserved for metering use. And these keys are
# additional to the ones included in the namespace. (list value)
#reserved_metadata_keys =

# Inspector to use for inspecting the hypervisor layer. (string value)
#hypervisor_inspector = libvirt

# Libvirt domain type. (string value)
# Allowed values: kvm, lxc, qemu, uml, xen
#libvirt_type = kvm

# Override the default libvirt URI (which is dependent on
# libvirt_type). (string value)
#libvirt_uri =

# Exchange name for Data Processing notifications. (string value)
#sahara_control_exchange = sahara

# Dispatcher to process data. (multi valued)
# Deprecated group/name - [collector]/dispatcher
#dispatcher = database

# Exchange name for Keystone notifications. (string value)
#keystone_control_exchange = keystone

# Number of items to request in each paginated Glance API request
# (parameter used by glanceclient). If this is less than or equal to
# 0, page size is not specified (default value in glanceclient is
# used). (integer value)
#glance_page_size = 0

# Exchange name for Glance notifications. (string value)
#glance_control_exchange = glance

# Exchange name for Ironic notifications. (string value)
#ironic_exchange = ironic

# Exchanges name to listen for notifications. (multi valued)
#http_control_exchanges = nova
#http_control_exchanges = glance
#http_control_exchanges = neutron
#http_control_exchanges = cinder

# Exchange name for Neutron notifications. (string value)
# Deprecated group/name - [DEFAULT]/quantum_control_exchange
#neutron_control_exchange = neutron

# Allow novaclient's debug log output. (boolean value)
```

```
#nova_http_log_debug = false

# Swift reseller prefix. Must be on par with reseller_prefix in proxy-
# server.conf. (string value)
#reseller_prefix = AUTH_

# 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. (string value)
#backdoor_port = <None>

# Print debugging output (set logging level to DEBUG instead of
# default WARNING level). (boolean value)
#debug = false

# Print more verbose output (set logging level to INFO instead of
# default WARNING level). (boolean value)
#verbose = false

# Log output to standard error. (boolean value)
#use_stderr = true

# 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.
# (string value)
# Deprecated group/name - [DEFAULT]/log_config
#log_config_append = <None>

# DEPRECATED. A logging.Formatter log message format string which may
# use any of the available logging.LogRecord attributes. This option
# is deprecated. Please use logging_context_format_string and
# logging_default_format_string instead. (string value)
#log_format = <None>

# Format string for %(asctime)s in log records. Default: %(default)s
# . (string value)
#log_date_format = %Y-%m-%d %H:%M:%S

# (Optional) Name of log file to output to. If no default is set,
# logging will go to stdout. (string value)
# Deprecated group/name - [DEFAULT]/logfile
#log_file = <None>

# (Optional) The base directory used for relative --log-file paths.
# (string value)
# Deprecated group/name - [DEFAULT]/logdir
#log_dir = <None>

# Use syslog for logging. Existing syslog format is DEPRECATED during
# I, and will change in J to honor RFC5424. (boolean value)
#use_syslog = false

# (Optional) Enables or disables syslog rfc5424 format for logging. If
# enabled, prefixes the MSG part of the syslog message with APP-NAME
# (RFC5424). The format without the APP-NAME is deprecated in I, and
```







```
# impl_zmq. (integer value)
#rpc_cast_timeout = 30

# Heartbeat frequency. (integer value)
#matchmaker_heartbeat_freq = 300

# Heartbeat time-to-live. (integer value)
#matchmaker_heartbeat_ttl = 600

# Size of RPC thread pool. (integer value)
#rpc_thread_pool_size = 64

# Driver or drivers to handle sending notifications. (multi valued)
#notification_driver =

# AMQP topic used for OpenStack notifications. (list value)
# Deprecated group/name - [rpc_notifier2]/topics
#notification_topics = notifications

# Seconds to wait for a response from a call. (integer value)
#rpc_response_timeout = 60

# A URL representing the messaging driver to use and its full
# configuration. If not set, we fall back to the rpc_backend option
# and driver specific configuration. (string value)
#transport_url = <None>

# The messaging driver to use, defaults to rabbit. Other drivers
# include qpid and zmq. (string value)
#rpc_backend = rabbit

# The default exchange under which topics are scoped. May be
# overridden by an exchange name specified in the transport_url
# option. (string value)
#control_exchange = openstack

[alarm]

#
# From ceilometer
#

# SSL Client certificate for REST notifier. (string value)
#rest_notifier_certificate_file =

# SSL Client private key for REST notifier. (string value)
#rest_notifier_certificate_key =

# Whether to verify the SSL Server certificate when calling alarm
# action. (boolean value)
#rest_notifier_ssl_verify = true

# Number of retries for REST notifier (integer value)
#rest_notifier_max_retries = 0

# Period of evaluation cycle, should be >= than configured pipeline
# interval for collection of underlying meters. (integer value)
# Deprecated group/name - [alarm]/threshold_evaluation_interval
#evaluation_interval = 60
```







```
# 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= (string value)
#mysql_sql_mode = TRADITIONAL

# Timeout before idle SQL connections are reaped. (integer value)
# Deprecated group/name - [DEFAULT]/sql_idle_timeout
# Deprecated group/name - [DATABASE]/sql_idle_timeout
# Deprecated group/name - [sql]/idle_timeout
#idle_timeout = 3600

# Minimum number of SQL connections to keep open in a pool. (integer
# value)
# Deprecated group/name - [DEFAULT]/sql_min_pool_size
# Deprecated group/name - [DATABASE]/sql_min_pool_size
#min_pool_size = 1

# Maximum number of SQL connections to keep open in a pool. (integer
# value)
# Deprecated group/name - [DEFAULT]/sql_max_pool_size
# Deprecated group/name - [DATABASE]/sql_max_pool_size
#max_pool_size = <None>

# Maximum number of database connection retries during startup. Set to
# -1 to specify an infinite retry count. (integer value)
# Deprecated group/name - [DEFAULT]/sql_max_retries
# Deprecated group/name - [DATABASE]/sql_max_retries
#max_retries = 10

# Interval between retries of opening a SQL connection. (integer
# value)
# Deprecated group/name - [DEFAULT]/sql_retry_interval
# Deprecated group/name - [DATABASE]/reconnect_interval
#retry_interval = 10

# If set, use this value for max_overflow with SQLAlchemy. (integer
# value)
# Deprecated group/name - [DEFAULT]/sql_max_overflow
# Deprecated group/name - [DATABASE]/sqlalchemy_max_overflow
#max_overflow = <None>

# Verbosity of SQL debugging information: 0=None, 100=Everything.
# (integer value)
# Deprecated group/name - [DEFAULT]/sql_connection_debug
#connection_debug = 0

# Add Python stack traces to SQL as comment strings. (boolean value)
# Deprecated group/name - [DEFAULT]/sql_connection_trace
#connection_trace = false

# If set, use this value for pool_timeout with SQLAlchemy. (integer
# value)
# Deprecated group/name - [DATABASE]/sqlalchemy_pool_timeout
#pool_timeout = <None>

# Enable the experimental use of database reconnect on connection
# lost. (boolean value)
#use_db_reconnect = false
```

```
# Seconds between retries of a database transaction. (integer value)
#db_retry_interval = 1

# If True, increases the interval between retries of a database
# operation up to db_max_retry_interval. (boolean value)
#db_inc_retry_interval = true

# If db_inc_retry_interval is set, the maximum seconds between retries
# of a database operation. (integer value)
#db_max_retry_interval = 10

# Maximum retries in case of connection error or deadlock error before
# error is raised. Set to -1 to specify an infinite retry count.
# (integer value)
#db_max_retries = 20

[dispatcher_file]

#
# From ceilometer
#

# Name and the location of the file to record meters. (string value)
#file_path = <None>

# The max size of the file. (integer value)
#max_bytes = 0

# The max number of the files to keep. (integer value)
#backup_count = 0

[event]

#
# From ceilometer
#

# Configuration file for event definitions. (string value)
#definitions_cfg_file = event_definitions.yaml

# Drop notifications if no event definition matches. (Otherwise, we
# convert them with just the default traits) (boolean value)
#drop_unmatched_notifications = false

# Store the raw notification for select priority levels (info and/or
# error). By default, raw details are not captured. (multi valued)
#store_raw =

[hardware]

#
# From ceilometer
#

# URL scheme to use for hardware nodes. (string value)
#url_scheme = snmp://
```



```
# SNMPd user name of all nodes running in the cloud. (string value)
#readonly_user_name = ro_snmp_user

# SNMPd password of all the nodes running in the cloud. (string value)
#readonly_user_password = password

[ipmi]

#
# From ceilometer
#

# Number of retries upon Intel Node Manager initialization failure
# (integer value)
#node_manager_init_retry = 3

# Tolerance of IPMI/NM polling failures before disable this pollster.
# Negative indicates retrying forever. (integer value)
#polling_retry = 3

[keystone_authtoken]

#
# From keystonemiddleware.auth_token
#

# Complete public Identity API endpoint. (string value)
#auth_uri = <None>

# API version of the admin Identity API endpoint. (string value)
#auth_version = <None>

# Do not handle authorization requests within the middleware, but
# delegate the authorization decision to downstream WSGI components.
# (boolean value)
#delay_auth_decision = false

# Request timeout value for communicating with Identity API server.
# (integer value)
#http_connect_timeout = <None>

# How many times are we trying to reconnect when communicating with
# Identity API Server. (integer value)
#http_request_max_retries = 3

# Env key for the swift cache. (string value)
#cache = <None>

# Required if identity server requires client certificate (string
# value)
#certfile = <None>

# Required if identity server requires client certificate (string
# value)
#keyfile = <None>

# A PEM encoded Certificate Authority to use when verifying HTTPs
# connections. Defaults to system CAs. (string value)
```

```
#cafile = <None>

# Verify HTTPS connections. (boolean value)
#insecure = false

# Directory used to cache files related to PKI tokens. (string value)
#signing_dir = <None>

# Optionally specify a list of memcached server(s) to use for caching.
# If left undefined, tokens will instead be cached in-process. (list
# value)
# Deprecated group/name - [DEFAULT]/memcache_servers
#memcached_servers = <None>

# 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. (integer
# value)
#token_cache_time = 300

# 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. (integer value)
#revocation_cache_time = 10

# (Optional) If defined, indicate whether token data should be
# authenticated or authenticated and encrypted. Acceptable values are
# MAC or ENCRYPT. 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. (string value)
#memcache_security_strategy = <None>

# (Optional, mandatory if memcache_security_strategy is defined) This
# string is used for key derivation. (string value)
#memcache_secret_key = <None>

# (Optional) Number of seconds memcached server is considered dead
# before it is tried again. (integer value)
#memcache_pool_dead_retry = 300

# (Optional) Maximum total number of open connections to every
# memcached server. (integer value)
#memcache_pool_maxsize = 10

# (Optional) Socket timeout in seconds for communicating with a
# memcache server. (integer value)
#memcache_pool_socket_timeout = 3

# (Optional) Number of seconds a connection to memcached is held
# unused in the pool before it is closed. (integer value)
#memcache_pool_unused_timeout = 60

# (Optional) Number of seconds that an operation will wait to get a
# memcache client connection from the pool. (integer value)
#memcache_pool_conn_get_timeout = 10

# (Optional) Use the advanced (eventlet safe) memcache client pool.
# The advanced pool will only work under python 2.x. (boolean value)
```

```
#memcache_use_advanced_pool = false

# (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. (boolean
# value)
#include_service_catalog = true

# 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. (string value)
#enforce_token_bind = permissive

# If true, the revocation list will be checked for cached tokens. This
# requires that PKI tokens are configured on the identity server.
# (boolean value)
#check_revocations_for_cached = false

# 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. (list
# value)
#hash_algorithms = md5

# Prefix to prepend at the beginning of the path. Deprecated, use
# identity_uri. (string value)
#auth_admin_prefix =

# Host providing the admin Identity API endpoint. Deprecated, use
# identity_uri. (string value)
#auth_host = 127.0.0.1

# Port of the admin Identity API endpoint. Deprecated, use
# identity_uri. (integer value)
#auth_port = 35357

# Protocol of the admin Identity API endpoint (http or https).
# Deprecated, use identity_uri. (string value)
#auth_protocol = https

# Complete admin Identity API endpoint. This should specify the
# unversioned root endpoint e.g. https://localhost:35357/ (string
# value)
#identity_uri = <None>

# This option is deprecated and may be removed in a future release.
# Single shared secret with the Keystone configuration used for
# bootstrapping a Keystone installation, or otherwise bypassing the
# normal authentication process. This option should not be used, use
# `admin_user` and `admin_password` instead. (string value)
#admin_token = <None>
```

```
# Service username. (string value)
#admin_user = <None>

# Service user password. (string value)
#admin_password = <None>

# Service tenant name. (string value)
#admin_tenant_name = admin

[matchmaker_redis]

#
# From oslo.messaging
#

# Host to locate redis. (string value)
#host = 127.0.0.1

# Use this port to connect to redis host. (integer value)
#port = 6379

# Password for Redis server (optional). (string value)
#password = <None>

[matchmaker_ring]

#
# From oslo.messaging
#

# Matchmaker ring file (JSON). (string value)
# Deprecated group/name - [DEFAULT]/matchmaker_ringfile
#ringfile = /etc/oslo/matchmaker_ring.json

[notification]

#
# From ceilometer
#

# Acknowledge message when event persistence fails. (boolean value)
# Deprecated group/name - [collector]/ack_on_event_error
#ack_on_event_error = true

# Save event details. (boolean value)
# Deprecated group/name - [collector]/store_events
#store_events = false

# WARNING: Ceilometer historically offered the ability to store events
# as meters. This usage is NOT advised as it can flood the metering
# database and cause performance degradation. This option disables the
# collection of non-metric meters and will be the default behavior in
# Liberty. (boolean value)
#disable_non_metric_meters = false

# Enable workload partitioning, allowing multiple notification agents
```



```
#ssl_ca_file =  
  
# Identifying certificate PEM file to present to clients (string  
# value)  
# Deprecated group/name - [amqp1]/ssl_cert_file  
#ssl_cert_file =  
  
# Private key PEM file used to sign cert_file certificate (string  
# value)  
# Deprecated group/name - [amqp1]/ssl_key_file  
#ssl_key_file =  
  
# Password for decrypting ssl_key_file (if encrypted) (string value)  
# Deprecated group/name - [amqp1]/ssl_key_password  
#ssl_key_password = <None>  
  
# Accept clients using either SSL or plain TCP (boolean value)  
# Deprecated group/name - [amqp1]/allow_insecure_clients  
#allow_insecure_clients = false  
  
[oslo_messaging_qpid]  
  
#  
# From oslo.messaging  
#  
  
# Use durable queues in AMQP. (boolean value)  
# Deprecated group/name - [DEFAULT]/rabbit_durable_queues  
#amqp_durable_queues = false  
  
# Auto-delete queues in AMQP. (boolean value)  
# Deprecated group/name - [DEFAULT]/amqp_auto_delete  
#amqp_auto_delete = false  
  
# Size of RPC connection pool. (integer value)  
# Deprecated group/name - [DEFAULT]/rpc_conn_pool_size  
#rpc_conn_pool_size = 30  
  
# Qpid broker hostname. (string value)  
# Deprecated group/name - [DEFAULT]/qpid_hostname  
#qpid_hostname = localhost  
  
# Qpid broker port. (integer value)  
# Deprecated group/name - [DEFAULT]/qpid_port  
#qpid_port = 5672  
  
# Qpid HA cluster host:port pairs. (list value)  
# Deprecated group/name - [DEFAULT]/qpid_hosts  
#qpid_hosts = $qpid_hostname:$qpid_port  
  
# Username for Qpid connection. (string value)  
# Deprecated group/name - [DEFAULT]/qpid_username  
#qpid_username =  
  
# Password for Qpid connection. (string value)  
# Deprecated group/name - [DEFAULT]/qpid_password  
#qpid_password =  
  
# Space separated list of SASL mechanisms to use for auth. (string
```

```
# value)
# Deprecated group/name - [DEFAULT]/qp_id_sasl_mechanisms
#qp_id_sasl_mechanisms =

# Seconds between connection keepalive heartbeats. (integer value)
# Deprecated group/name - [DEFAULT]/qp_id_heartbeat
#qp_id_heartbeat = 60

# Transport to use, either 'tcp' or 'ssl'. (string value)
# Deprecated group/name - [DEFAULT]/qp_id_protocol
#qp_id_protocol = tcp

# Whether to disable the Nagle algorithm. (boolean value)
# Deprecated group/name - [DEFAULT]/qp_id_tcp_nodelay
#qp_id_tcp_nodelay = true

# The number of prefetched messages held by receiver. (integer value)
# Deprecated group/name - [DEFAULT]/qp_id_receiver_capacity
#qp_id_receiver_capacity = 1

# The qp_id topology version to use. Version 1 is what was originally
# used by impl_qp_id. Version 2 includes some backwards-incompatible
# changes that allow broker federation to work. Users should update
# to version 2 when they are able to take everything down, as it
# requires a clean break. (integer value)
# Deprecated group/name - [DEFAULT]/qp_id_topology_version
#qp_id_topology_version = 1

[oslo_messaging_rabbit]

#
# From oslo.messaging
#

# Use durable queues in AMQP. (boolean value)
# Deprecated group/name - [DEFAULT]/rabbit_durable_queues
#amqp_durable_queues = false

# Auto-delete queues in AMQP. (boolean value)
# Deprecated group/name - [DEFAULT]/amqp_auto_delete
#amqp_auto_delete = false

# Size of RPC connection pool. (integer value)
# Deprecated group/name - [DEFAULT]/rpc_conn_pool_size
#rpc_conn_pool_size = 30

# 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. (string value)
# Deprecated group/name - [DEFAULT]/kombu_ssl_version
#kombu_ssl_version =

# SSL key file (valid only if SSL enabled). (string value)
# Deprecated group/name - [DEFAULT]/kombu_ssl_keyfile
#kombu_ssl_keyfile =

# SSL cert file (valid only if SSL enabled). (string value)
# Deprecated group/name - [DEFAULT]/kombu_ssl_certfile
#kombu_ssl_certfile =
```

```
# SSL certification authority file (valid only if SSL enabled).
# (string value)
# Deprecated group/name - [DEFAULT]/kombu_ssl_ca_certs
#kombu_ssl_ca_certs =

# How long to wait before reconnecting in response to an AMQP consumer
# cancel notification. (floating point value)
# Deprecated group/name - [DEFAULT]/kombu_reconnect_delay
#kombu_reconnect_delay = 1.0

# The RabbitMQ broker address where a single node is used. (string
# value)
# Deprecated group/name - [DEFAULT]/rabbit_host
#rabbit_host = localhost

# The RabbitMQ broker port where a single node is used. (integer
# value)
# Deprecated group/name - [DEFAULT]/rabbit_port
#rabbit_port = 5672

# RabbitMQ HA cluster host:port pairs. (list value)
# Deprecated group/name - [DEFAULT]/rabbit_hosts
#rabbit_hosts = $rabbit_host:$rabbit_port

# Connect over SSL for RabbitMQ. (boolean value)
# Deprecated group/name - [DEFAULT]/rabbit_use_ssl
#rabbit_use_ssl = false

# The RabbitMQ userid. (string value)
# Deprecated group/name - [DEFAULT]/rabbit_userid
#rabbit_userid = guest

# The RabbitMQ password. (string value)
# Deprecated group/name - [DEFAULT]/rabbit_password
#rabbit_password = guest

# The RabbitMQ login method. (string value)
# Deprecated group/name - [DEFAULT]/rabbit_login_method
#rabbit_login_method = AMQPPLAIN

# The RabbitMQ virtual host. (string value)
# Deprecated group/name - [DEFAULT]/rabbit_virtual_host
#rabbit_virtual_host = /

# How frequently to retry connecting with RabbitMQ. (integer value)
#rabbit_retry_interval = 1

# How long to backoff for between retries when connecting to RabbitMQ.
# (integer value)
# Deprecated group/name - [DEFAULT]/rabbit_retry_backoff
#rabbit_retry_backoff = 2

# Maximum number of RabbitMQ connection retries. Default is 0
# (infinite retry count). (integer value)
# Deprecated group/name - [DEFAULT]/rabbit_max_retries
#rabbit_max_retries = 0

# Use HA queues in RabbitMQ (x-ha-policy: all). If you change this
# option, you must wipe the RabbitMQ database. (boolean value)
```



```
# Deprecated group/name - [DEFAULT]/rabbit_ha_queues
#rabbit_ha_queues = false

# Number of seconds after which the Rabbit broker is considered down
# if heartbeat's keep-alive fails (0 disables the heartbeat, >0
# enables it. Enabling heartbeats requires kombu>=3.0.7 and
# amqp>=1.4.0). EXPERIMENTAL (integer value)
#heartbeat_timeout_threshold = 0

# How often times during the heartbeat_timeout_threshold we check the
# heartbeat. (integer value)
#heartbeat_rate = 2

# Deprecated, use rpc_backend=kombu+memory or rpc_backend=fake
# (boolean value)
# Deprecated group/name - [DEFAULT]/fake_rabbit
#fake_rabbit = false

[oslo_policy]

#
# From oslo.policy
#

# The JSON file that defines policies. (string value)
# Deprecated group/name - [DEFAULT]/policy_file
#policy_file = policy.json

# Default rule. Enforced when a requested rule is not found. (string
# value)
# Deprecated group/name - [DEFAULT]/policy_default_rule
#policy_default_rule = default

# 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. (multi valued)
# Deprecated group/name - [DEFAULT]/policy_dirs
#policy_dirs = policy.d

[polling]

#
# From ceilometer
#

# 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 partitioning_group_prefix a
# disjoint subset of pollsters should be loaded. (string value)
# Deprecated group/name - [central]/partitioning_group_prefix
#partitioning_group_prefix = <None>

[publisher]

#
```

```
# From ceilometer
#

# Secret value for signing messages. Set value empty if signing is not
# required to avoid computational overhead. (string value)
# Deprecated group/name - [DEFAULT]/metering_secret
# Deprecated group/name - [publisher_rpc]/metering_secret
# Deprecated group/name - [publisher]/metering_secret
#telemetry_secret = change this for valid signing

[publisher_notifier]

#
# From ceilometer
#

# The topic that ceilometer uses for metering notifications. (string
# value)
#metering_topic = metering

# The topic that ceilometer uses for event notifications. (string
# value)
#event_topic = event

# The driver that ceilometer uses for metering notifications. (string
# value)
# Deprecated group/name - [DEFAULT]/metering_driver
#telemetry_driver = messagingv2

[publisher_rpc]

#
# From ceilometer
#

# The topic that ceilometer uses for metering messages. (string value)
# Deprecated group/name - [DEFAULT]/metering_topic
#metering_topic = metering

[rgw_admin_credentials]

#
# From ceilometer
#

# Access key for Radosgw Admin. (string value)
#access_key = <None>

# Secret key for Radosgw Admin. (string value)
#secret_key = <None>

[service_credentials]

#
# From ceilometer
#
```

```

# User name to use for OpenStack service access. (string value)
# Deprecated group/name - [DEFAULT]/os_username
#os_username = ceilometer

# Password to use for OpenStack service access. (string value)
# Deprecated group/name - [DEFAULT]/os_password
#os_password = admin

# Tenant ID to use for OpenStack service access. (string value)
# Deprecated group/name - [DEFAULT]/os_tenant_id
#os_tenant_id =

# Tenant name to use for OpenStack service access. (string value)
# Deprecated group/name - [DEFAULT]/os_tenant_name
#os_tenant_name = admin

# Certificate chain for SSL validation. (string value)
#os_cacert = <None>

# Auth URL to use for OpenStack service access. (string value)
# Deprecated group/name - [DEFAULT]/os_auth_url
#os_auth_url = http://localhost:5000/v2.0

# Region name to use for OpenStack service endpoints. (string value)
# Deprecated group/name - [DEFAULT]/os_region_name
#os_region_name = <None>

# Type of endpoint in Identity service catalog to use for
# communication with OpenStack services. (string value)
#os_endpoint_type = publicURL

# Disables X.509 certificate validation when an SSL connection to
# Identity Service is established. (boolean value)
#insecure = false

[service_types]

#
# From ceilometer
#

# Kwapi service type. (string value)
#kwapi = energy

# Glance service type. (string value)
#glance = image

# Neutron service type. (string value)
#neutron = network

# Nova service type. (string value)
#nova = compute

# Radosgw service type. (string value)
#radosgw = object-store

# Swift service type. (string value)
#swift = object-store

```

```
[vmware]

#
# From ceilometer
#

# IP address of the VMware Vsphere host. (string value)
#host_ip =

# Port of the VMware Vsphere host. (integer value)
#host_port = 443

# Username of VMware Vsphere. (string value)
#host_username =

# Password of VMware Vsphere. (string value)
#host_password =

# Number of times a VMware Vsphere API may be retried. (integer value)
#api_retry_count = 10

# Sleep time in seconds for polling an ongoing async task. (floating
# point value)
#task_poll_interval = 0.5

# Optional vim service WSDL location e.g
# http://<server>/vimService.wsdl. Optional over-ride to default
# location for bug work-arounds. (string value)
#wSDL_location = <None>

[xenapi]

#
# From ceilometer
#

# URL for connection to XenServer/Xen Cloud Platform. (string value)
#connection_url = <None>

# Username for connection to XenServer/Xen Cloud Platform. (string
# value)
#connection_username = root

# Password for connection to XenServer/Xen Cloud Platform. (string
# value)
#connection_password = <None>

# Timeout in seconds for XenAPI login. (integer value)
#login_timeout = 10
```

## **event\_definitions.yaml**

The `event_definitions.yaml` file defines how events received from other OpenStack components should be translated to Telemetry events.

This file provides a standard set of events and corresponding traits that may be of interest. This file can be modified to add and drop traits that operators may find useful.

```

---
- event_type: compute.instance.*
  traits: &instance_traits
    tenant_id:
      fields: payload.tenant_id
    user_id:
      fields: payload.user_id
    instance_id:
      fields: payload.instance_id
    host:
      fields: publisher_id
    plugin:
      name: split
      parameters:
        segment: 1
        max_split: 1
    service:
      fields: publisher_id
      plugin: split
    memory_mb:
      type: int
      fields: payload.memory_mb
    disk_gb:
      type: int
      fields: payload.disk_gb
    root_gb:
      type: int
      fields: payload.root_gb
    ephemeral_gb:
      type: int
      fields: payload.ephemeral_gb
    vcpus:
      type: int
      fields: payload.vcpus
    instance_type_id:
      type: int
      fields: payload.instance_type_id
    instance_type:
      fields: payload.instance_type
    state:
      fields: payload.state
    os_architecture:
      fields: payload.image_meta.'org.openstack_1__architecture'
    os_version:
      fields: payload.image_meta.'org.openstack_1__os_version'
    os_distro:
      fields: payload.image_meta.'org.openstack_1__os_distro'
    launched_at:
      type: datetime
      fields: payload.launched_at
    deleted_at:
      type: datetime
      fields: payload.deleted_at
- event_type: compute.instance.exists
  traits:
    <<: *instance_traits
    audit_period_beginning:

```

```

    type: datetime
    fields: payload.audit_period_beginning
  audit_period_ending:
    type: datetime
    fields: payload.audit_period_ending
- event_type: ['volume.exists', 'volume.create.*', 'volume.delete.*', 'volume.
resize.*', 'volume.attach.*', 'volume.detach.*', 'volume.update.*', 'snapshot.
exists', 'snapshot.create.*', 'snapshot.delete.*', 'snapshot.update.*']
  traits: &cinder_traits
    user_id:
      fields: payload.user_id
    project_id:
      fields: payload.tenant_id
    availability_zone:
      fields: payload.availability_zone
    display_name:
      fields: payload.display_name
    replication_status:
      fields: payload.replication_status
    status:
      fields: payload.status
    created_at:
      fields: payload.created_at
- event_type: ['volume.exists', 'volume.create.*', 'volume.delete.*', 'volume.
resize.*', 'volume.attach.*', 'volume.detach.*', 'volume.update.*']
  traits:
    <<: *cinder_traits
    resource_id:
      fields: payload.volume_id
    host:
      fields: payload.host
    size:
      fields: payload.size
    type:
      fields: payload.volume_type
    replication_status:
      fields: payload.replication_status
- event_type: ['snapshot.exists', 'snapshot.create.*', 'snapshot.delete.*',
'snapshot.update.*']
  traits:
    <<: *cinder_traits
    resource_id:
      fields: payload.snapshot_id
    volume_id:
      fields: payload.volume_id
- event_type: ['image.update', 'image.upload', 'image.delete']
  traits: &glance_crud
    project_id:
      fields: payload.owner
    resource_id:
      fields: payload.id
    name:
      fields: payload.name
    status:
      fields: payload.status
    created_at:
      fields: payload.created_at
    user_id:
      fields: payload.owner
    deleted_at:

```



```

typeURI:
  fields: payload.typeURI
id:
  fields: payload.id
action:
  fields: payload.action
eventType:
  fields: payload.eventType
eventTime:
  fields: payload.eventTime
outcome:
  fields: payload.outcome
initiator_typeURI:
  fields: payload.initiator.typeURI
initiator_id:
  fields: payload.initiator.id
initiator_name:
  fields: payload.initiator.name
initiator_host_agent:
  fields: payload.initiator.host.agent
initiator_host_addr:
  fields: payload.initiator.host.address
target_typeURI:
  fields: payload.target.typeURI
target_id:
  fields: payload.target.id
observer_typeURI:
  fields: payload.observer.typeURI
observer_id:
  fields: payload.observer.id
- event_type: objectstore.http.request
traits: &objectstore_request
typeURI:
  fields: payload.typeURI
id:
  fields: payload.id
action:
  fields: payload.action
eventType:
  fields: payload.eventType
eventTime:
  fields: payload.eventTime
outcome:
  fields: payload.outcome
initiator_typeURI:
  fields: payload.initiator.typeURI
initiator_id:
  fields: payload.initiator.id
initiator_project_id:
  fields: payload.initiator.project_id
target_typeURI:
  fields: payload.target.typeURI
target_id:
  fields: payload.target.id
target_action:
  fields: payload.target.action
target_metadata_path:
  fields: payload.target.metadata.path
target_metadata_version:
  fields: payload.target.metadata.version

```



```
target_metadata_container:  
  fields: payload.target.metadata.container  
target_metadata_object:  
  fields: payload.target.metadata.object  
observer_id:  
  fields: payload.observer.id  
- event_type: magnetodb.table.*  
traits: &kv_store  
  resource_id:  
    fields: payload.table_uuid  
  user_id:  
    fields: _context_user_id  
  project_id:  
    fields: _context_tenant  
- event_type: ['network.*', 'subnet.*', 'port.*', 'router.*', 'floatingip.  
*', 'pool.*', 'vip.*', 'member.*', 'health_monitor.*', 'firewall.*',  
'firewall_policy.*', 'firewall_rule.*', 'vpnservice.*', 'ipsecpolicy.*',  
'ikepolicy.*', 'ipsec_site_connection.*']  
traits: &network_traits  
  user_id:  
    fields: _context_user_id  
  project_id:  
    fields: _context_tenant_id  
- event_type: network.*  
traits:  
  <<: *network_traits  
  resource_id:  
    fields: ['payload.network.id', 'payload.id']  
- event_type: subnet.*  
traits:  
  <<: *network_traits  
  resource_id:  
    fields: ['payload.subnet.id', 'payload.id']  
- event_type: port.*  
traits:  
  <<: *network_traits  
  resource_id:  
    fields: ['payload.port.id', 'payload.id']  
- event_type: router.*  
traits:  
  <<: *network_traits  
  resource_id:  
    fields: ['payload.router.id', 'payload.id']  
- event_type: floatingip.*  
traits:  
  <<: *network_traits  
  resource_id:  
    fields: ['payload.floatingip.id', 'payload.id']  
- event_type: pool.*  
traits:  
  <<: *network_traits  
  resource_id:  
    fields: ['payload.pool.id', 'payload.id']  
- event_type: vip.*  
traits:  
  <<: *network_traits  
  resource_id:  
    fields: ['payload.vip.id', 'payload.id']  
- event_type: member.*  
traits:
```







## event\_pipeline.yaml

Event pipelines describe a coupling between notification event\_types and the corresponding sinks for publication of the event data. They are defined in the event\_pipeline.yaml file.

This file can be modified to adjust which notifications to capture and the and where to publish the events.

```

---
sources:
  - name: event_source
    events:
      - "*"
    sinks:
      - event_sink
sinks:
  - name: event_sink
    transformers:
    triggers:
    publishers:
      - direct://
  
```

## policy.json

The policy.json file defines additional access controls that apply to the Telemetry service.

```

{
  "context_is_admin": "role:admin",
  "context_is_project": "project_id:!(target.project_id)s",
  "context_is_owner": "user_id:!(target.user_id)s",
  "segregation": "rule:context_is_admin",
  "default": ""
}
  
```

# New, updated and deprecated options in Kilo for Telemetry

**Table 12.35. New options**

Option = default value	(Type) Help string
[DEFAULT] api_workers = 1	(IntOpt) Number of workers for Ceilometer API server.
[DEFAULT] event_pipeline_cfg_file = event_pipeline.yaml	(StrOpt) Configuration file for event pipeline definition.
[DEFAULT] magnetodb_control_exchange = magnetodb	(StrOpt) Exchange name for Magnetodb notifications.
[DEFAULT] polling_namespaces = ['compute', 'central']	(MultiChoicesOpt) Polling namespace(s) to be used while resource polling
[DEFAULT] pollster_list = []	(MultiChoicesOpt) List of pollsters (or wildcard templates) to be used while polling
[DEFAULT] reserved_metadata_keys =	(ListOpt) List of metadata keys reserved for metering use. And these keys are additional to the ones included in the namespace.

Option = default value	(Type) Help string
[DEFAULT] shuffle_time_before_polling_task = 0	(IntOpt) To reduce large requests at same time to Nova or other components from different compute agents, shuffle start time of polling task.
[DEFAULT] sql_expire_samples_only = False	(BoolOpt) Indicates if expirer expires only samples. If set true, expired samples will be deleted, but residual resource and meter definition data will remain.
[DEFAULT] swift_control_exchange = swift	(StrOpt) Exchange name for Swift notifications.
[DEFAULT] zaqar_control_exchange = zaqar	(StrOpt) Exchange name for Messaging service notifications.
[alarms] gnocchi_url = http://localhost:8041	(StrOpt) URL to Gnocchi.
[collector] requeue_event_on_dispatcher_error = False	(BoolOpt) Requeue the event on the collector event queue when the collector fails to dispatch it.
[coordination] check_watchers = 10.0	(FloatOpt) Number of seconds between checks to see if group membership has changed
[database] db2nosql_resource_id_maxlen = 512	(IntOpt) The max length of resources id in DB2 nosql, the value should be larger than len(hostname) * 2 as compute node's resource id is <hostname>_<nodename>.
[database] event_connection = None	(StrOpt) The connection string used to connect to the event database. (if unset, connection is used)
[database] event_time_to_live = -1	(IntOpt) Number of seconds that events are kept in the database for (<= 0 means forever).
[database] metering_time_to_live = -1	(IntOpt) Number of seconds that samples are kept in the database for (<= 0 means forever).
[database] mongodbd_replica_set =	(StrOpt) The name of the replica set which is used to connect to MongoDB database. If it is set, MongoReplicaSetClient will be used instead of MongoClient.
[dispatcher_http] cadf_only = False	(BoolOpt) The flag that indicates if only cadf message should be posted. If false, all meters will be posted.
[dispatcher_http] event_target = None	(StrOpt) The target for event data where the http request will be sent to. If this is not set, it will default to same as Sample target.
[dispatcher_http] target =	(StrOpt) The target where the http request will be sent. If this is not set, no data will be posted. For example: target = http://hostname:1234/path
[dispatcher_http] timeout = 5	(IntOpt) The max time in seconds to wait for a request to timeout.
[event] store_raw = []	(MultiStrOpt) Store the raw notification for select priority levels (info and/or error). By default, raw details are not captured.
[ipmi] polling_retry = 3	(IntOpt) Tolerance of IPMI/NM polling failures before disable this pollster. Negative indicates retrying forever.
[notification] disable_non_metric_meters = False	(BoolOpt) WARNING: Ceilometer historically offered the ability to store events as meters. This usage is NOT advised as it can flood the metering database and cause performance degradation. This option disables the collection of non-metric meters and will be the default behavior in Liberty.
[notification] workload_partitioning = False	(BoolOpt) Enable workload partitioning, allowing multiple notification agents to be run simultaneously.
[oslo_concurrency] disable_process_locking = False	(BoolOpt) Enables or disables inter-process locks.
[oslo_concurrency] lock_path = None	(StrOpt) 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.

Option = default value	(Type) Help string
[oslo_messaging_amqp] allow_insecure_clients = False	(BoolOpt) Accept clients using either SSL or plain TCP
[oslo_messaging_amqp] broadcast_prefix = broadcast	(StrOpt) address prefix used when broadcasting to all servers
[oslo_messaging_amqp] container_name = None	(StrOpt) Name for the AMQP container
[oslo_messaging_amqp] group_request_prefix = unicast	(StrOpt) address prefix when sending to any server in group
[oslo_messaging_amqp] idle_timeout = 0	(IntOpt) Timeout for inactive connections (in seconds)
[oslo_messaging_amqp] server_request_prefix = exclusive	(StrOpt) address prefix used when sending to a specific server
[oslo_messaging_amqp] ssl_ca_file =	(StrOpt) CA certificate PEM file to verify server certificate
[oslo_messaging_amqp] ssl_cert_file =	(StrOpt) Identifying certificate PEM file to present to clients
[oslo_messaging_amqp] ssl_key_file =	(StrOpt) Private key PEM file used to sign cert_file certificate
[oslo_messaging_amqp] ssl_key_password = None	(StrOpt) Password for decrypting ssl_key_file (if encrypted)
[oslo_messaging_amqp] trace = False	(BoolOpt) Debug: dump AMQP frames to stdout
[oslo_messaging_qpuid] amqp_auto_delete = False	(BoolOpt) Auto-delete queues in AMQP.
[oslo_messaging_qpuid] amqp_durable_queues = False	(BoolOpt) Use durable queues in AMQP.
[oslo_messaging_qpuid] qpuid_heartbeat = 60	(IntOpt) Seconds between connection keepalive heartbeats.
[oslo_messaging_qpuid] qpuid_hostname = localhost	(StrOpt) Qpid broker hostname.
[oslo_messaging_qpuid] qpuid_hosts = \$qpuid_hostname:\$qpuid_port	(ListOpt) Qpid HA cluster host:port pairs.
[oslo_messaging_qpuid] qpuid_password =	(StrOpt) Password for Qpid connection.
[oslo_messaging_qpuid] qpuid_port = 5672	(IntOpt) Qpid broker port.
[oslo_messaging_qpuid] qpuid_protocol = tcp	(StrOpt) Transport to use, either 'tcp' or 'ssl'.
[oslo_messaging_qpuid] qpuid_receiver_capacity = 1	(IntOpt) The number of prefetched messages held by receiver.
[oslo_messaging_qpuid] qpuid_sasl_mechanisms =	(StrOpt) Space separated list of SASL mechanisms to use for auth.
[oslo_messaging_qpuid] qpuid_tcp_nodelay = True	(BoolOpt) Whether to disable the Nagle algorithm.
[oslo_messaging_qpuid] qpuid_topology_version = 1	(IntOpt) The qpid topology version to use. Version 1 is what was originally used by impl_qpuid. Version 2 includes some backwards-incompatible changes that allow broker federation to work. Users should update to version 2 when they are able to take everything down, as it requires a clean break.
[oslo_messaging_qpuid] qpuid_username =	(StrOpt) Username for Qpid connection.
[oslo_messaging_qpuid] rpc_conn_pool_size = 30	(IntOpt) Size of RPC connection pool.
[oslo_messaging_rabbit] amqp_auto_delete = False	(BoolOpt) Auto-delete queues in AMQP.
[oslo_messaging_rabbit] amqp_durable_queues = False	(BoolOpt) Use durable queues in AMQP.
[oslo_messaging_rabbit] fake_rabbit = False	(BoolOpt) Deprecated, use rpc_backend=kombu+memory or rpc_backend=fake
[oslo_messaging_rabbit] heartbeat_rate = 2	(IntOpt) How often times during the heartbeat_timeout_threshold we check the heartbeat.
[oslo_messaging_rabbit] heartbeat_timeout_threshold = 0	(IntOpt) Number of seconds after which the Rabbit broker is considered down if heartbeat's keep-alive fails (0 disables the heartbeat, >0 enables it. Enabling heartbeats requires kombu>=3.0.7 and amqp>=1.4.0). EXPERIMENTAL

<b>Option = default value</b>	<b>(Type) Help string</b>
[oslo_messaging_rabbit] kombu_reconnect_delay = 1.0	(FloatOpt) How long to wait before reconnecting in response to an AMQP consumer cancel notification.
[oslo_messaging_rabbit] kombu_ssl_ca_certs =	(StrOpt) SSL certification authority file (valid only if SSL enabled).
[oslo_messaging_rabbit] kombu_ssl_certfile =	(StrOpt) SSL cert file (valid only if SSL enabled).
[oslo_messaging_rabbit] kombu_ssl_keyfile =	(StrOpt) SSL key file (valid only if SSL enabled).
[oslo_messaging_rabbit] kombu_ssl_version =	(StrOpt) 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.
[oslo_messaging_rabbit] rabbit_ha_queues = False	(BoolOpt) Use HA queues in RabbitMQ (x-ha-policy: all). If you change this option, you must wipe the RabbitMQ database.
[oslo_messaging_rabbit] rabbit_host = localhost	(StrOpt) The RabbitMQ broker address where a single node is used.
[oslo_messaging_rabbit] rabbit_hosts = \$rabbit_host:\$rabbit_port	(ListOpt) RabbitMQ HA cluster host:port pairs.
[oslo_messaging_rabbit] rabbit_login_method = AMQ-PLAIN	(StrOpt) The RabbitMQ login method.
[oslo_messaging_rabbit] rabbit_max_retries = 0	(IntOpt) Maximum number of RabbitMQ connection retries. Default is 0 (infinite retry count).
[oslo_messaging_rabbit] rabbit_password = guest	(StrOpt) The RabbitMQ password.
[oslo_messaging_rabbit] rabbit_port = 5672	(IntOpt) The RabbitMQ broker port where a single node is used.
[oslo_messaging_rabbit] rabbit_retry_backoff = 2	(IntOpt) How long to backoff for between retries when connecting to RabbitMQ.
[oslo_messaging_rabbit] rabbit_retry_interval = 1	(IntOpt) How frequently to retry connecting with RabbitMQ.
[oslo_messaging_rabbit] rabbit_use_ssl = False	(BoolOpt) Connect over SSL for RabbitMQ.
[oslo_messaging_rabbit] rabbit_userid = guest	(StrOpt) The RabbitMQ userid.
[oslo_messaging_rabbit] rabbit_virtual_host = /	(StrOpt) The RabbitMQ virtual host.
[oslo_messaging_rabbit] rpc_conn_pool_size = 30	(IntOpt) Size of RPC connection pool.
[oslo_middleware] max_request_body_size = 114688	(IntOpt) The maximum body size for each request, in bytes.
[oslo_policy] policy_default_rule = default	(StrOpt) Default rule. Enforced when a requested rule is not found.
[oslo_policy] policy_dirs = ['policy.d']	(MultiStrOpt) 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.
[oslo_policy] policy_file = policy.json	(StrOpt) The JSON file that defines policies.
[polling] partitioning_group_prefix = None	(StrOpt) 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 partitioning_group_prefix a disjoint subset of pollsters should be loaded.
[publisher] telemetry_secret = change this for valid signing	(StrOpt) Secret value for signing messages. Set value empty if signing is not required to avoid computational overhead.
[publisher_notifier] event_topic = event	(StrOpt) The topic that ceilometer uses for event notifications.















Service	Default port	Used by
iSCSI target	3260	OpenStack Block Storage. Required.
MySQL database service	3306	Most OpenStack components.
Message Broker (AMQP traffic)	5672	OpenStack Block Storage, Networking, Orchestration, and Compute.

On some deployments, the default port used by a service may fall within the defined local port range of a host. To check a host's local port range:

```
$ sysctl -a | grep ip_local_port_range
```

If a service's default port falls within this range, run the following program to check if the port has already been assigned to another application:

```
$ lsof -i :PORT
```

Configure the service to use a different port if the default port is already being used by another application.









