



Red Hat OpenStack Platform 10

Configuration Reference

Configuring Red Hat OpenStack Platform environments

Red Hat OpenStack Platform 10 Configuration Reference

Configuring Red Hat OpenStack Platform environments

OpenStack Documentation Team
Red Hat Customer Content Services
rhos-docs@redhat.com

Legal Notice

Copyright © 2017 Red Hat, Inc.

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

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

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

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

Table of Contents

CHAPTER 1. COMMON CONFIGURATIONS	3
1.1. COMMON CONFIGURATION OPTIONS	3
CHAPTER 2. BARE METAL	14
2.1. BARE METAL CONFIGURATION OPTIONS	14
CHAPTER 3. BLOCK STORAGE	47
3.1. BLOCK STORAGE CONFIGURATION OPTIONS	47
CHAPTER 4. IDENTITY SERVICE	114
4.1. IDENTITY SERVICE CONFIGURATION OPTIONS	114
CHAPTER 5. NETWORKING	137
5.1. NETWORKING CONFIGURATION OPTIONS	137
CHAPTER 6. OBJECT STORAGE	169
6.1. OBJECT STORAGE CONFIGURATION OPTIONS	169
CHAPTER 7. TELEMETRY	227
7.1. TELEMETRY CONFIGURATION OPTIONS	227
7.2. TELEMETRY ALARMING CONFIGURATION OPTIONS	241

CHAPTER 1. COMMON CONFIGURATIONS

This chapter describes the common configurations for shared service and libraries.

1.1. COMMON CONFIGURATION OPTIONS

1.1.1. Description of Configuration Options

The following tables provide a comprehensive list of the common configuration options.

Table 1.1. Description of AMQP configuration options

Configuration option = Default value	Description
[DEFAULT]	
control_exchange = openstack	(String) The default exchange under which topics are scoped. May be overridden by an exchange name specified in the <code>transport_url</code> option.
default_publisher_id = None	(String) Default publisher_id for outgoing notifications
transport_url = None	(String) 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 1.2. Description of authorization token configuration options

Configuration option = Default value	Description
[keystone_authtoken]	
admin_password = None	(String) Service user password.
admin_tenant_name = admin	(String) Service tenant name.
admin_token = None	(String) 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 <code>admin_user</code> and <code>admin_password</code> instead.
admin_user = None	(String) Service username.

Configuration option = Default value	Description
auth_admin_prefix =	(String) Prefix to prepend at the beginning of the path. Deprecated, use <code>identity_uri</code> .
auth_host = 127.0.0.1	(String) Host providing the admin Identity API endpoint. Deprecated, use <code>identity_uri</code> .
auth_port = 35357	(Integer) Port of the admin Identity API endpoint. Deprecated, use <code>identity_uri</code> .
auth_protocol = https	(String) Protocol of the admin Identity API endpoint. Deprecated, use <code>identity_uri</code> .
auth_section = None	(Unknown) Config Section from which to load plugin specific options
auth_type = None	(Unknown) Authentication type to load
auth_uri = None	(String) Complete "public" Identity API endpoint. This endpoint should not be an "admin" endpoint, as it should be accessible by all end users. Unauthenticated clients are redirected to this endpoint to authenticate. Although this endpoint should ideally be unversioned, client support in the wild varies. If you're using a versioned v2 endpoint here, then this should <i>not</i> be the same endpoint the service user utilizes for validating tokens, because normal end users may not be able to reach that endpoint.
auth_version = None	(String) API version of the admin Identity API endpoint.
cache = None	(String) Request environment key where the Swift cache object is stored. When <code>auth_token</code> middleware is deployed with a Swift cache, use this option to have the middleware share a caching backend with swift. Otherwise, use the <code>memcached_servers</code> option instead.
cafile = None	(String) A PEM encoded Certificate Authority to use when verifying HTTPs connections. Defaults to system CAs.
certfile = None	(String) Required if identity server requires client certificate
check_revocations_for_cached = False	(Boolean) 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 = False	(Boolean) Do not handle authorization requests within the middleware, but delegate the authorization decision to downstream WSGI components.

Configuration option = Default value	Description
enforce_token_bind = permissive	(String) 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 = md5	(List) 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.
http_connect_timeout = None	(Integer) Request timeout value for communicating with Identity API server.
http_request_max_retries = 3	(Integer) How many times are we trying to reconnect when communicating with Identity API Server.
identity_uri = None	(String) Complete admin Identity API endpoint. This should specify the unversioned root endpoint e.g. https://localhost:35357/
include_service_catalog = True	(Boolean) (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.
insecure = False	(Boolean) Verify HTTPS connections.
keyfile = None	(String) Required if identity server requires client certificate
memcache_pool_conn_get_timeout = 10	(Integer) (Optional) Number of seconds that an operation will wait to get a memcached client connection from the pool.
memcache_pool_dead_retry = 300	(Integer) (Optional) Number of seconds memcached server is considered dead before it is tried again.
memcache_pool_maxsize = 10	(Integer) (Optional) Maximum total number of open connections to every memcached server.

Configuration option = Default value	Description
memcache_pool_socket_timeout = 3	(Integer) (Optional) Socket timeout in seconds for communicating with a memcached server.
memcache_pool_unused_timeout = 60	(Integer) (Optional) Number of seconds a connection to memcached is held unused in the pool before it is closed.
memcache_secret_key = None	(String) (Optional, mandatory if memcache_security_strategy is defined) This string is used for key derivation.
memcache_security_strategy = None	(String) (Optional) If defined, indicate whether token data should be authenticated or authenticated and encrypted. If MAC, token data is authenticated (with HMAC) in the cache. If ENCRYPT, token data is encrypted and authenticated in the cache. If the value is not one of these options or empty, auth_token will raise an exception on initialization.
memcache_use_advanced_pool = False	(Boolean) (Optional) Use the advanced (eventlet safe) memcached client pool. The advanced pool will only work under python 2.x.
memcached_servers = None	(List) Optionally specify a list of memcached server(s) to use for caching. If left undefined, tokens will instead be cached in-process.
region_name = None	(String) The region in which the identity server can be found.
revocation_cache_time = 10	(Integer) Determines the frequency at which the list of revoked tokens is retrieved from the Identity service (in seconds). A high number of revocation events combined with a low cache duration may significantly reduce performance. Only valid for PKI tokens.
signing_dir = None	(String) Directory used to cache files related to PKI tokens.
token_cache_time = 300	(Integer) 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 1.3. Description of authentication configuration options

Configuration option = Default value	Description
[DEFAULT]	
auth_strategy = keystone	(String) This determines the strategy to use for authentication: keystone or noauth2. 'noauth2' is designed for testing only, as it does no actual credential checking. 'noauth2' provides administrative credentials only if 'admin' is specified as the username.

Table 1.4. Description of database configuration options

Configuration option = Default value	Description
[DEFAULT]	
db_driver = SERVICE.db	(String) DEPRECATED: The driver to use for database access
[database]	
backend = sqlalchemy	(String) The back end to use for the database.
connection = None	(String) The SQLAlchemy connection string to use to connect to the database.
connection_debug = 0	(Integer) Verbosity of SQL debugging information: 0=None, 100=Everything.
connection_trace = False	(Boolean) Add Python stack traces to SQL as comment strings.
db_inc_retry_interval = True	(Boolean) If True, increases the interval between retries of a database operation up to db_max_retry_interval.
db_max_retries = 20	(Integer) Maximum retries in case of connection error or deadlock error before error is raised. Set to -1 to specify an infinite retry count.
db_max_retry_interval = 10	(Integer) If db_inc_retry_interval is set, the maximum seconds between retries of a database operation.
db_retry_interval = 1	(Integer) Seconds between retries of a database transaction.

Configuration option = Default value	Description
idle_timeout = 3600	(Integer) Timeout before idle SQL connections are reaped.
max_overflow = 50	(Integer) If set, use this value for max_overflow with SQLAlchemy.
max_pool_size = None	(Integer) Maximum number of SQL connections to keep open in a pool.
max_retries = 10	(Integer) Maximum number of database connection retries during startup. Set to -1 to specify an infinite retry count.
min_pool_size = 1	(Integer) Minimum number of SQL connections to keep open in a pool.
mysql_sql_mode = TRADITIONAL	(String) 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=
pool_timeout = None	(Integer) If set, use this value for pool_timeout with SQLAlchemy.
retry_interval = 10	(Integer) Interval between retries of opening a SQL connection.
slave_connection = None	(String) The SQLAlchemy connection string to use to connect to the slave database.
sqlite_db = oslo.sqlite	(String) The file name to use with SQLite.
sqlite_synchronous = True	(Boolean) If True, SQLite uses synchronous mode.
use_db_reconnect = False	(Boolean) Enable the experimental use of database reconnect on connection lost.
use_tpool = False	(Boolean) Enable the experimental use of thread pooling for all DB API calls

Table 1.5. Description of common logging configuration options

Configuration option = Default value	Description
[DEFAULT]	

Configuration option = Default value	Description
debug = False	(Boolean) If set to true, the logging level will be set to DEBUG instead of the default INFO level.
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, requests.packages.urllib3.util.retry=WARN, urllib3.util.retry=WARN, keystonemiddleware=WARN, routes.middleware=WARN, stevedore=WARN, taskflow=WARN, keystoneauth=WARN, oslo.cache=INFO, dogpile.core.dogpile=INFO	(List) List of package logging levels in logger=LEVEL pairs. This option is ignored if log_config_append is set.
fatal_deprecations = False	(Boolean) Enables or disables fatal status of deprecations.
fatal_exception_format_errors = False	(Boolean) Make exception message format errors fatal
instance_format = "[instance: %(uuid)s] "	(String) The format for an instance that is passed with the log message.
instance_uuid_format = "[instance: %(uuid)s] "	(String) The format for an instance UUID that is passed with the log message.
log_config_append = None	(String) The name of a logging configuration file. This file is appended to any existing logging configuration files. For details about logging configuration files, see the Python logging module documentation. Note that when logging configuration files are used then all logging configuration is set in the configuration file and other logging configuration options are ignored (for example, logging_context_format_string).
log_date_format = %Y-%m-%d %H:%M:%S	(String) Defines the format string for %(asctime)s in log records. Default: %(default)s . This option is ignored if log_config_append is set.
log_dir = None	(String) (Optional) The base directory used for relative log_file paths. This option is ignored if log_config_append is set.
log_file = None	(String) (Optional) Name of log file to send logging output to. If no default is set, logging will go to stderr as defined by use_stderr. This option is ignored if log_config_append is set.

Configuration option = Default value	Description
logging_context_format_string = % (asctime)s.%(msecs)03d %(process)d % (levelname)s %(name)s [%(request_id)s % (user_identity)s] %(instance)s%(message)s	(String) Format string to use for log messages with context.
logging_debug_format_suffix = % (funcName)s %(pathname)s:%(lineno)d	(String) Additional data to append to log message when logging level for the message is DEBUG.
logging_default_format_string = % (asctime)s.%(msecs)03d %(process)d % (levelname)s %(name)s [-] %(instance)s% (message)s	(String) Format string to use for log messages when context is undefined.
logging_exception_prefix = %(asctime)s.% (msecs)03d %(process)d ERROR %(name)s %(instance)s	(String) Prefix each line of exception output with this format.
logging_user_identity_format = %(user)s % (tenant)s %(domain)s %(user_domain)s % (project_domain)s	(String) Defines the format string for % (user_identity)s that is used in logging_context_format_string.
publish_errors = False	(Boolean) Enables or disables publication of error events.
syslog_log_facility = LOG_USER	(String) Syslog facility to receive log lines. This option is ignored if log_config_append is set.
use_stderr = True	(Boolean) Log output to standard error. This option is ignored if log_config_append is set.
use_syslog = False	(Boolean) Use syslog for logging. Existing syslog format is DEPRECATED and will be changed later to honor RFC5424. This option is ignored if log_config_append is set.
verbose = True	(Boolean) DEPRECATED: If set to false, the logging level will be set to WARNING instead of the default INFO level.
watch_log_file = False	(Boolean) Uses logging handler designed to watch file system. When log file is moved or removed this handler will open a new log file with specified path instantaneously. It makes sense only if log_file option is specified and Linux platform is used. This option is ignored if log_config_append is set.

Table 1.6. Description of policy configuration options

Configuration option = Default value	Description
[oslo_policy]	
policy_default_rule = default	(String) Default rule. Enforced when a requested rule is not found.
policy_dirs = ['policy.d']	(Multi-valued) Directories where policy configuration files are stored. They can be relative to any directory in the search path defined by the config_dir option, or absolute paths. The file defined by policy_file must exist for these directories to be searched. Missing or empty directories are ignored.
policy_file = policy.json	(String) The JSON file that defines policies.

Table 1.7. Description of RPC configuration options

Configuration option = Default value	Description
[DEFAULT]	
notification_format = both	(String) Specifies which notification format shall be used by nova.
rpc_backend = rabbit	(String) The messaging driver to use, defaults to rabbit. Other drivers include amqp and zmq.
rpc_cast_timeout = -1	(Integer) Seconds to wait before a cast expires (TTL). The default value of -1 specifies an infinite linger period. The value of 0 specifies no linger period. Pending messages shall be discarded immediately when the socket is closed. Only supported by impl_zmq.
rpc_conn_pool_size = 30	(Integer) Size of RPC connection pool.
rpc_poll_timeout = 1	(Integer) The default number of seconds that poll should wait. Poll raises timeout exception when timeout expired.
rpc_response_timeout = 60	(Integer) Seconds to wait for a response from a call.
[cells]	

Configuration option = Default value	Description
rpc_driver_queue_base = cells.intercell	(String) RPC driver queue base When sending a message to another cell by JSON-ifying the message and making an RPC cast to 'process_message', a base queue is used. This option defines the base queue name to be used when communicating between cells. Various topics by message type will be appended to this. Possible values: * The base queue name to be used when communicating between cells. Services which consume this: * nova-cells Related options: * None
[oslo_concurrency]	
disable_process_locking = False	(Boolean) Enables or disables inter-process locks.
lock_path = None	(String) 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.
[oslo_messaging]	
event_stream_topic = neutron_lbaas_event	(String) topic name for receiving events from a queue
[oslo_messaging_amqp]	
allow_insecure_clients = False	(Boolean) Accept clients using either SSL or plain TCP
broadcast_prefix = broadcast	(String) address prefix used when broadcasting to all servers
container_name = None	(String) Name for the AMQP container
group_request_prefix = unicast	(String) address prefix when sending to any server in group
idle_timeout = 0	(Integer) Timeout for inactive connections (in seconds)
password =	(String) Password for message broker authentication
sasl_config_dir =	(String) Path to directory that contains the SASL configuration
sasl_config_name =	(String) Name of configuration file (without .conf suffix)

Configuration option = Default value	Description
sasl_mechanisms =	(String) Space separated list of acceptable SASL mechanisms
server_request_prefix = exclusive	(String) address prefix used when sending to a specific server
ssl_ca_file =	(String) CA certificate PEM file to verify server certificate
ssl_cert_file =	(String) Identifying certificate PEM file to present to clients
ssl_key_file =	(String) Private key PEM file used to sign cert_file certificate
ssl_key_password = None	(String) Password for decrypting ssl_key_file (if encrypted)
trace = False	(Boolean) Debug: dump AMQP frames to stdout
username =	(String) User name for message broker authentication
[oslo_messaging_notifications]	
driver = []	(Multi-valued) The Drivers(s) to handle sending notifications. Possible values are messaging, messagingv2, routing, log, test, noop
topics = notifications	(List) AMQP topic used for OpenStack notifications.
transport_url = None	(String) A URL representing the messaging driver to use for notifications. If not set, we fall back to the same configuration used for RPC.
[upgrade_levels]	
baseapi = None	(String) Set a version cap for messages sent to the base api in any service

CHAPTER 2. 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`.

2.1. BARE METAL CONFIGURATION OPTIONS

2.1.1. Description of Configuration Options

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

Table 2.1. Description of agent configuration options

Configuration option = Default value	Description
[agent]	
agent_api_version = v1	(String) API version to use for communicating with the ramdisk agent.
deploy_logs_collect = on_failure	(String) Whether Ironic should collect the deployment logs on deployment failure (<code>on_failure</code>), always or never.
deploy_logs_local_path = /var/log/ironic/deploy	(String) The path to the directory where the logs should be stored, used when the <code>deploy_logs_storage_backend</code> is configured to "local".
deploy_logs_storage_backend = local	(String) The name of the storage backend where the logs will be stored.
deploy_logs_swift_container = ironic_deploy_logs_container	(String) The name of the Swift container to store the logs, used when the <code>deploy_logs_storage_backend</code> is configured to "swift".
deploy_logs_swift_days_to_expire = 30	(Integer) Number of days before a log object is marked as expired in Swift. If None, the logs will be kept forever or until manually deleted. Used when the <code>deploy_logs_storage_backend</code> is configured to "swift".
manage_agent_boot = True	(Boolean) Whether Ironic will manage booting of the agent ramdisk. If set to False, you will need to configure your mechanism to allow booting the agent ramdisk.
memory_consumed_by_agent = 0	(Integer) The memory size in MiB consumed by agent when it is booted on a bare metal node. This is used for checking if the image can be downloaded and deployed on the bare metal node after booting agent ramdisk. This may be set according to the memory consumed by the agent ramdisk image.

Configuration option = Default value	Description
post_deploy_get_power_state_retries = 6	(Integer) Number of times to retry getting power state to check if bare metal node has been powered off after a soft power off.
post_deploy_get_power_state_retry_interval = 5	(Integer) Amount of time (in seconds) to wait between polling power state after trigger soft poweroff.
stream_raw_images = True	(Boolean) Whether the agent ramdisk should stream raw images directly onto the disk or not. By streaming raw images directly onto the disk the agent ramdisk will not spend time copying the image to a tmpfs partition (therefore consuming less memory) prior to writing it to the disk. Unless the disk where the image will be copied to is really slow, this option should be set to True. Defaults to True.

Table 2.2. Description of AMT configuration options

Configuration option = Default value	Description
[amt]	
action_wait = 10	(Integer) Amount of time (in seconds) to wait, before retrying an AMT operation
awake_interval = 60	(Integer) Time interval (in seconds) for successive awake call to AMT interface, this depends on the IdleTimeout setting on AMT interface. AMT Interface will go to sleep after 60 seconds of inactivity by default. IdleTimeout=0 means AMT will not go to sleep at all. Setting awake_interval=0 will disable awake call.
max_attempts = 3	(Integer) Maximum number of times to attempt an AMT operation, before failing
protocol = http	(String) Protocol used for AMT endpoint

Table 2.3. Description of API configuration options

Configuration option = Default value	Description
[api]	
api_workers = None	(Integer) Number of workers for OpenStack Ironic API service. The default is equal to the number of CPUs available if that can be determined, else a default worker count of 1 is returned.
enable_ssl_api = False	(Boolean) Enable the integrated stand-alone API to service requests via HTTPS instead of HTTP. If there is a front-end service performing HTTPS offloading from the service, this option should be False; note, you will want to change public API endpoint to represent SSL termination URL with 'public_endpoint' option.
host_ip = 0.0.0.0	(String) The IP address on which ironic-api listens.
max_limit = 1000	(Integer) The maximum number of items returned in a single response from a collection resource.
port = 6385	(Port number) The TCP port on which ironic-api listens.
public_endpoint = None	(String) Public URL to use when building the links to the API resources (for example, " https://ironic.rocks:6384 "). If None the links will be built using the request's host URL. If the API is operating behind a proxy, you will want to change this to represent the proxy's URL. Defaults to None.
ramdisk_heartbeat_timeout = 300	(Integer) Maximum interval (in seconds) for agent heartbeats.
restrict_lookup = True	(Boolean) Whether to restrict the lookup API to only nodes in certain states.
[oslo_middleware]	
enable_proxy_headers_parsing = False	(Boolean) Whether the application is behind a proxy or not. This determines if the middleware should parse the headers or not.
max_request_body_size = 114688	(Integer) The maximum body size for each request, in bytes.
secure_proxy_ssl_header = X-Forwarded-Proto	(String) DEPRECATED: The HTTP Header that will be used to determine what the original request protocol scheme was, even if it was hidden by a SSL termination proxy.
[oslo_versionedobjects]	

Configuration option = Default value	Description
fatal_exception_format_errors = False	(Boolean) Make exception message format errors fatal

Table 2.4. Description of audit configuration options

Configuration option = Default value	Description
[audit]	
audit_map_file = /etc/ironic/ironic_api_audit_map.conf	(String) Path to audit map file for ironic-api service. Used only when API audit is enabled.
enabled = False	(Boolean) Enable auditing of API requests (for ironic-api service).
ignore_req_list = None	(String) Comma separated list of Ironic REST API HTTP methods to be ignored during audit. For example: auditing will not be done on any GET or POST requests if this is set to "GET,POST". It is used only when API audit is enabled.
namespace = openstack	(String) namespace prefix for generated id
[audit_middleware_notifications]	
driver = None	(String) The Driver to handle sending notifications. Possible values are messaging, messagingv2, routing, log, test, noop. If not specified, then value from oslo_messaging_notifications conf section is used.
topics = None	(List) List of AMQP topics used for OpenStack notifications. If not specified, then value from oslo_messaging_notifications conf section is used.
transport_url = None	(String) A URL representing messaging driver to use for notification. If not specified, we fall back to the same configuration used for RPC.

Table 2.5. Description of Cisco UCS configuration options

Configuration option = Default value	Description
[cimc]	

Configuration option = Default value	Description
action_interval = 10	(Integer) Amount of time in seconds to wait in between power operations
max_retry = 6	(Integer) Number of times a power operation needs to be retried
[cisco_ucs]	
action_interval = 5	(Integer) Amount of time in seconds to wait in between power operations
max_retry = 6	(Integer) Number of times a power operation needs to be retried

Table 2.6. Description of common configuration options

Configuration option = Default value	Description
[DEFAULT]	
bindir = /usr/local/bin	(String) Directory where ironic binaries are installed.
debug_tracebacks_in_api = False	(Boolean) Return server tracebacks in the API response for any error responses. WARNING: this is insecure and should not be used in a production environment.
default_network_interface = None	(String) Default network interface to be used for nodes that do not have network_interface field set. A complete list of network interfaces present on your system may be found by enumerating the "ironic.hardware.interfaces.network" endpoint.
enabled_drivers = pxe_ipmitool	(List) Specify the list of drivers to load during service initialization. Missing drivers, or drivers which fail to initialize, will prevent the conductor service from starting. The option default is a recommended set of production-oriented drivers. A complete list of drivers present on your system may be found by enumerating the "ironic.drivers" endpoint. An example may be found in the developer documentation online.

Configuration option = Default value	Description
enabled_network_interfaces = flat, noop	(List) Specify the list of network interfaces to load during service initialization. Missing network interfaces, or network interfaces which fail to initialize, will prevent the conductor service from starting. The option default is a recommended set of production-oriented network interfaces. A complete list of network interfaces present on your system may be found by enumerating the "ironic.hardware.interfaces.network" endpoint. This value must be the same on all ironic-conductor and ironic-api services, because it is used by ironic-api service to validate a new or updated node's network_interface value.
executor_thread_pool_size = 64	(Integer) Size of executor thread pool.
fatal_exception_format_errors = False	(Boolean) Used if there is a formatting error when generating an exception message (a programming error). If True, raise an exception; if False, use the unformatted message.
force_raw_images = True	(Boolean) If True, convert backing images to "raw" disk image format.
grub_config_template = \$pybasedir/common/grub_conf.template	(String) Template file for grub configuration file.
hash_distribution_replicas = 1	(Integer) [Experimental Feature] Number of hosts to map onto each hash partition. Setting this to more than one will cause additional conductor services to prepare deployment environments and potentially allow the Ironic cluster to recover more quickly if a conductor instance is terminated.
hash_partition_exponent = 5	(Integer) Exponent to determine number of hash partitions to use when distributing load across conductors. Larger values will result in more even distribution of load and less load when rebalancing the ring, but more memory usage. Number of partitions per conductor is $(2^{\text{hash_partition_exponent}})$. This determines the granularity of rebalancing: given 10 hosts, and an exponent of the 2, there are 40 partitions in the ring. A few thousand partitions should make rebalancing smooth in most cases. The default is suitable for up to a few hundred conductors. Too many partitions has a CPU impact.
hash_ring_reset_interval = 180	(Integer) Interval (in seconds) between hash ring resets.
host = localhost	(String) Name of this node. This can be an opaque identifier. It is not necessarily a hostname, FQDN, or IP address. However, the node name must be valid within an AMQP key.

Configuration option = Default value	Description
isolinux_bin = /usr/lib/syslinux/isolinux.bin	(String) Path to isolinux binary file.
isolinux_config_template = \$pybasedir/common/isolinux_config.template	(String) Template file for isolinux configuration file.
my_ip = 127.0.0.1	(String) IP address of this host. If unset, will determine the IP programmatically. If unable to do so, will use "127.0.0.1".
notification_level = None	(String) Specifies the minimum level for which to send notifications. If not set, no notifications will be sent. The default is for this option to be unset.
parallel_image_downloads = False	(Boolean) Run image downloads and raw format conversions in parallel.
pybasedir = /usr/lib/python/site-packages/ironic/ironic	(String) Directory where the ironic python module is installed.
rootwrap_config = /etc/ironic/rootwrap.conf	(String) Path to the rootwrap configuration file to use for running commands as root.
state_path = \$pybasedir	(String) Top-level directory for maintaining ironic's state.
tempdir = /tmp	(String) Temporary working directory, default is Python temp dir.
[ironic_lib]	
fatal_exception_format_errors = False	(Boolean) Make exception message format errors fatal.
root_helper = sudo ironic-rootwrap /etc/ironic/rootwrap.conf	(String) Command that is prefixed to commands that are run as root. If not specified, no commands are run as root.

Table 2.7. Description of conductor configuration options

Configuration option = Default value	Description
[conductor]	
api_url = None	(String) URL of Ironic API service. If not set ironic can get the current value from the keystone service catalog.

Configuration option = Default value	Description
automated_clean = True	(Boolean) Enables or disables automated cleaning. Automated cleaning is a configurable set of steps, such as erasing disk drives, that are performed on the node to ensure it is in a baseline state and ready to be deployed to. This is done after instance deletion as well as during the transition from a "manageable" to "available" state. When enabled, the particular steps performed to clean a node depend on which driver that node is managed by; see the individual driver's documentation for details. NOTE: The introduction of the cleaning operation causes instance deletion to take significantly longer. In an environment where all tenants are trusted (eg, because there is only one tenant), this option could be safely disabled.
check_provision_state_interval = 60	(Integer) Interval between checks of provision timeouts, in seconds.
clean_callback_timeout = 1800	(Integer) Timeout (seconds) to wait for a callback from the ramdisk doing the cleaning. If the timeout is reached the node will be put in the "clean failed" provision state. Set to 0 to disable timeout.
configdrive_swift_container = ironic_configdrive_container	(String) Name of the Swift container to store config drive data. Used when configdrive_use_swift is True.
configdrive_use_swift = False	(Boolean) Whether to upload the config drive to Swift.
deploy_callback_timeout = 1800	(Integer) Timeout (seconds) to wait for a callback from a deploy ramdisk. Set to 0 to disable timeout.
force_power_state_during_sync = True	(Boolean) During sync_power_state, should the hardware power state be set to the state recorded in the database (True) or should the database be updated based on the hardware state (False).
heartbeat_interval = 10	(Integer) Seconds between conductor heart beats.
heartbeat_timeout = 60	(Integer) Maximum time (in seconds) since the last check-in of a conductor. A conductor is considered inactive when this time has been exceeded.
inspect_timeout = 1800	(Integer) Timeout (seconds) for waiting for node inspection. 0 - unlimited.
node_locked_retry_attempts = 3	(Integer) Number of attempts to grab a node lock.
node_locked_retry_interval = 1	(Integer) Seconds to sleep between node lock attempts.

Configuration option = Default value	Description
periodic_max_workers = 8	(Integer) Maximum number of worker threads that can be started simultaneously by a periodic task. Should be less than RPC thread pool size.
power_state_sync_max_retries = 3	(Integer) During sync_power_state failures, limit the number of times Ironic should try syncing the hardware node power state with the node power state in DB
send_sensor_data = False	(Boolean) Enable sending sensor data message via the notification bus
send_sensor_data_interval = 600	(Integer) Seconds between conductor sending sensor data message to ceilometer via the notification bus.
send_sensor_data_types = ALL	(List) 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.
sync_local_state_interval = 180	(Integer) 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.
sync_power_state_interval = 60	(Integer) Interval between syncing the node power state to the database, in seconds.
workers_pool_size = 100	(Integer) The size of the workers greenthread pool. Note that 2 threads will be reserved by the conductor itself for handling heart beats and periodic tasks.

Table 2.8. Description of console configuration options

Configuration option = Default value	Description
[console]	
subprocess_checking_interval = 1	(Integer) Time interval (in seconds) for checking the status of console subprocess.
subprocess_timeout = 10	(Integer) Time (in seconds) to wait for the console subprocess to start.

Configuration option = Default value	Description
terminal = shellinaboxd	(String) Path to serial console terminal program. Used only by Shell In A Box console.
terminal_cert_dir = None	(String) Directory containing the terminal SSL cert (PEM) for serial console access. Used only by Shell In A Box console.
terminal_pid_dir = None	(String) Directory for holding terminal pid files. If not specified, the temporary directory will be used.

Table 2.9. Description of logging configuration options

Configuration option = Default value	Description
[DEFAULT]	
pecan_debug = False	(Boolean) Enable pecan debug mode. WARNING: this is insecure and should not be used in a production environment.

Table 2.10. Description of deploy configuration options

Configuration option = Default value	Description
[deploy]	
continue_if_disk_secure_erase_fails = False	(Boolean) Defines what to do if an ATA secure erase operation fails during cleaning in the Ironic Python Agent. If False, the cleaning operation will fail and the node will be put in clean failed state. If True, shred will be invoked and cleaning will continue.
erase_devices_metadata_priority = None	(Integer) Priority to run in-band clean step that erases metadata from devices, via the Ironic Python Agent ramdisk. If unset, will use the priority set in the ramdisk (defaults to 99 for the GenericHardwareManager). If set to 0, will not run during cleaning.
erase_devices_priority = None	(Integer) 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.

Configuration option = Default value	Description
http_root = /httpboot	(String) ironic-conductor node's HTTP root path.
http_url = None	(String) ironic-conductor node's HTTP server URL. Example: http://192.1.2.3:8080
power_off_after_deploy_failure = True	(Boolean) Whether to power off a node after deploy failure. Defaults to True.
shred_final_overwrite_with_zeros = True	(Boolean) Whether to write zeros to a node's block devices after writing random data. This will write zeros to the device even when <code>deploy.shred_random_overwrite_iterations</code> is 0. This option is only used if a device could not be ATA Secure Erased. Defaults to True.
shred_random_overwrite_iterations = 1	(Integer) During shred, overwrite all block devices N times with random data. This is only used if a device could not be ATA Secure Erased. Defaults to 1.

Table 2.11. Description of DHCP configuration options

Configuration option = Default value	Description
[dhcp]	
dhcp_provider = neutron	(String) DHCP provider to use. "neutron" uses Neutron, and "none" uses a no-op provider.

Table 2.12. Description of disk partitioner configuration options

Configuration option = Default value	Description
[disk_partitioner]	
check_device_interval = 1	(Integer) After Ironic has completed creating the partition table, it continues to check for activity on the attached iSCSI device status at this interval prior to copying the image to the node, in seconds
check_device_max_retries = 20	(Integer) The maximum number of times to check that the device is not accessed by another process. If the device is still busy after that, the disk partitioning will be treated as having failed.

Configuration option = Default value	Description
[disk_utils]	
bios_boot_partition_size = 1	(Integer) Size of BIOS Boot partition in MiB when configuring GPT partitioned systems for local boot in BIOS.
dd_block_size = 1M	(String) Block size to use when writing to the nodes disk.
efi_system_partition_size = 200	(Integer) Size of EFI system partition in MiB when configuring UEFI systems for local boot.
iscsi_verify_attempts = 3	(Integer) Maximum attempts to verify an iSCSI connection is active, sleeping 1 second between attempts.

Table 2.13. Description of DRAC configuration options

Configuration option = Default value	Description
[drac]	
query_raid_config_job_status_interval = 120	(Integer) Interval (in seconds) between periodic RAID job status checks to determine whether the asynchronous RAID configuration was successfully finished or not.

Table 2.14. Description of glance configuration options

Configuration option = Default value	Description
[glance]	
allowed_direct_url_schemes =	(List) A list of URL schemes that can be downloaded directly via the direct_url. Currently supported schemes: [file].
auth_section = None	(Unknown) Config Section from which to load plugin specific options
auth_strategy = keystone	(String) Authentication strategy to use when connecting to glance.

Configuration option = Default value	Description
auth_type = None	(Unknown) Authentication type to load
cafile = None	(String) PEM encoded Certificate Authority to use when verifying HTTPs connections.
certfile = None	(String) PEM encoded client certificate cert file
glance_api_insecure = False	(Boolean) Allow to perform insecure SSL (https) requests to glance.
glance_api_servers = None	(List) A list of the glance api servers available to ironic. Prefix with https:// for SSL-based glance API servers. Format is [hostname IP]:port.
glance_cafile = None	(String) Optional path to a CA certificate bundle to be used to validate the SSL certificate served by glance. It is used when glance_api_insecure is set to False.
glance_host = \$my_ip	(String) Default glance hostname or IP address.
glance_num_retries = 0	(Integer) Number of retries when downloading an image from glance.
glance_port = 9292	(Port number) Default glance port.
glance_protocol = http	(String) Default protocol to use when connecting to glance. Set to https for SSL.
insecure = False	(Boolean) Verify HTTPS connections.
keyfile = None	(String) PEM encoded client certificate key file
swift_account = None	(String) The account that Glance uses to communicate with Swift. The format is "AUTH_uuid". "uuid" is the UUID for the account configured in the glance-api.conf. Required for temporary URLs when Glance backend is Swift. For example: "AUTH_a422b2-91f3-2f46-74b7-d7c9e8958f5d30". Swift temporary URL format: "endpoint_url/api_version/[account/]container/object_id"
swift_api_version = v1	(String) The Swift API version to create a temporary URL for. Defaults to "v1". Swift temporary URL format: "endpoint_url/api_version/[account/]container/object_id"

Configuration option = Default value	Description
swift_container = glance	(String) The Swift container Glance is configured to store its images in. Defaults to "glance", which is the default in glance-api.conf. Swift temporary URL format: "endpoint_url/api_version/[account/]container/object_id"
swift_endpoint_url = None	(String) The "endpoint" (scheme, hostname, optional port) for the Swift URL of the form "endpoint_url/api_version/[account/]container/object_id". Do not include trailing "/". For example, use " https://swift.example.com ". If using RADOS Gateway, endpoint may also contain /swift path; if it does not, it will be appended. Required for temporary URLs.
swift_store_multiple_containers_seed = 0	(Integer) This should match a config by the same name in the Glance configuration file. When set to 0, a single-tenant store will only use one container to store all images. When set to an integer value between 1 and 32, a single-tenant store will use multiple containers to store images, and this value will determine how many containers are created.
swift_temp_url_cache_enabled = False	(Boolean) Whether to cache generated Swift temporary URLs. Setting it to true is only useful when an image caching proxy is used. Defaults to False.
swift_temp_url_duration = 1200	(Integer) The length of time in seconds that the temporary URL will be valid for. Defaults to 20 minutes. If some deploys get a 401 response code when trying to download from the temporary URL, try raising this duration. This value must be greater than or equal to the value for swift_temp_url_expected_download_start_delay
swift_temp_url_expected_download_start_delay = 0	(Integer) This is the delay (in seconds) from the time of the deploy request (when the Swift temporary URL is generated) to when the IPA ramdisk starts up and URL is used for the image download. This value is used to check if the Swift temporary URL duration is large enough to let the image download begin. Also if temporary URL caching is enabled this will determine if a cached entry will still be valid when the download starts. swift_temp_url_duration value must be greater than or equal to this option's value. Defaults to 0.
swift_temp_url_key = None	(String) The secret token given to Swift to allow temporary URL downloads. Required for temporary URLs.
temp_url_endpoint_type = swift	(String) Type of endpoint to use for temporary URLs. If the Glance backend is Swift, use "swift"; if it is CEPH with RADOS gateway, use "radosgw".

Configuration option = Default value	Description
timeout = None	(Integer) Timeout value for http requests

Table 2.15. Description of iLO configuration options

Configuration option = Default value	Description
[ilo]	
ca_file = None	(String) CA certificate file to validate iLO.
clean_priority_clear_secure_boot_keys = 0	(Integer) Priority for clear_secure_boot_keys clean step. This step is not enabled by default. It can be enabled to clear all secure boot keys enrolled with iLO.
clean_priority_erase_devices = None	(Integer) DEPRECATED: Priority for erase devices clean step. If unset, it defaults to 10. If set to 0, the step will be disabled and will not run during cleaning. This configuration option is duplicated by [deploy] erase_devices_priority, please use that instead.
clean_priority_reset_bios_to_default = 10	(Integer) Priority for reset_bios_to_default clean step.
clean_priority_reset_ilo = 0	(Integer) Priority for reset_ilo clean step.
clean_priority_reset_ilo_credential = 30	(Integer) Priority for reset_ilo_credential clean step. This step requires "ilo_change_password" parameter to be updated in nodes's driver_info with the new password.
clean_priority_reset_secure_boot_keys_to_default = 20	(Integer) Priority for reset_secure_boot_keys clean step. This step will reset the secure boot keys to manufacturing defaults.
client_port = 443	(Port number) Port to be used for iLO operations
client_timeout = 60	(Integer) Timeout (in seconds) for iLO operations
default_boot_mode = auto	(String) Default boot mode to be used in provisioning when "boot_mode" capability is not provided in the "properties/capabilities" of the node. The default is "auto" for backward compatibility. When "auto" is specified, default boot mode will be selected based on boot mode settings on the system.

Configuration option = Default value	Description
power_retry = 6	(Integer) Number of times a power operation needs to be retried
power_wait = 2	(Integer) Amount of time in seconds to wait in between power operations
swift_ilo_container = ironic_ilo_container	(String) The Swift iLO container to store data.
swift_object_expiry_timeout = 900	(Integer) Amount of time in seconds for Swift objects to auto-expire.
use_web_server_for_images = False	(Boolean) Set this to True to use http web server to host floppy images and generated boot ISO. This requires http_root and http_url to be configured in the [deploy] section of the config file. If this is set to False, then Ironic will use Swift to host the floppy images and generated boot_iso.

Table 2.16. Description of inspector configuration options

Configuration option = Default value	Description
[inspector]	
auth_section = None	(Unknown) Config Section from which to load plugin specific options
auth_type = None	(Unknown) Authentication type to load
cafile = None	(String) PEM encoded Certificate Authority to use when verifying HTTPS connections.
certfile = None	(String) PEM encoded client certificate cert file
enabled = False	(Boolean) whether to enable inspection using ironic-inspector
insecure = False	(Boolean) Verify HTTPS connections.
keyfile = None	(String) PEM encoded client certificate key file
service_url = None	(String) ironic-inspector HTTP endpoint. If this is not set, the service catalog will be used.
status_check_period = 60	(Integer) period (in seconds) to check status of nodes on inspection

Configuration option = Default value	Description
timeout = None	(Integer) Timeout value for http requests

Table 2.17. Description of IPMI configuration options

Configuration option = Default value	Description
[ipmi]	
min_command_interval = 5	(Integer) Minimum time, in seconds, between IPMI operations sent to a server. There is a risk with some hardware that setting this too low may cause the BMC to crash. Recommended setting is 5 seconds.
retry_timeout = 60	(Integer) Maximum time in seconds to retry IPMI operations. There is a tradeoff when setting this value. Setting this too low may cause older BMCs to crash and require a hard reset. However, setting too high can cause the sync power state periodic task to hang when there are slow or unresponsive BMCs.

Table 2.18. Description of iRMC configuration options

Configuration option = Default value	Description
[irmc]	
auth_method = basic	(String) Authentication method to be used for iRMC operations
client_timeout = 60	(Integer) Timeout (in seconds) for iRMC operations
port = 443	(Port number) Port to be used for iRMC operations
remote_image_server = None	(String) IP of remote image server
remote_image_share_name = share	(String) share name of remote_image_server
remote_image_share_root = /remote_image_share_root	(String) Ironic conductor node's "NFS" or "CIFS" root path
remote_image_share_type = CIFS	(String) Share type of virtual media

Configuration option = Default value	Description
remote_image_user_domain =	(String) Domain name of remote_image_user_name
remote_image_user_name = None	(String) User name of remote_image_server
remote_image_user_password = None	(String) Password of remote_image_user_name
sensor_method = ipmitool	(String) Sensor data retrieval method.
snmp_community = public	(String) SNMP community. Required for versions "v1" and "v2c"
snmp_port = 161	(Port number) SNMP port
snmp_security = None	(String) SNMP security name. Required for version "v3"
snmp_version = v2c	(String) SNMP protocol version

Table 2.19. Description of iSCSI configuration options

Configuration option = Default value	Description
[iscsi]	
portal_port = 3260	(Port number) The port number on which the iSCSI portal listens for incoming connections.

Table 2.20. Description of keystone configuration options

Configuration option = Default value	Description
[keystone]	
region_name = None	(String) The region used for getting endpoints of OpenStack services.

Table 2.21. Description of metrics statsd configuration options

Configuration option = Default value	Description
[metrics_statsd]	
agent_statsd_host = localhost	(String) Host for the agent ramdisk to use with the statsd backend. This must be accessible from networks the agent is booted on.
agent_statsd_port = 8125	(Port number) Port for the agent ramdisk to use with the statsd backend.
statsd_host = localhost	(String) Host for use with the statsd backend.
statsd_port = 8125	(Port number) Port to use with the statsd backend.

Table 2.22. Description of metrics configuration options specific to statsd backend

Configuration option = Default value	Description
[metrics]	
agent_backend = noop	(String) Backend for the agent ramdisk to use for metrics. Default possible backends are "noop" and "statsd".
agent_global_prefix = None	(String) Prefix all metric names sent by the agent ramdisk with this value. The format of metric names is [global_prefix.][uuid.] [host_name.]prefix.metric_name.
agent_prepend_host = False	(Boolean) Prepend the hostname to all metric names sent by the agent ramdisk. The format of metric names is [global_prefix.][uuid.] [host_name.]prefix.metric_name.
agent_prepend_host_reverse = True	(Boolean) Split the prepended host value by "." and reverse it for metrics sent by the agent ramdisk (to better match the reverse hierarchical form of domain names).
agent_prepend_uuid = False	(Boolean) Prepend the node's ironic uuid to all metric names sent by the agent ramdisk. The format of metric names is [global_prefix.][uuid.] [host_name.]prefix.metric_name.
backend = noop	(String) Backend to use for the metrics system.

Configuration option = Default value	Description
global_prefix = None	(String) Prefix all metric names with this value. By default, there is no global prefix. The format of metric names is [global_prefix.] [host_name.]prefix.metric_name.
prepend_host = False	(Boolean) Prepend the hostname to all metric names. The format of metric names is [global_prefix.] [host_name.]prefix.metric_name.
prepend_host_reverse = True	(Boolean) Split the prepended host value by "." and reverse it (to better match the reverse hierarchical form of domain names).

Table 2.23. Description of neutron configuration options

Configuration option = Default value	Description
[neutron]	
auth_section = None	(Unknown) Config Section from which to load plugin specific options
auth_strategy = keystone	(String) Authentication strategy to use when connecting to neutron. Running neutron in noauth mode (related to but not affected by this setting) is insecure and should only be used for testing.
auth_type = None	(Unknown) Authentication type to load
cafile = None	(String) PEM encoded Certificate Authority to use when verifying HTTPs connections.
certfile = None	(String) PEM encoded client certificate cert file
cleaning_network_uuid = None	(String) Neutron network UUID for the ramdisk to be booted into for cleaning nodes. Required for "neutron" network interface. It is also required if cleaning nodes when using "flat" network interface or "neutron" DHCP provider.
insecure = False	(Boolean) Verify HTTPS connections.
keyfile = None	(String) PEM encoded client certificate key file
port_setup_delay = 0	(Integer) Delay value to wait for Neutron agents to setup sufficient DHCP configuration for port.

Configuration option = Default value	Description
provisioning_network_uuid = None	(String) Neutron network UUID for the ramdisk to be booted into for provisioning nodes. Required for "neutron" network interface.
retries = 3	(Integer) Client retries in the case of a failed request.
timeout = None	(Integer) Timeout value for http requests
url = None	(String) URL for connecting to neutron. Default value translates to 'http://\$my_ip:9696' when auth_strategy is 'noauth', and to discovery from Keystone catalog when auth_strategy is 'keystone'.
url_timeout = 30	(Integer) Timeout value for connecting to neutron in seconds.

Table 2.24. Description of OneView configuration options

Configuration option = Default value	Description
[oneview]	
allow_insecure_connections = False	(Boolean) Option to allow insecure connection with OneView.
enable_periodic_tasks = True	(Boolean) Whether to enable the periodic tasks for OneView driver be aware when OneView hardware resources are taken and released by Ironic or OneView users and proactively manage nodes in clean fail state according to Dynamic Allocation model of hardware resources allocation in OneView.
manager_url = None	(String) URL where OneView is available.
max_polling_attempts = 12	(Integer) Max connection retries to check changes on OneView.
password = None	(String) OneView password to be used.
periodic_check_interval = 300	(Integer) Period (in seconds) for periodic tasks to be executed when enable_periodic_tasks=True.
tls_cacert_file = None	(String) Path to CA certificate.
username = None	(String) OneView username to be used.

Table 2.25. Description of PXE configuration options

Configuration option = Default value	Description
[pxe]	
default_ephemeral_format = ext4	(String) Default file system format for ephemeral partition, if one is created.
image_cache_size = 20480	(Integer) Maximum size (in MiB) of cache for master images, including those in use.
image_cache_ttl = 10080	(Integer) Maximum TTL (in minutes) for old master images in cache.
images_path = /var/lib/ironic/images/	(String) On the ironic-conductor node, directory where images are stored on disk.
instance_master_path = /var/lib/ironic/master_images	(String) On the ironic-conductor node, directory where master instance images are stored on disk. Setting to <None> disables image caching.
ip_version = 4	(String) The IP version that will be used for PXE booting. Defaults to 4. EXPERIMENTAL
ipxe_boot_script = \$pybasedir/drivers/modules/boot.ipxe	(String) On ironic-conductor node, the path to the main iPXE script file.
ipxe_enabled = False	(Boolean) Enable iPXE boot.
ipxe_timeout = 0	(Integer) Timeout value (in seconds) for downloading an image via iPXE. Defaults to 0 (no timeout)
ipxe_use_swift = False	(Boolean) Download deploy images directly from swift using temporary URLs. If set to false (default), images are downloaded to the ironic-conductor node and served over its local HTTP server. Applicable only when 'ipxe_enabled' option is set to true.
pxe_append_params = nofb nomodeset vga=normal	(String) Additional append parameters for baremetal PXE boot.
pxe_bootfile_name = pxelinux.0	(String) Bootfile DHCP parameter.
pxe_config_template = \$pybasedir/drivers/modules/pxe_config.template	(String) On ironic-conductor node, template file for PXE configuration.
tftp_master_path = /tftpboot/master_images	(String) On ironic-conductor node, directory where master TFTP images are stored on disk. Setting to <None> disables image caching.

Configuration option = Default value	Description
tftp_root = /tftpboot	(String) ironic-conductor node's TFTP root path. The ironic-conductor must have read/write access to this path.
tftp_server = \$my_ip	(String) IP address of ironic-conductor node's TFTP server.
uefi_pxe_bootfile_name = bootx64.efi	(String) Bootfile DHCP parameter for UEFI boot mode.
uefi_pxe_config_template = \$pybasedir/drivers/modules/pxe_grub_config.template	(String) On ironic-conductor node, template file for PXE configuration for UEFI boot loader.

Table 2.26. Description of Redis configuration options

Configuration option = Default value	Description
[matchmaker_redis]	
check_timeout = 20000	(Integer) Time in ms to wait before the transaction is killed.
host = 127.0.0.1	(String) DEPRECATED: Host to locate redis. Replaced by [DEFAULT]/transport_url
password =	(String) DEPRECATED: Password for Redis server (optional). Replaced by [DEFAULT]/transport_url
port = 6379	(Port number) DEPRECATED: Use this port to connect to redis host. Replaced by [DEFAULT]/transport_url
sentinel_hosts =	(List) DEPRECATED: List of Redis Sentinel hosts (fault tolerance mode) e.g. [host:port, host1:port ...] Replaced by [DEFAULT]/transport_url
socket_timeout = 10000	(Integer) Timeout in ms on blocking socket operations
wait_timeout = 2000	(Integer) Time in ms to wait between connection attempts.

Table 2.27. Description of SeaMicro configuration options

Configuration option = Default value	Description
[seamicro]	
action_timeout = 10	(Integer) Seconds to wait for power action to be completed
max_retry = 3	(Integer) Maximum retries for SeaMicro operations

Table 2.28. Description of service catalog configuration options

Configuration option = Default value	Description
[service_catalog]	
auth_section = None	(Unknown) Config Section from which to load plugin specific options
auth_type = None	(Unknown) Authentication type to load
cafile = None	(String) PEM encoded Certificate Authority to use when verifying HTTPs connections.
certfile = None	(String) PEM encoded client certificate cert file
insecure = False	(Boolean) Verify HTTPS connections.
keyfile = None	(String) PEM encoded client certificate key file
timeout = None	(Integer) Timeout value for http requests

Table 2.29. Description of SNMP configuration options

Configuration option = Default value	Description
[snmp]	
power_timeout = 10	(Integer) Seconds to wait for power action to be completed
reboot_delay = 0	(Integer) Time (in seconds) to sleep between when rebooting (powering off and on again)

Table 2.30. Description of SSH configuration options

Configuration option = Default value	Description
[ssh]	
get_vm_name_attempts = 3	(Integer) Number of attempts to try to get VM name used by the host that corresponds to a node's MAC address.
get_vm_name_retry_interval = 3	(Integer) Number of seconds to wait between attempts to get VM name used by the host that corresponds to a node's MAC address.
libvirt_uri = qemu:///system	(String) libvirt URI.

Table 2.31. Description of swift configuration options

Configuration option = Default value	Description
[swift]	
auth_section = None	(Unknown) Config Section from which to load plugin specific options
auth_type = None	(Unknown) Authentication type to load
cafile = None	(String) PEM encoded Certificate Authority to use when verifying HTTPs connections.
certfile = None	(String) PEM encoded client certificate cert file
insecure = False	(Boolean) Verify HTTPS connections.
keyfile = None	(String) PEM encoded client certificate key file
swift_max_retries = 2	(Integer) Maximum number of times to retry a Swift request, before failing.
timeout = None	(Integer) Timeout value for http requests

Table 2.32. Description of VirtualBox configuration options

Configuration option = Default value	Description
[virtualbox]	
port = 18083	(Port number) Port on which VirtualBox web service is listening.

2.1.2. New, Updated, and Deprecated Options in Newton for Bare Metal Service

Table 2.33. New options

Option = default value	(Type) Help string
[DEFAULT] default_network_interface = None	(StrOpt) Default network interface to be used for nodes that do not have network_interface field set. A complete list of network interfaces present on your system may be found by enumerating the "ironic.hardware.interfaces.network" endpoint.
[DEFAULT] enabled_network_interfaces = flat, noop	(ListOpt) Specify the list of network interfaces to load during service initialization. Missing network interfaces, or network interfaces which fail to initialize, will prevent the conductor service from starting. The option default is a recommended set of production-oriented network interfaces. A complete list of network interfaces present on your system may be found by enumerating the "ironic.hardware.interfaces.network" endpoint. This value must be the same on all ironic-conductor and ironic-api services, because it is used by ironic-api service to validate a new or updated node's network_interface value.
[DEFAULT] notification_level = None	(StrOpt) Specifies the minimum level for which to send notifications. If not set, no notifications will be sent. The default is for this option to be unset.
[agent] deploy_logs_collect = on_failure	(StrOpt) Whether Ironic should collect the deployment logs on deployment failure (on_failure), always or never.
[agent] deploy_logs_local_path = /var/log/ironic/deploy	(StrOpt) The path to the directory where the logs should be stored, used when the deploy_logs_storage_backend is configured to "local".
[agent] deploy_logs_storage_backend = local	(StrOpt) The name of the storage backend where the logs will be stored.
[agent] deploy_logs_swift_container = ironic_deploy_logs_container	(StrOpt) The name of the Swift container to store the logs, used when the deploy_logs_storage_backend is configured to "swift".

Option = default value	(Type) Help string
[agent] deploy_logs_swift_days_to_expire = 30	(IntOpt) Number of days before a log object is marked as expired in Swift. If None, the logs will be kept forever or until manually deleted. Used when the <code>deploy_logs_storage_backend</code> is configured to "swift".
[api] ramdisk_heartbeat_timeout = 300	(IntOpt) Maximum interval (in seconds) for agent heartbeats.
[api] restrict_lookup = True	(BoolOpt) Whether to restrict the lookup API to only nodes in certain states.
[audit] audit_map_file = /etc/ironic/ironic_api_audit_map.conf	(StrOpt) Path to audit map file for ironic-api service. Used only when API audit is enabled.
[audit] enabled = False	(BoolOpt) Enable auditing of API requests (for ironic-api service).
[audit] ignore_req_list = None	(StrOpt) Comma separated list of Ironic REST API HTTP methods to be ignored during audit. For example: auditing will not be done on any GET or POST requests if this is set to "GET,POST". It is used only when API audit is enabled.
[audit] namespace = openstack	(StrOpt) namespace prefix for generated id
[audit_middleware_notifications] driver = None	(StrOpt) The Driver to handle sending notifications. Possible values are messaging, messagingv2, routing, log, test, noop. If not specified, then value from <code>oslo_messaging_notifications</code> conf section is used.
[audit_middleware_notifications] topics = None	(ListOpt) List of AMQP topics used for OpenStack notifications. If not specified, then value from <code>oslo_messaging_notifications</code> conf section is used.
[audit_middleware_notifications] transport_url = None	(StrOpt) A URL representing messaging driver to use for notification. If not specified, we fall back to the same configuration used for RPC.
[deploy] continue_if_disk_secure_erase_fails = False	(BoolOpt) Defines what to do if an ATA secure erase operation fails during cleaning in the Ironic Python Agent. If False, the cleaning operation will fail and the node will be put in clean failed state. If True, shred will be invoked and cleaning will continue.
[deploy] erase_devices_metadata_priority = None	(IntOpt) Priority to run in-band clean step that erases metadata from devices, via the Ironic Python Agent ramdisk. If unset, will use the priority set in the ramdisk (defaults to 99 for the <code>GenericHardwareManager</code>). If set to 0, will not run during cleaning.

Option = default value	(Type) Help string
[deploy] power_off_after_deploy_failure = True	(BoolOpt) Whether to power off a node after deploy failure. Defaults to True.
[deploy] shred_final_overwrite_with_zeros = True	(BoolOpt) Whether to write zeros to a node's block devices after writing random data. This will write zeros to the device even when <code>deploy.shred_random_overwrite_iterations</code> is 0. This option is only used if a device could not be ATA Secure Erased. Defaults to True.
[deploy] shred_random_overwrite_iterations = 1	(IntOpt) During shred, overwrite all block devices N times with random data. This is only used if a device could not be ATA Secure Erased. Defaults to 1.
[drac] query_raid_config_job_status_interval = 120	(IntOpt) Interval (in seconds) between periodic RAID job status checks to determine whether the asynchronous RAID configuration was successfully finished or not.
[glance] auth_section = None	(Opt) Config Section from which to load plugin specific options
[glance] auth_type = None	(Opt) Authentication type to load
[glance] cafile = None	(StrOpt) PEM encoded Certificate Authority to use when verifying HTTPs connections.
[glance] certfile = None	(StrOpt) PEM encoded client certificate cert file
[glance] insecure = False	(BoolOpt) Verify HTTPS connections.
[glance] keyfile = None	(StrOpt) PEM encoded client certificate key file
[glance] timeout = None	(IntOpt) Timeout value for http requests
[ilo] ca_file = None	(StrOpt) CA certificate file to validate iLO.
[ilo] default_boot_mode = auto	(StrOpt) Default boot mode to be used in provisioning when "boot_mode" capability is not provided in the "properties/capabilities" of the node. The default is "auto" for backward compatibility. When "auto" is specified, default boot mode will be selected based on boot mode settings on the system.
[inspector] auth_section = None	(Opt) Config Section from which to load plugin specific options
[inspector] auth_type = None	(Opt) Authentication type to load

Option = default value	(Type) Help string
[inspector] cafile = None	(StrOpt) PEM encoded Certificate Authority to use when verifying HTTPs connections.
[inspector] certfile = None	(StrOpt) PEM encoded client certificate cert file
[inspector] insecure = False	(BoolOpt) Verify HTTPS connections.
[inspector] keyfile = None	(StrOpt) PEM encoded client certificate key file
[inspector] timeout = None	(IntOpt) Timeout value for http requests
[iscsi] portal_port = 3260	(PortOpt) The port number on which the iSCSI portal listens for incoming connections.
[metrics] agent_backend = noop	(StrOpt) Backend for the agent ramdisk to use for metrics. Default possible backends are "noop" and "statsd".
[metrics] agent_global_prefix = None	(StrOpt) Prefix all metric names sent by the agent ramdisk with this value. The format of metric names is [global_prefix.][uuid.] [host_name.]prefix.metric_name.
[metrics] agent_prepend_host = False	(BoolOpt) Prepend the hostname to all metric names sent by the agent ramdisk. The format of metric names is [global_prefix.][uuid.] [host_name.]prefix.metric_name.
[metrics] agent_prepend_host_reverse = True	(BoolOpt) Split the prepended host value by "." and reverse it for metrics sent by the agent ramdisk (to better match the reverse hierarchical form of domain names).
[metrics] agent_prepend_uuid = False	(BoolOpt) Prepend the node's ironic uuid to all metric names sent by the agent ramdisk. The format of metric names is [global_prefix.][uuid.] [host_name.]prefix.metric_name.
[metrics] backend = noop	(StrOpt) Backend to use for the metrics system.
[metrics] global_prefix = None	(StrOpt) Prefix all metric names with this value. By default, there is no global prefix. The format of metric names is [global_prefix.] [host_name.]prefix.metric_name.
[metrics] prepend_host = False	(BoolOpt) Prepend the hostname to all metric names. The format of metric names is [global_prefix.] [host_name.]prefix.metric_name.

Option = default value	(Type) Help string
[metrics] prepend_host_reverse = True	(BoolOpt) Split the prepended host value by "." and reverse it (to better match the reverse hierarchical form of domain names).
[metrics_statsd] agent_statsd_host = localhost	(StrOpt) Host for the agent ramdisk to use with the statsd backend. This must be accessible from networks the agent is booted on.
[metrics_statsd] agent_statsd_port = 8125	(PortOpt) Port for the agent ramdisk to use with the statsd backend.
[metrics_statsd] statsd_host = localhost	(StrOpt) Host for use with the statsd backend.
[metrics_statsd] statsd_port = 8125	(PortOpt) Port to use with the statsd backend.
[neutron] auth_section = None	(Opt) Config Section from which to load plugin specific options
[neutron] auth_type = None	(Opt) Authentication type to load
[neutron] cafile = None	(StrOpt) PEM encoded Certificate Authority to use when verifying HTTPs connections.
[neutron] certfile = None	(StrOpt) PEM encoded client certificate cert file
[neutron] insecure = False	(BoolOpt) Verify HTTPS connections.
[neutron] keyfile = None	(StrOpt) PEM encoded client certificate key file
[neutron] port_setup_delay = 0	(IntOpt) Delay value to wait for Neutron agents to setup sufficient DHCP configuration for port.
[neutron] provisioning_network_uuid = None	(StrOpt) Neutron network UUID for the ramdisk to be booted into for provisioning nodes. Required for "neutron" network interface.
[neutron] timeout = None	(IntOpt) Timeout value for http requests
[oneview] enable_periodic_tasks = True	(BoolOpt) Whether to enable the periodic tasks for OneView driver be aware when OneView hardware resources are taken and released by Ironic or OneView users and proactively manage nodes in clean fail state according to Dynamic Allocation model of hardware resources allocation in OneView.
[oneview] periodic_check_interval = 300	(IntOpt) Period (in seconds) for periodic tasks to be executed when enable_periodic_tasks=True.

Option = default value	(Type) Help string
[pxe] ipxe_use_swift = False	(BoolOpt) Download deploy images directly from swift using temporary URLs. If set to false (default), images are downloaded to the ironic-conductor node and served over its local HTTP server. Applicable only when 'ipxe_enabled' option is set to true.
[service_catalog] auth_section = None	(Opt) Config Section from which to load plugin specific options
[service_catalog] auth_type = None	(Opt) Authentication type to load
[service_catalog] cafile = None	(StrOpt) PEM encoded Certificate Authority to use when verifying HTTPs connections.
[service_catalog] certfile = None	(StrOpt) PEM encoded client certificate cert file
[service_catalog] insecure = False	(BoolOpt) Verify HTTPS connections.
[service_catalog] keyfile = None	(StrOpt) PEM encoded client certificate key file
[service_catalog] timeout = None	(IntOpt) Timeout value for http requests
[swift] auth_section = None	(Opt) Config Section from which to load plugin specific options
[swift] auth_type = None	(Opt) Authentication type to load
[swift] cafile = None	(StrOpt) PEM encoded Certificate Authority to use when verifying HTTPs connections.
[swift] certfile = None	(StrOpt) PEM encoded client certificate cert file
[swift] insecure = False	(BoolOpt) Verify HTTPS connections.
[swift] keyfile = None	(StrOpt) PEM encoded client certificate key file
[swift] timeout = None	(IntOpt) Timeout value for http requests

Table 2.34. New default values

Option	Previous default value	New default value
[neutron] url	http://\$my_ip:9696	None
[pxe] uefi_pxe_bootfile_name	elilo.efi	bootx64.efi

Option	Previous default value	New default value
[pxe] uefi_pxe_config_template	\$pybasedir/drivers/modules/elilo_efi_pxe_config.template	\$pybasedir/drivers/modules/pxe_grub_config.template

Table 2.35. Deprecated options

Deprecated option	New Option
[DEFAULT] use_syslog	None
[agent] heartbeat_timeout	[api] ramdisk_heartbeat_timeout
[deploy] erase_devices_iterations	[deploy] shred_random_overwrite_iterations
[keystone_auth token] cafile	[glance] cafile
[keystone_auth token] cafile	[neutron] cafile
[keystone_auth token] cafile	[service_catalog] cafile
[keystone_auth token] cafile	[swift] cafile
[keystone_auth token] cafile	[inspector] cafile
[keystone_auth token] certfile	[service_catalog] certfile
[keystone_auth token] certfile	[neutron] certfile
[keystone_auth token] certfile	[glance] certfile
[keystone_auth token] certfile	[inspector] certfile
[keystone_auth token] certfile	[swift] certfile
[keystone_auth token] insecure	[glance] insecure
[keystone_auth token] insecure	[inspector] insecure
[keystone_auth token] insecure	[swift] insecure
[keystone_auth token] insecure	[service_catalog] insecure
[keystone_auth token] insecure	[neutron] insecure
[keystone_auth token] keyfile	[inspector] keyfile

Deprecated option	New Option
[keystone_authtoken] keyfile	[swift] keyfile
[keystone_authtoken] keyfile	[neutron] keyfile
[keystone_authtoken] keyfile	[glance] keyfile
[keystone_authtoken] keyfile	[service_catalog] keyfile

CHAPTER 3. BLOCK STORAGE

The OpenStack Block Storage service provides persistent storage for Compute instances, working with many different storage drivers that you can configure.

3.1. BLOCK STORAGE CONFIGURATION OPTIONS

3.1.1. Description of Configuration Options

The following tables provide a comprehensive list of the Block Storage configuration options.

Table 3.1. Description of API configuration options

Configuration option = Default value	Description
[DEFAULT]	
api_rate_limit = True	(Boolean) Enables or disables rate limit of the API.
az_cache_duration = 3600	(Integer) Cache volume availability zones in memory for the provided duration in seconds
backend_host = None	(String) Backend override of host value.
default_timeout = 31536000	(Integer) 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.
enable_v1_api = True	(Boolean) DEPRECATED: Deploy v1 of the Cinder API.
enable_v2_api = True	(Boolean) DEPRECATED: Deploy v2 of the Cinder API.
enable_v3_api = True	(Boolean) Deploy v3 of the Cinder API.
extra_capabilities = {}	(String) User defined capabilities, a JSON formatted string specifying key/value pairs. The key/value pairs can be used by the CapabilitiesFilter to select between backends when requests specify volume types. For example, specifying a service level or the geographical location of a backend, then creating a volume type to allow the user to select by these different properties.
ignore_pool_full_threshold = False	(Boolean) Force LUN creation even if the full threshold of pool is reached. By default, the value is False.

Configuration option = Default value	Description
management_ips =	(String) List of Management IP addresses (separated by commas)
message_ttl = 2592000	(Integer) message minimum life in seconds.
osapi_max_limit = 1000	(Integer) The maximum number of items that a collection resource returns in a single response
osapi_max_request_body_size = 114688	(Integer) Max size for body of a request
osapi_volume_base_URL = None	(String) Base URL that will be presented to users in links to the OpenStack Volume API
osapi_volume_ext_list =	(List) Specify list of extensions to load when using <code>osapi_volume_extension</code> option with <code>cinder.api.contrib.select_extensions</code>
osapi_volume_extension = ['cinder.api.contrib.standard_extensions']	(Multi-valued) osapi volume extension to load
osapi_volume_listen = 0.0.0.0	(String) IP address on which OpenStack Volume API listens
osapi_volume_listen_port = 8776	(Port number) Port on which OpenStack Volume API listens
osapi_volume_use_ssl = False	(Boolean) Wraps the socket in a SSL context if True is set. A certificate file and key file must be specified.
osapi_volume_workers = None	(Integer) Number of workers for OpenStack Volume API service. The default is equal to the number of CPUs available.
per_volume_size_limit = -1	(Integer) Max size allowed per volume, in gigabytes
public_endpoint = None	(String) Public url to use for versions endpoint. The default is None, which will use the request's <code>host_url</code> attribute to populate the URL base. If Cinder is operating behind a proxy, you will want to change this to represent the proxy's URL.
query_volume_filters = name, status, metadata, availability_zone, bootable, group_id	(List) Volume filter options which non-admin user could use to query volumes. Default values are: <code>['name', 'status', 'metadata', 'availability_zone', 'bootable', 'group_id']</code>
transfer_api_class = cinder.transfer.api.API	(String) The full class name of the volume transfer API class
volume_api_class = cinder.volume.api.API	(String) The full class name of the volume API class to use

Configuration option = Default value	Description
volume_name_prefix = openstack-	(String) Prefix before volume name to differentiate DISCO volume created through openstack and the other ones
volume_name_template = volume-%s	(String) Template string to be used to generate volume names
volume_number_multiplier = -1.0	(Floating point) Multiplier used for weighing volume number. Negative numbers mean to spread vs stack.
volume_transfer_key_length = 16	(Integer) The number of characters in the autogenerated auth key.
volume_transfer_salt_length = 8	(Integer) The number of characters in the salt.
[oslo_middleware]	
enable_proxy_headers_parsing = False	(Boolean) Whether the application is behind a proxy or not. This determines if the middleware should parse the headers or not.
max_request_body_size = 114688	(Integer) The maximum body size for each request, in bytes.
secure_proxy_ssl_header = X-Forwarded-Proto	(String) DEPRECATED: The HTTP Header that will be used to determine what the original request protocol scheme was, even if it was hidden by a SSL termination proxy.
[oslo_versionedobjects]	
fatal_exception_format_errors = False	(Boolean) Make exception message format errors fatal

Table 3.2. Description of authorization configuration options

Configuration option = Default value	Description
[DEFAULT]	
auth_strategy = keystone	(String) The strategy to use for auth. Supports noauth or keystone.

Table 3.3. Description of Ceph backup driver configuration options

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

Table 3.4. Description of GCS backup driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
backup_gcs_block_size = 32768	(Integer) The size in bytes that changes are tracked for incremental backups. backup_gcs_object_size has to be multiple of backup_gcs_block_size.
backup_gcs_bucket = None	(String) The GCS bucket to use.
backup_gcs_bucket_location = US	(String) Location of GCS bucket.
backup_gcs_credential_file = None	(String) Absolute path of GCS service account credential file.
backup_gcs_enable_progress_timer = True	(Boolean) Enable or Disable the timer to send the periodic progress notifications to Ceilometer when backing up the volume to the GCS backend storage. The default value is True to enable the timer.
backup_gcs_num_retries = 3	(Integer) Number of times to retry.

Configuration option = Default value	Description
backup_gcs_object_size = 52428800	(Integer) The size in bytes of GCS backup objects.
backup_gcs_project_id = None	(String) Owner project id for GCS bucket.
backup_gcs_proxy_url = None	(URI) URL for http proxy access.
backup_gcs_reader_chunk_size = 2097152	(Integer) GCS object will be downloaded in chunks of bytes.
backup_gcs_retry_error_codes = 429	(List) List of GCS error codes.
backup_gcs_storage_class = NEARLINE	(String) Storage class of GCS bucket.
backup_gcs_user_agent = gcscinder	(String) Http user-agent string for gcs api.
backup_gcs_writer_chunk_size = 2097152	(Integer) GCS object will be uploaded in chunks of bytes. Pass in a value of -1 if the file is to be uploaded as a single chunk.

Table 3.5. Description of NFS backup driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
backup_container = None	(String) Custom directory to use for backups.
backup_enable_progress_timer = True	(Boolean) Enable or Disable the timer to send the periodic progress notifications to Ceilometer when backing up the volume to the backend storage. The default value is True to enable the timer.
backup_file_size = 1999994880	(Integer) The maximum size in bytes of the files used to hold backups. If the volume being backed up exceeds this size, then it will be backed up into multiple files. backup_file_size must be a multiple of backup_sha_block_size_bytes.
backup_mount_options = None	(String) Mount options passed to the NFS client. See NFS man page for details.
backup_mount_point_base = \$state_path/backup_mount	(String) Base dir containing mount point for NFS share.
backup_sha_block_size_bytes = 32768	(Integer) The size in bytes that changes are tracked for incremental backups. backup_file_size has to be multiple of backup_sha_block_size_bytes.

Configuration option = Default value	Description
backup_share = None	(String) NFS share in hostname:path, ipv4addr:path, or "[ipv6addr]:path" format.

Table 3.6. Description of POSIX backup driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
backup_container = None	(String) Custom directory to use for backups.
backup_enable_progress_timer = True	(Boolean) Enable or Disable the timer to send the periodic progress notifications to Ceilometer when backing up the volume to the backend storage. The default value is True to enable the timer.
backup_file_size = 1999994880	(Integer) The maximum size in bytes of the files used to hold backups. If the volume being backed up exceeds this size, then it will be backed up into multiple files. backup_file_size must be a multiple of backup_sha_block_size_bytes.
backup_posix_path = \$state_path/backup	(String) Path specifying where to store backups.
backup_sha_block_size_bytes = 32768	(Integer) The size in bytes that changes are tracked for incremental backups. backup_file_size has to be multiple of backup_sha_block_size_bytes.

Table 3.7. Description of Swift backup driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
backup_swift_auth = per_user	(String) Swift authentication mechanism
backup_swift_auth_version = 1	(String) Swift authentication version. Specify "1" for auth 1.0, or "2" for auth 2.0 or "3" for auth 3.0
backup_swift_block_size = 32768	(Integer) The size in bytes that changes are tracked for incremental backups. backup_swift_object_size has to be multiple of backup_swift_block_size.

Configuration option = Default value	Description
backup_swift_ca_cert_file = None	(String) Location of the CA certificate file to use for swift client requests.
backup_swift_container = volumebackups	(String) The default Swift container to use
backup_swift_enable_progress_timer = True	(Boolean) 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.
backup_swift_key = None	(String) Swift key for authentication
backup_swift_object_size = 52428800	(Integer) The size in bytes of Swift backup objects
backup_swift_project = None	(String) Swift project/account name. Required when connecting to an auth 3.0 system
backup_swift_project_domain = None	(String) Swift project domain name. Required when connecting to an auth 3.0 system
backup_swift_retry_attempts = 3	(Integer) The number of retries to make for Swift operations
backup_swift_retry_backoff = 2	(Integer) The backoff time in seconds between Swift retries
backup_swift_tenant = None	(String) Swift tenant/account name. Required when connecting to an auth 2.0 system
backup_swift_url = None	(String) The URL of the Swift endpoint
backup_swift_user = None	(String) Swift user name
backup_swift_user_domain = None	(String) Swift user domain name. Required when connecting to an auth 3.0 system
keystone_catalog_info = identity:Identity Service:publicURL	(String) Info to match when looking for keystone in the service catalog. Format is: separated values of the form: <service_type>:<service_name>:<endpoint_type> - Only used if backup_swift_auth_url is unset
swift_catalog_info = object-store:swift:publicURL	(String) Info to match when looking for swift in the service catalog. Format is: separated values of the form: <service_type>:<service_name>:<endpoint_type> - Only used if backup_swift_url is unset

Table 3.8. Description of IBM Tivoli Storage Manager backup driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
backup_tsm_compression = True	(Boolean) Enable or Disable compression for backups
backup_tsm_password = password	(String) TSM password for the running username
backup_tsm_volume_prefix = backup	(String) Volume prefix for the backup id when backing up to TSM

Table 3.9. Description of backups configuration options

Configuration option = Default value	Description
[DEFAULT]	
backup_api_class = cinder.backup.api.API	(String) The full class name of the volume backup API class
backup_compression_algorithm = zlib	(String) Compression algorithm (None to disable)
backup_driver = cinder.backup.drivers.swift	(String) Driver to use for backups.
backup_manager = cinder.backup.manager.BackupManager	(String) Full class name for the Manager for volume backup
backup_metadata_version = 2	(Integer) Backup metadata version to be used when backing up volume metadata. If this number is bumped, make sure the service doing the restore supports the new version.
backup_name_template = backup-%s	(String) Template string to be used to generate backup names
backup_object_number_per_notification = 10	(Integer) The number of chunks or objects, for which one Ceilometer notification will be sent
backup_service_inithost_offload = True	(Boolean) Offload pending backup delete during backup service startup. If false, the backup service will remain down until all pending backups are deleted.
backup_timer_interval = 120	(Integer) Interval, in seconds, between two progress notifications reporting the backup status

Configuration option = Default value	Description
backup_use_same_host = False	(Boolean) Backup services use same backend.
backup_use_temp_snapshot = False	(Boolean) If this is set to True, the backup_use_temp_snapshot path will be used during the backup. Otherwise, it will use backup_use_temp_volume path.
snapshot_check_timeout = 3600	(Integer) How long we check whether a snapshot is finished before we give up
snapshot_name_template = snapshot-%s	(String) Template string to be used to generate snapshot names
snapshot_same_host = True	(Boolean) Create volume from snapshot at the host where snapshot resides

Table 3.10. Description of block device configuration options

Configuration option = Default value	Description
[DEFAULT]	
available_devices =	(List) List of all available devices

Table 3.11. Description of CloudByte volume driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
cb_account_name = None	(String) CloudByte storage specific account name. This maps to a project name in OpenStack.
cb_add_qosgroup = {'latency': '15', 'iops': '10', 'graceallowed': 'false', 'iopscontrol': 'true', 'memlimit': '0', 'throughput': '0', 'tpcontrol': 'false', 'networkspeed': '0'}	(Dict) These values will be used for CloudByte storage's addQos API call.
cb_apikey = None	(String) Driver will use this API key to authenticate against the CloudByte storage's management interface.

Configuration option = Default value	Description
cb_auth_group = None	(String) This corresponds to the discovery authentication group in CloudByte storage. Chap users are added to this group. Driver uses the first user found for this group. Default value is None.
cb_confirm_volume_create_retries = 3	(Integer) Will confirm a successful volume creation in CloudByte storage by making this many number of attempts.
cb_confirm_volume_create_retry_interval = 5	(Integer) A retry value in seconds. Will be used by the driver to check if volume creation was successful in CloudByte storage.
cb_confirm_volume_delete_retries = 3	(Integer) Will confirm a successful volume deletion in CloudByte storage by making this many number of attempts.
cb_confirm_volume_delete_retry_interval = 5	(Integer) A retry value in seconds. Will be used by the driver to check if volume deletion was successful in CloudByte storage.
cb_create_volume = {'compression': 'off', 'deduplication': 'off', 'blocklength': '512B', 'sync': 'always', 'protocoltype': 'ISCSI', 'recordsize': '16k'}	(Dict) These values will be used for CloudByte storage's createVolume API call.
cb_tsm_name = None	(String) This corresponds to the name of Tenant Storage Machine (TSM) in CloudByte storage. A volume will be created in this TSM.
cb_update_file_system = compression, sync, noofcopies, readonly	(List) These values will be used for CloudByte storage's updateFileSystem API call.
cb_update_qos_group = iops, latency, graceallowed	(List) These values will be used for CloudByte storage's updateQosGroup API call.

Table 3.12. Description of common configuration options

Configuration option = Default value	Description
[DEFAULT]	
allow_availability_zone_fallback = False	(Boolean) If the requested Cinder availability zone is unavailable, fall back to the value of <code>default_availability_zone</code> , then <code>storage_availability_zone</code> , instead of failing.

Configuration option = Default value	Description
chap = disabled	(String) CHAP authentication mode, effective only for iscsi (disabled enabled)
chap_password =	(String) Password for specified CHAP account name.
chap_username =	(String) CHAP user name.
chiscsi_conf = /etc/chelsio-iscsi/chiscsi.conf	(String) Chiscsi (CXT) global defaults configuration file
cinder_internal_tenant_project_id = None	(String) ID of the project which will be used as the Cinder internal tenant.
cinder_internal_tenant_user_id = None	(String) ID of the user to be used in volume operations as the Cinder internal tenant.
cluster = None	(String) Name of this cluster. Used to group volume hosts that share the same backend configurations to work in HA Active-Active mode. Active-Active is not yet supported.
compute_api_class = cinder.compute.nova.API	(String) The full class name of the compute API class to use
connection_type = iscsi	(String) Connection type to the IBM Storage Array
consistencygroup_api_class = cinder.consistencygroup.api.API	(String) The full class name of the consistencygroup API class
default_availability_zone = None	(String) Default availability zone for new volumes. If not set, the storage_availability_zone option value is used as the default for new volumes.
default_group_type = None	(String) Default group type to use
default_volume_type = None	(String) Default volume type to use
driver_client_cert = None	(String) The path to the client certificate for verification, if the driver supports it.
driver_client_cert_key = None	(String) The path to the client certificate key for verification, if the driver supports it.
driver_data_namespace = None	(String) Namespace for driver private data values to be saved in.
driver_ssl_cert_path = None	(String) Can be used to specify a non default path to a CA_BUNDLE file or directory with certificates of trusted CAs, which will be used to validate the backend

Configuration option = Default value	Description
driver_ssl_cert_verify = False	(Boolean) If set to True the http client will validate the SSL certificate of the backend endpoint.
enable_force_upload = False	(Boolean) Enables the Force option on <code>upload_to_image</code> . This enables running <code>upload_volume</code> on in-use volumes for backends that support it.
enable_new_services = True	(Boolean) Services to be added to the available pool on create
enable_unsupported_driver = False	(Boolean) Set this to True when you want to allow an unsupported driver to start. Drivers that haven't maintained a working CI system and testing are marked as unsupported until CI is working again. This also marks a driver as deprecated and may be removed in the next release.
end_time = None	(String) If this option is specified then the end time specified is used instead of the end time of the last completed audit period.
enforce_multipath_for_image_xfer = False	(Boolean) If this is set to True, attachment of volumes for image transfer will be aborted when <code>multipathd</code> is not running. Otherwise, it will fallback to single path.
executor_thread_pool_size = 64	(Integer) Size of executor thread pool.
fatal_exception_format_errors = False	(Boolean) Make exception message format errors fatal.
group_api_class = cinder.group.api.API	(String) The full class name of the group API class
host = localhost	(String) Name of this node. This can be an opaque identifier. It is not necessarily a host name, FQDN, or IP address.
iet_conf = /etc/iet/ietd.conf	(String) IET configuration file
iscsi_secondary_ip_addresses =	(List) The list of secondary IP addresses of the iSCSI daemon
max_over_subscription_ratio = 20.0	(Floating point) 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. The ratio has to be a minimum of 1.0.

Configuration option = Default value	Description
monkey_patch = False	(Boolean) Enable monkey patching
monkey_patch_modules =	(List) List of modules/decorators to monkey patch
my_ip = 10.0.0.1	(String) IP address of this host
no_snapshot_gb_quota = False	(Boolean) Whether snapshots count against gigabyte quota
num_shell_tries = 3	(Integer) Number of times to attempt to run flakey shell commands
os_privileged_user_auth_url = None	(String) Auth URL associated with the OpenStack privileged account.
os_privileged_user_name = None	(String) OpenStack privileged account username. Used for requests to other services (such as Nova) that require an account with special rights.
os_privileged_user_password = None	(String) Password associated with the OpenStack privileged account.
os_privileged_user_tenant = None	(String) Tenant name associated with the OpenStack privileged account.
periodic_fuzzy_delay = 60	(Integer) Range, in seconds, to randomly delay when starting the periodic task scheduler to reduce stampeding. (Disable by setting to 0)
periodic_interval = 60	(Integer) Interval, in seconds, between running periodic tasks
replication_api_class = cinder.replication.api.API	(String) The full class name of the volume replication API class
replication_device = None	(Unknown) Multi opt of dictionaries to represent a replication target device. This option may be specified multiple times in a single config section to specify multiple replication target devices. Each entry takes the standard dict config form: replication_device = target_device_id: <required>,key1:value1,key2:value2...
report_discard_supported = False	(Boolean) Report to clients of Cinder that the backend supports discard (aka. trim/unmap). This will not actually change the behavior of the backend or the client directly, it will only notify that it can be used.
report_interval = 10	(Integer) Interval, in seconds, between nodes reporting state to datastore

Configuration option = Default value	Description
reserved_percentage = 0	(Integer) The percentage of backend capacity is reserved
rootwrap_config = /etc/cinder/rootwrap.conf	(String) Path to the rootwrap configuration file to use for running commands as root
send_actions = False	(Boolean) Send the volume and snapshot create and delete notifications generated in the specified period.
service_down_time = 60	(Integer) Maximum time since last check-in for a service to be considered up
ssh_hosts_key_file = \$state_path/ssh_known_hosts	(String) File containing SSH host keys for the systems with which Cinder needs to communicate. OPTIONAL: Default=\$state_path/ssh_known_hosts
start_time = None	(String) If this option is specified then the start time specified is used instead of the start time of the last completed audit period.
state_path = /var/lib/cinder	(String) Top-level directory for maintaining cinder's state
storage_availability_zone = nova	(String) Availability zone of this node
storage_protocol = iscsi	(String) Protocol for transferring data between host and storage back-end.
strict_ssh_host_key_policy = False	(Boolean) Option to enable strict host key checking. When set to "True" Cinder will only connect to systems with a host key present in the configured "ssh_hosts_key_file". When set to "False" the host key will be saved upon first connection and used for subsequent connections. Default=False
suppress_requests_ssl_warnings = False	(Boolean) Suppress requests library SSL certificate warnings.
tcp_keepalive = True	(Boolean) Sets the value of TCP_KEEPALIVE (True/False) for each server socket.
tcp_keepalive_count = None	(Integer) Sets the value of TCP_KEEPCNT for each server socket. Not supported on OS X.
tcp_keepalive_interval = None	(Integer) Sets the value of TCP_KEEPINTVL in seconds for each server socket. Not supported on OS X.
until_refresh = 0	(Integer) Count of reservations until usage is refreshed

Configuration option = Default value	Description
use_chap_auth = False	(Boolean) Option to enable/disable CHAP authentication for targets.
use_forwarded_for = False	(Boolean) Treat X-Forwarded-For as the canonical remote address. Only enable this if you have a sanitizing proxy.
[key_manager]	
api_class = castellan.key_manager.barbican_key_manager.BarbicanKeyManager	(String) The full class name of the key manager API class
fixed_key = None	(String) Fixed key returned by key manager, specified in hex

Table 3.13. Description of Compute configuration options

Configuration option = Default value	Description
[DEFAULT]	
nova_api_insecure = False	(Boolean) Allow to perform insecure SSL requests to nova
nova_ca_certificates_file = None	(String) Location of ca certificates file to use for nova client requests.
nova_catalog_admin_info = compute:Compute Service:adminURL	(String) Same as nova_catalog_info, but for admin endpoint.
nova_catalog_info = compute:Compute Service:publicURL	(String) Match this value when searching for nova in the service catalog. Format is: separated values of the form: <service_type>:<service_name>:<endpoint_type>
nova_endpoint_admin_template = None	(String) Same as nova_endpoint_template, but for admin endpoint.
nova_endpoint_template = None	(String) Override service catalog lookup with template for nova endpoint e.g. http://localhost:8774/v2/%(project_id)s
os_region_name = None	(String) Region name of this node

Table 3.14. Description of Coordination configuration options

Configuration option = Default value	Description
[coordination]	
backend_url = file://\$state_path	(String) The backend URL to use for distributed coordination.
heartbeat = 1.0	(Floating point) Number of seconds between heartbeats for distributed coordination.
initial_reconnect_backoff = 0.1	(Floating point) Initial number of seconds to wait after failed reconnection.
max_reconnect_backoff = 60.0	(Floating point) Maximum number of seconds between sequential reconnection retries.

Table 3.15. Description of logging configuration options

Configuration option = Default value	Description
[DEFAULT]	
trace_flags = None	(List) List of options that control which trace info is written to the DEBUG log level to assist developers. Valid values are method and api.

Table 3.16. Description of Dell Storage Center volume driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
dell_sc_api_port = 3033	(Port number) Dell API port
dell_sc_server_folder = openstack	(String) Name of the server folder to use on the Storage Center
dell_sc_ssn = 64702	(Integer) Storage Center System Serial Number
dell_sc_verify_cert = False	(Boolean) Enable HTTPS SC certificate verification

Configuration option = Default value	Description
dell_sc_volume_folder = openstack	(String) Name of the volume folder to use on the Storage Center
dell_server_os = Red Hat Linux 6.x	(String) Server OS type to use when creating a new server on the Storage Center.
excluded_domain_ip = None	(Unknown) Domain IP to be excluded from iSCSI returns.
secondary_san_ip =	(String) IP address of secondary DSM controller
secondary_san_login = Admin	(String) Secondary DSM user name
secondary_san_password =	(String) Secondary DSM user password name
secondary_sc_api_port = 3033	(Port number) Secondary Dell API port

Table 3.17. Description of EMC SIO volume driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
sio_max_over_subscription_ratio = 10.0	(Floating point) max_over_subscription_ratio setting for the ScaleIO driver. This replaces the general max_over_subscription_ratio which has no effect in this driver. Maximum value allowed for ScaleIO is 10.0.
sio_protection_domain_id = None	(String) Protection Domain ID.
sio_protection_domain_name = None	(String) Protection Domain name.
sio_rest_server_port = 443	(String) REST server port.
sio_round_volume_capacity = True	(Boolean) Round up volume capacity.
sio_server_certificate_path = None	(String) Server certificate path.
sio_storage_pool_id = None	(String) Storage Pool ID.
sio_storage_pool_name = None	(String) Storage Pool name.
sio_storage_pools = None	(String) Storage Pools.

Configuration option = Default value	Description
sio_unmap_volume_before_deletion = False	(Boolean) Unmap volume before deletion.
sio_verify_server_certificate = False	(Boolean) Verify server certificate.

Table 3.18. Description of EMC configuration options

Configuration option = Default value	Description
[DEFAULT]	
check_max_pool_luns_threshold = False	(Boolean) Report free_capacity_gb as 0 when the limit to maximum number of pool LUNs is reached. By default, the value is False.
cinder_emc_config_file = /etc/cinder/cinder_emc_config.xml	(String) use this file for cinder emc plugin config data
destroy_empty_storage_group = False	(Boolean) To destroy storage group when the last LUN is removed from it. By default, the value is False.
force_delete_lun_in_storagegroup = False	(Boolean) Delete a LUN even if it is in Storage Groups. By default, the value is False.
initiator_auto_deregistration = False	(Boolean) Automatically deregister initiators after the related storage group is destroyed. By default, the value is False.
initiator_auto_registration = False	(Boolean) Automatically register initiators. By default, the value is False.
io_port_list = None	(List) Comma separated iSCSI or FC ports to be used in Nova or Cinder.
iscsi_initiators = None	(String) Mapping between hostname and its iSCSI initiator IP addresses.
max_luns_per_storage_group = 255	(Integer) Default max number of LUNs in a storage group. By default, the value is 255.
naviseccli_path = None	(String) Naviseccli Path.
storage_vnx_authentication_type = global	(String) VNX authentication scope type. By default, the value is global.
storage_vnx_pool_names = None	(List) Comma-separated list of storage pool names to be used.

Configuration option = Default value	Description
storage_vnx_security_file_dir = None	(String) Directory path that contains the VNX security file. Make sure the security file is generated first.

Table 3.19. Description of EMC XtremIO volume driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
xtremio_array_busy_retry_count = 5	(Integer) Number of retries in case array is busy
xtremio_array_busy_retry_interval = 5	(Integer) Interval between retries in case array is busy
xtremio_cluster_name =	(String) XMS cluster id in multi-cluster environment
xtremio_volumes_per_glance_cache = 100	(Integer) Number of volumes created from each cached glance image

Table 3.20. Description of Dell EqualLogic volume driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
eqlx_chap_login = admin	(String) Existing CHAP account name. Note that this option is deprecated in favour of "chap_username" as specified in cinder/volume/driver.py and will be removed in next release.
eqlx_chap_password = password	(String) Password for specified CHAP account name. Note that this option is deprecated in favour of "chap_password" as specified in cinder/volume/driver.py and will be removed in the next release
eqlx_cli_max_retries = 5	(Integer) Maximum retry count for reconnection. Default is 5.

Configuration option = Default value	Description
eqlx_cli_timeout = 30	(Integer) Timeout for the Group Manager cli command execution. Default is 30. Note that this option is deprecated in favour of "ssh_conn_timeout" as specified in cinder/volume/drivers/san/san.py and will be removed in M release.
eqlx_group_name = group-0	(String) Group name to use for creating volumes. Defaults to "group-0".
eqlx_pool = default	(String) Pool in which volumes will be created. Defaults to "default".
eqlx_use_chap = False	(Boolean) 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.

Table 3.21. Description of Eternus volume driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
cinder_eternus_config_file = /etc/cinder/cinder_fujitsu_eternus_dx.xml	(String) config file for cinder eternus_dx volume driver

Table 3.22. Description of IBM FlashSystem volume driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
flashsystem_connection_protocol = FC	(String) Connection protocol should be FC. (Default is FC.)
flashsystem_iscsi_portid = 0	(Integer) Default iSCSI Port ID of FlashSystem. (Default port is 0.)
flashsystem_multihostmap_enabled = True	(Boolean) Allows vdisk to multi host mapping. (Default is True)

Configuration option = Default value	Description
flashsystem_multipath_enabled = False	(Boolean) DEPRECATED: This option no longer has any affect. It is deprecated and will be removed in the next release.

Table 3.23. Description of HDS HNAS iSCSI and NFS driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
hds_hnas_iscsi_config_file = /opt/hds/hnas/cinder_iscsi_conf.xml	(String) DEPRECATED: Legacy configuration file for HNAS iSCSI Cinder plugin. This is not needed if you fill all configuration on cinder.conf
hds_hnas_nfs_config_file = /opt/hds/hnas/cinder_nfs_conf.xml	(String) DEPRECATED: Legacy configuration file for HNAS NFS Cinder plugin. This is not needed if you fill all configuration on cinder.conf

Table 3.24. Description of Hitachi storage volume driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
hitachi_add_chap_user = False	(Boolean) Add CHAP user
hitachi_async_copy_check_interval = 10	(Integer) Interval to check copy asynchronously
hitachi_auth_method = None	(String) iSCSI authentication method
hitachi_auth_password = HBSD-CHAP-password	(String) iSCSI authentication password
hitachi_auth_user = HBSD-CHAP-user	(String) iSCSI authentication username
hitachi_copy_check_interval = 3	(Integer) Interval to check copy
hitachi_copy_speed = 3	(Integer) Copy speed of storage system
hitachi_default_copy_method = FULL	(String) Default copy method of storage system

Configuration option = Default value	Description
hitachi_group_range = None	(String) Range of group number
hitachi_group_request = False	(Boolean) Request for creating HostGroup or iSCSI Target
hitachi_horcm_add_conf = True	(Boolean) Add to HORCM configuration
hitachi_horcm_numbers = 200,201	(String) Instance numbers for HORCM
hitachi_horcm_password = None	(String) Password of storage system for HORCM
hitachi_horcm_resource_lock_timeout = 600	(Integer) Timeout until a resource lock is released, in seconds. The value must be between 0 and 7200.
hitachi_horcm_user = None	(String) Username of storage system for HORCM
hitachi_ldev_range = None	(String) Range of logical device of storage system
hitachi_pool_id = None	(Integer) Pool ID of storage system
hitachi_serial_number = None	(String) Serial number of storage system
hitachi_target_ports = None	(String) Control port names for HostGroup or iSCSI Target
hitachi_thin_pool_id = None	(Integer) Thin pool ID of storage system
hitachi_unit_name = None	(String) Name of an array unit
hitachi_zoning_request = False	(Boolean) Request for FC Zone creating HostGroup
hnas_chap_enabled = True	(Boolean) Whether the chap authentication is enabled in the iSCSI target or not.
hnas_cluster_admin_ip0 = None	(String) The IP of the HNAS cluster admin. Required only for HNAS multi-cluster setups.
hnas_mgmt_ip0 = None	(IP) Management IP address of HNAS. This can be any IP in the admin address on HNAS or the SMU IP.
hnas_password = None	(String) HNAS password.
hnas_ssc_cmd = ssc	(String) Command to communicate to HNAS.
hnas_ssh_port = 22	(Port number) Port to be used for SSH authentication.
hnas_ssh_private_key = None	(String) Path to the SSH private key used to authenticate in HNAS SMU.

Configuration option = Default value	Description
hnas_svc0_hdp = None	(String) Service 0 HDP
hnas_svc0_iscsi_ip = None	(IP) Service 0 iSCSI IP
hnas_svc0_volume_type = None	(String) Service 0 volume type
hnas_svc1_hdp = None	(String) Service 1 HDP
hnas_svc1_iscsi_ip = None	(IP) Service 1 iSCSI IP
hnas_svc1_volume_type = None	(String) Service 1 volume type
hnas_svc2_hdp = None	(String) Service 2 HDP
hnas_svc2_iscsi_ip = None	(IP) Service 2 iSCSI IP
hnas_svc2_volume_type = None	(String) Service 2 volume type
hnas_svc3_hdp = None	(String) Service 3 HDP
hnas_svc3_iscsi_ip = None	(IP) Service 3 iSCSI IP
hnas_svc3_volume_type = None	(String) Service 3 volume type
hnas_username = None	(String) HNAS username.

Table 3.25. Description of HPE 3PAR Fibre Channel and iSCSI drivers configuration options

Configuration option = Default value	Description
[DEFAULT]	
hpe3par_api_url =	(String) 3PAR WSAPI Server Url like <a href="https://<3par ip>:8080/api/v1">https://<3par ip>:8080/api/v1
hpe3par_cpg = OpenStack	(List) List of the CPG(s) to use for volume creation
hpe3par_cpg_snap =	(String) The CPG to use for Snapshots for volumes. If empty the userCPG will be used.
hpe3par_debug = False	(Boolean) Enable HTTP debugging to 3PAR
hpe3par_iscsi_chap_enabled = False	(Boolean) Enable CHAP authentication for iSCSI connections.

Configuration option = Default value	Description
hpe3par_iscsi_ips =	(List) List of target iSCSI addresses to use.
hpe3par_password =	(String) 3PAR password for the user specified in hpe3par_username
hpe3par_snapshot_expiration =	(String) The time in hours when a snapshot expires and is deleted. This must be larger than expiration
hpe3par_snapshot_retention =	(String) The time in hours to retain a snapshot. You can't delete it before this expires.
hpe3par_username =	(String) 3PAR username with the 'edit' role

Table 3.26. Description of Huawei storage driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
cinder_huawei_conf_file = /etc/cinder/cinder_huawei_conf.xml	(String) The configuration file for the Cinder Huawei driver.
hypermetro_devices = None	(String) The remote device hypermetro will use.
metro_domain_name = None	(String) The remote metro device domain name.
metro_san_address = None	(String) The remote metro device request url.
metro_san_password = None	(String) The remote metro device san password.
metro_san_user = None	(String) The remote metro device san user.
metro_storage_pools = None	(String) The remote metro device pool names.

Table 3.27. Description of IBM Storage driver configuration options

Configuration option = Default value	Description
[DEFAULT]	

Configuration option = Default value	Description
proxy = storage.proxy.IBMStorageProxy	(String) Proxy driver that connects to the IBM Storage Array
san_clustername =	(String) Cluster name to use for creating volumes
san_ip =	(String) IP address of SAN controller
san_login = admin	(String) Username for SAN controller
san_password =	(String) Password for SAN controller

Table 3.28. Description of images configuration options

Configuration option = Default value	Description
[DEFAULT]	
allowed_direct_url_schemes =	(List) A list of url schemes that can be downloaded directly via the direct_url. Currently supported schemes: [file].
glance_api_insecure = False	(Boolean) Allow to perform insecure SSL (https) requests to glance (https will be used but cert validation will not be performed).
glance_api_servers = None	(List) A list of the URLs of glance API servers available to cinder ([http[s]://][hostname[ip]:port]). If protocol is not specified it defaults to http.
glance_api_ssl_compression = False	(Boolean) 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.
glance_api_version = 1	(Integer) Version of the glance API to use
glance_ca_certificates_file = None	(String) Location of ca certificates file to use for glance client requests.
glance_catalog_info = image:glance:publicURL	(String) Info to match when looking for glance in the service catalog. Format is: separated values of the form: <service_type>:<service_name>:<endpoint_type> - Only used if glance_api_servers are not provided.

Configuration option = Default value	Description
glance_core_properties = checksum, container_format, disk_format, image_name, image_id, min_disk, min_ram, name, size	(List) Default core properties of image
glance_num_retries = 0	(Integer) Number retries when downloading an image from glance
glance_request_timeout = None	(Integer) http/https timeout value for glance operations. If no value (None) is supplied here, the glanceclient default value is used.
image_conversion_dir = \$state_path/conversion	(String) Directory used for temporary storage during image conversion
image_upload_use_cinder_backend = False	(Boolean) If set to True, upload-to-image in raw format will create a cloned volume and register its location to the image service, instead of uploading the volume content. The cinder backend and locations support must be enabled in the image service, and glance_api_version must be set to 2.
image_upload_use_internal_tenant = False	(Boolean) If set to True, the image volume created by upload-to-image will be placed in the internal tenant. Otherwise, the image volume is created in the current context's tenant.
image_volume_cache_enabled = False	(Boolean) Enable the image volume cache for this backend.
image_volume_cache_max_count = 0	(Integer) Max number of entries allowed in the image volume cache. 0 => unlimited.
image_volume_cache_max_size_gb = 0	(Integer) Max size of the image volume cache for this backend in GB. 0 => unlimited.
use_multipath_for_image_xfer = False	(Boolean) Do we attach/detach volumes in cinder using multipath for volume to image and image to volume transfers?

Table 3.29. Description of LVM configuration options

Configuration option = Default value	Description
[DEFAULT]	

Configuration option = Default value	Description
lvm_conf_file = /etc/cinder/lvm.conf	(String) 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).
lvm_max_over_subscription_ratio = 1.0	(Floating point) max_over_subscription_ratio setting for the LVM driver. If set, this takes precedence over the general max_over_subscription_ratio option. If None, the general option is used.
lvm_mirrors = 0	(Integer) If >0, create LVs with multiple mirrors. Note that this requires lvm_mirrors + 2 PVs with available space
lvm_suppress_fd_warnings = False	(Boolean) Suppress leaked file descriptor warnings in LVM commands.
lvm_type = default	(String) Type of LVM volumes to deploy; (default, thin, or auto). Auto defaults to thin if thin is supported.
volume_group = cinder-volumes	(String) Name for the VG that will contain exported volumes

Table 3.30. Description of NAS configuration options

Configuration option = Default value	Description
[DEFAULT]	
nas_host =	(String) IP address or Hostname of NAS system.
nas_login = admin	(String) User name to connect to NAS system.
nas_mount_options = None	(String) Options used to mount the storage backend file system where Cinder volumes are stored.
nas_password =	(String) Password to connect to NAS system.
nas_private_key =	(String) Filename of private key to use for SSH authentication.

Configuration option = Default value	Description
nas_secure_file_operations = auto	(String) 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.
nas_secure_file_permissions = auto	(String) 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.
nas_share_path =	(String) 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 .
nas_ssh_port = 22	(Port number) SSH port to use to connect to NAS system.

Table 3.31. Description of NetApp 7-Mode iSCSI driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
netapp_login = None	(String) Administrative user account name used to access the storage system or proxy server.
netapp_partner_backend_name = None	(String) 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	(String) Password for the administrative user account specified in the netapp_login option.
netapp_pool_name_search_pattern = (.+)	(String) This option is used to restrict provisioning to the specified pools. Specify the value of this option to be a regular expression which will be applied to the names of objects from the storage backend which represent pools in Cinder. This option is only utilized when the storage protocol is configured to use iSCSI or FC.

Configuration option = Default value	Description
netapp_replication_aggregate_map = None	(Unknown) Multi opt of dictionaries to represent the aggregate mapping between source and destination back ends when using whole back end replication. For every source aggregate associated with a cinder pool (NetApp FlexVol), you would need to specify the destination aggregate on the replication target device. A replication target device is configured with the configuration option replication_device. Specify this option as many times as you have replication devices. Each entry takes the standard dict config form: netapp_replication_aggregate_map = backend_id: <name_of_replication_device_section>,src_aggr_name1:dest_aggr_name1,src_aggr_name2:dest_aggr_name2,...
netapp_server_hostname = None	(String) The hostname (or IP address) for the storage system or proxy server.
netapp_server_port = None	(Integer) 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_size_multiplier = 1.2	(Floating point) The quantity to be multiplied by the requested volume size to ensure enough space is available on the virtual storage server (Vserver) to fulfill the volume creation request. Note: this option is deprecated and will be removed in favor of "reserved_percentage" in the Mitaka release.
netapp_snapmirror_quiesce_timeout = 3600	(Integer) The maximum time in seconds to wait for existing SnapMirror transfers to complete before aborting during a failover.
netapp_storage_family = ontap_cluster	(String) 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	(String) The storage protocol to be used on the data path with the storage system.
netapp_transport_type = http	(String) The transport protocol used when communicating with the storage system or proxy server.
netapp_vfiler = None	(String) 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.

Table 3.32. Description of NetApp 7-Mode NFS driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
expiry_thres_minutes = 720	(Integer) 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	(String) Administrative user account name used to access the storage system or proxy server.
netapp_partner_backend_name = None	(String) 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	(String) Password for the administrative user account specified in the netapp_login option.
netapp_pool_name_search_pattern = (.+)	(String) This option is used to restrict provisioning to the specified pools. Specify the value of this option to be a regular expression which will be applied to the names of objects from the storage backend which represent pools in Cinder. This option is only utilized when the storage protocol is configured to use iSCSI or FC.
netapp_replication_aggregate_map = None	(Unknown) Multi opt of dictionaries to represent the aggregate mapping between source and destination back ends when using whole back end replication. For every source aggregate associated with a cinder pool (NetApp FlexVol), you would need to specify the destination aggregate on the replication target device. A replication target device is configured with the configuration option replication_device. Specify this option as many times as you have replication devices. Each entry takes the standard dict config form: netapp_replication_aggregate_map = backend_id: <name_of_replication_device_section>,src_aggr_name1:dest_aggr_name1,src_aggr_name2:dest_aggr_name2,...
netapp_server_hostname = None	(String) The hostname (or IP address) for the storage system or proxy server.

Configuration option = Default value	Description
netapp_server_port = None	(Integer) 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_snapmirror_quiesce_timeout = 3600	(Integer) The maximum time in seconds to wait for existing SnapMirror transfers to complete before aborting during a failover.
netapp_storage_family = ontap_cluster	(String) The storage family type used on the storage system; valid values are <code>ontap_7mode</code> for using Data ONTAP operating in 7-Mode, <code>ontap_cluster</code> for using clustered Data ONTAP, or <code>eseries</code> for using E-Series.
netapp_storage_protocol = None	(String) The storage protocol to be used on the data path with the storage system.
netapp_transport_type = http	(String) The transport protocol used when communicating with the storage system or proxy server.
netapp_vfiler = None	(String) 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	(Integer) 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	(Integer) When the percentage of available space on an NFS share has reached the percentage specified by this option, the driver will stop clearing files from the NFS image cache that have not been accessed in the last M minutes, where M is the value of the <code>expiry_thres_minutes</code> configuration option.

Table 3.33. Description of NetApp cDOT iSCSI driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
netapp_login = None	(String) Administrative user account name used to access the storage system or proxy server.

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

Configuration option = Default value	Description
netapp_size_multiplier = 1.2	(Floating point) The quantity to be multiplied by the requested volume size to ensure enough space is available on the virtual storage server (Vserver) to fulfill the volume creation request. Note: this option is deprecated and will be removed in favor of "reserved_percentage" in the Mitaka release.
netapp_snapmirror_quiesce_timeout = 3600	(Integer) The maximum time in seconds to wait for existing SnapMirror transfers to complete before aborting during a failover.
netapp_storage_family = ontap_cluster	(String) 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	(String) The storage protocol to be used on the data path with the storage system.
netapp_transport_type = http	(String) The transport protocol used when communicating with the storage system or proxy server.
netapp_vserver = None	(String) This option specifies the virtual storage server (Vserver) name on the storage cluster on which provisioning of block storage volumes should occur.

Table 3.34. Description of NetApp cDOT NFS driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
expiry_thres_minutes = 720	(Integer) 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_copyoffload_tool_path = None	(String) 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.

Configuration option = Default value	Description
netapp_host_type = None	(String) This option defines the type of operating system for all initiators that can access a LUN. This information is used when mapping LUNs to individual hosts or groups of hosts.
netapp_host_type = None	(String) This option defines the type of operating system for all initiators that can access a LUN. This information is used when mapping LUNs to individual hosts or groups of hosts.
netapp_login = None	(String) Administrative user account name used to access the storage system or proxy server.
netapp_lun_ostype = None	(String) This option defines the type of operating system that will access a LUN exported from Data ONTAP; it is assigned to the LUN at the time it is created.
netapp_partner_backend_name = None	(String) 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	(String) Password for the administrative user account specified in the netapp_login option.
netapp_pool_name_search_pattern = (.+)	(String) This option is used to restrict provisioning to the specified pools. Specify the value of this option to be a regular expression which will be applied to the names of objects from the storage backend which represent pools in Cinder. This option is only utilized when the storage protocol is configured to use iSCSI or FC.
netapp_replication_aggregate_map = None	(Unknown) Multi opt of dictionaries to represent the aggregate mapping between source and destination back ends when using whole back end replication. For every source aggregate associated with a cinder pool (NetApp FlexVol), you would need to specify the destination aggregate on the replication target device. A replication target device is configured with the configuration option replication_device. Specify this option as many times as you have replication devices. Each entry takes the standard dict config form: netapp_replication_aggregate_map = backend_id: <name_of_replication_device_section>,src_aggr_name1:dest_aggr_name1,src_aggr_name2:dest_aggr_name2,...
netapp_server_hostname = None	(String) The hostname (or IP address) for the storage system or proxy server.

Configuration option = Default value	Description
netapp_server_port = None	(Integer) 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_snapmirror_quiesce_timeout = 3600	(Integer) The maximum time in seconds to wait for existing SnapMirror transfers to complete before aborting during a failover.
netapp_storage_family = ontap_cluster	(String) The storage family type used on the storage system; valid values are <code>ontap_7mode</code> for using Data ONTAP operating in 7-Mode, <code>ontap_cluster</code> for using clustered Data ONTAP, or <code>eseries</code> for using E-Series.
netapp_storage_protocol = None	(String) The storage protocol to be used on the data path with the storage system.
netapp_transport_type = http	(String) The transport protocol used when communicating with the storage system or proxy server.
netapp_vserver = None	(String) This option specifies the virtual storage server (Vserver) name on the storage cluster on which provisioning of block storage volumes should occur.
thres_avl_size_perc_start = 20	(Integer) 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	(Integer) When the percentage of available space on an NFS share has reached the percentage specified by this option, the driver will stop clearing files from the NFS image cache that have not been accessed in the last M minutes, where M is the value of the <code>expiry_thres_minutes</code> configuration option.

Table 3.35. Description of NetApp E-Series driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
netapp_controller_ips = None	(String) This option is only utilized when the storage family is configured to <code>eseries</code> . This option is used to restrict provisioning to the specified controllers. Specify the value of this option to be a comma separated list of controller hostnames or IP addresses to be used for provisioning.

Configuration option = Default value	Description
netapp_enable_multiattach = False	(Boolean) This option specifies whether the driver should allow operations that require multiple attachments to a volume. An example would be live migration of servers that have volumes attached. When enabled, this backend is limited to 256 total volumes in order to guarantee volumes can be accessed by more than one host.
netapp_host_type = None	(String) This option defines the type of operating system for all initiators that can access a LUN. This information is used when mapping LUNs to individual hosts or groups of hosts.
netapp_login = None	(String) Administrative user account name used to access the storage system or proxy server.
netapp_partner_backend_name = None	(String) 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	(String) Password for the administrative user account specified in the netapp_login option.
netapp_pool_name_search_pattern = (.+)	(String) This option is used to restrict provisioning to the specified pools. Specify the value of this option to be a regular expression which will be applied to the names of objects from the storage backend which represent pools in Cinder. This option is only utilized when the storage protocol is configured to use iSCSI or FC.
netapp_replication_aggregate_map = None	(Unknown) Multi opt of dictionaries to represent the aggregate mapping between source and destination back ends when using whole back end replication. For every source aggregate associated with a cinder pool (NetApp FlexVol), you would need to specify the destination aggregate on the replication target device. A replication target device is configured with the configuration option replication_device. Specify this option as many times as you have replication devices. Each entry takes the standard dict config form: netapp_replication_aggregate_map = backend_id: <name_of_replication_device_section>,src_aggr_name1:dest_aggr_name1,src_aggr_name2:dest_aggr_name2,...
netapp_sa_password = None	(String) Password for the NetApp E-Series storage array.
netapp_server_hostname = None	(String) The hostname (or IP address) for the storage system or proxy server.

Configuration option = Default value	Description
netapp_server_port = None	(Integer) 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_snapmirror_quiesce_timeout = 3600	(Integer) The maximum time in seconds to wait for existing SnapMirror transfers to complete before aborting during a failover.
netapp_storage_family = ontap_cluster	(String) The storage family type used on the storage system; valid values are <code>ontap_7mode</code> for using Data ONTAP operating in 7-Mode, <code>ontap_cluster</code> for using clustered Data ONTAP, or <code>eseries</code> for using E-Series.
netapp_transport_type = http	(String) The transport protocol used when communicating with the storage system or proxy server.
netapp_webservice_path = /devmgr/v2	(String) 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.

Table 3.36. Description of Nexenta driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
nexenta_blocksize = 4096	(Integer) Block size for datasets
nexenta_chunksize = 32768	(Integer) NexentaEdge iSCSI LUN object chunk size
nexenta_client_address =	(String) NexentaEdge iSCSI Gateway client address for non-VIP service
nexenta_dataset_compression = on	(String) Compression value for new ZFS folders.
nexenta_dataset_dedup = off	(String) Deduplication value for new ZFS folders.
nexenta_dataset_description =	(String) Human-readable description for the folder.
nexenta_host =	(String) IP address of Nexenta SA

Configuration option = Default value	Description
nexenta_iscsi_target_portal_port = 3260	(Integer) Nexenta target portal port
nexenta_mount_point_base = \$state_path/mnt	(String) Base directory that contains NFS share mount points
nexenta_nbd_symlinks_dir = /dev/disk/by-path	(String) NexentaEdge logical path of directory to store symbolic links to NBDs
nexenta_password = nexenta	(String) Password to connect to Nexenta SA
nexenta_rest_port = 8080	(Integer) HTTP port to connect to Nexenta REST API server
nexenta_rest_protocol = auto	(String) Use http or https for REST connection (default auto)
nexenta_rrmgr_compression = 0	(Integer) Enable stream compression, level 1..9. 1 - gives best speed; 9 - gives best compression.
nexenta_rrmgr_connections = 2	(Integer) Number of TCP connections.
nexenta_rrmgr_tcp_buf_size = 4096	(Integer) TCP Buffer size in KiloBytes.
nexenta_shares_config = /etc/cinder/nfs_shares	(String) File with the list of available nfs shares
nexenta_sparse = False	(Boolean) Enables or disables the creation of sparse datasets
nexenta_sparsed_volumes = True	(Boolean) Enables or disables the creation of volumes as sparsed files that take no space. If disabled (False), volume is created as a regular file, which takes a long time.
nexenta_target_group_prefix = cinder/	(String) Prefix for iSCSI target groups on SA
nexenta_target_prefix = iqn.1986-03.com.sun:02:cinder-	(String) IQN prefix for iSCSI targets
nexenta_user = admin	(String) User name to connect to Nexenta SA
nexenta_volume = cinder	(String) SA Pool that holds all volumes

Table 3.37. Description of os-brick configuration options

Configuration option = Default value	Description
[privsep_osbrick]	
capabilities = []	(Unknown) List of Linux capabilities retained by the privsep daemon.
group = None	(String) Group that the privsep daemon should run as.
helper_command = None	(String) Command to invoke to start the privsep daemon if not using the "fork" method. If not specified, a default is generated using "sudo privsep-helper" and arguments designed to recreate the current configuration. This command must accept suitable --privsep_context and --privsep_sock_path arguments.
user = None	(String) User that the privsep daemon should run as.

Table 3.38. Description of profiler configuration options

Configuration option = Default value	Description
[profiler]	
connection_string = messaging://	(String) Connection string for a notifier backend. Default value is messaging:// which sets the notifier to oslo_messaging. Examples of possible values: <ul style="list-style-type: none"> ● messaging://: use oslo_messaging driver for sending notifications.
enabled = False	(Boolean) Enables the profiling for all services on this node. Default value is False (fully disable the profiling feature). Possible values: <ul style="list-style-type: none"> ● True: Enables the feature ● False: Disables the feature. The profiling cannot be started via this project operations. If the profiling is triggered by another project, this project part will be empty.

Configuration option = Default value	Description
hmac_keys = SECRET_KEY	<p>(String) Secret key(s) to use for encrypting context data for performance profiling. This string value should have the following format: <key1>[,<key2>,...<keyn>], where each key is some random string. A user who triggers the profiling via the REST API has to set one of these keys in the headers of the REST API call to include profiling results of this node for this particular project.</p> <p>Both "enabled" flag and "hmac_keys" config options should be set to enable profiling. Also, to generate correct profiling information across all services at least one key needs to be consistent between OpenStack projects. This ensures it can be used from client side to generate the trace, containing information from all possible resources.</p>
trace_sqlalchemy = False	<p>(Boolean) Enables SQL requests profiling in services. Default value is False (SQL requests won't be traced).</p> <p>Possible values:</p> <ul style="list-style-type: none"> • True: Enables SQL requests profiling. Each SQL query will be part of the trace and can be analyzed by how much time was spent for that. • False: Disables SQL requests profiling. The spent time is only shown on a higher level of operations. Single SQL queries cannot be analyzed this way.

Table 3.39. Description of Pure Storage driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
pure_api_token = None	(String) REST API authorization token.
pure_automatic_max_oversubscription_ratio = True	(Boolean) Automatically determine an oversubscription ratio based on the current total data reduction values. If used this calculated value will override the max_over_subscription_ratio config option.

Configuration option = Default value	Description
pure_eradicate_on_delete = False	(Boolean) When enabled, all Pure volumes, snapshots, and protection groups will be eradicated at the time of deletion in Cinder. Data will NOT be recoverable after a delete with this set to True! When disabled, volumes and snapshots will go into pending eradication state and can be recovered.
pure_replica_interval_default = 900	(Integer) Snapshot replication interval in seconds.
pure_replica_retention_long_term_default = 7	(Integer) Retain snapshots per day on target for this time (in days.)
pure_replica_retention_long_term_per_day_default = 3	(Integer) Retain how many snapshots for each day.
pure_replica_retention_short_term_default = 14400	(Integer) Retain all snapshots on target for this time (in seconds.)

Table 3.40. Description of Quobyte USP volume driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
quobyte_client_cfg = None	(String) Path to a Quobyte Client configuration file.
quobyte_mount_point_base = \$state_path/mnt	(String) Base dir containing the mount point for the Quobyte volume.
quobyte_qcow2_volumes = True	(Boolean) Create volumes as QCOW2 files rather than raw files.
quobyte_sparsed_volumes = True	(Boolean) 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.
quobyte_volume_url = None	(String) URL to the Quobyte volume e.g., quobyte://<DIR host>/<volume name>

Table 3.41. Description of quota configuration options

Configuration option = Default value	Description
[DEFAULT]	
max_age = 0	(Integer) Number of seconds between subsequent usage refreshes
quota_backup_gigabytes = 1000	(Integer) Total amount of storage, in gigabytes, allowed for backups per project
quota_backups = 10	(Integer) Number of volume backups allowed per project
quota_consistencygroups = 10	(Integer) Number of consistencygroups allowed per project
quota_driver = cinder.quota.DbQuotaDriver	(String) Default driver to use for quota checks
quota_gigabytes = 1000	(Integer) Total amount of storage, in gigabytes, allowed for volumes and snapshots per project
quota_groups = 10	(Integer) Number of groups allowed per project
quota_snapshots = 10	(Integer) Number of volume snapshots allowed per project
quota_volumes = 10	(Integer) Number of volumes allowed per project
reservation_expire = 86400	(Integer) Number of seconds until a reservation expires
use_default_quota_class = True	(Boolean) Enables or disables use of default quota class with default quota.

Table 3.42. Description of Redis configuration options

Configuration option = Default value	Description
[matchmaker_redis]	
check_timeout = 20000	(Integer) Time in ms to wait before the transaction is killed.
host = 127.0.0.1	(String) DEPRECATED: Host to locate redis. Replaced by [DEFAULT]/transport_url
password =	(String) DEPRECATED: Password for Redis server (optional). Replaced by [DEFAULT]/transport_url

Configuration option = Default value	Description
port = 6379	(Port number) DEPRECATED: Use this port to connect to redis host. Replaced by [DEFAULT]/transport_url
sentinel_hosts =	(List) DEPRECATED: List of Redis Sentinel hosts (fault tolerance mode) e.g. [host:port, host1:port ...] Replaced by [DEFAULT]/transport_url
socket_timeout = 10000	(Integer) Timeout in ms on blocking socket operations
wait_timeout = 2000	(Integer) Time in ms to wait between connection attempts.

Table 3.43. Description of SAN configuration options

Configuration option = Default value	Description
[DEFAULT]	
san_clustername =	(String) Cluster name to use for creating volumes
san_ip =	(String) IP address of SAN controller
san_is_local = False	(Boolean) Execute commands locally instead of over SSH; use if the volume service is running on the SAN device
san_login = admin	(String) Username for SAN controller
san_password =	(String) Password for SAN controller
san_private_key =	(String) Filename of private key to use for SSH authentication
san_ssh_port = 22	(Port number) SSH port to use with SAN
san_thin_provision = True	(Boolean) Use thin provisioning for SAN volumes?
ssh_conn_timeout = 30	(Integer) SSH connection timeout in seconds
ssh_max_pool_conn = 5	(Integer) Maximum ssh connections in the pool
ssh_min_pool_conn = 1	(Integer) Minimum ssh connections in the pool

Table 3.44. Description of Scality SOFS volume driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
scality_sofs_config = None	(String) Path or URL to Scality SOFS configuration file
scality_sofs_mount_point = \$state_path/scality	(String) Base dir where Scality SOFS shall be mounted
scality_sofs_volume_dir = cinder/volumes	(String) Path from Scality SOFS root to volume dir

Table 3.45. Description of scheduler configuration options

Configuration option = Default value	Description
[DEFAULT]	
filter_function = None	(String) 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.
goodness_function = None	(String) 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.
scheduler_default_filters = AvailabilityZoneFilter, CapacityFilter, CapabilitiesFilter	(List) Which filter class names to use for filtering hosts when not specified in the request.
scheduler_default_weighers = CapacityWeigher	(List) Which weigher class names to use for weighing hosts.
scheduler_driver = cinder.scheduler.filter_scheduler.FilterScheduler	(String) Default scheduler driver to use
scheduler_host_manager = cinder.scheduler.host_manager.HostManager	(String) The scheduler host manager class to use
scheduler_json_config_location =	(String) Absolute path to scheduler configuration JSON file.
scheduler_manager = cinder.scheduler.manager.SchedulerManager	(String) Full class name for the Manager for scheduler

Configuration option = Default value	Description
scheduler_max_attempts = 3	(Integer) Maximum number of attempts to schedule a volume
scheduler_weight_handler = cinder.scheduler.weights.OrderedHostWeightHandler	(String) Which handler to use for selecting the host/pool after weighing

Table 3.46. Description of SCST volume driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
scst_target_driver = iscsi	(String) SCST target implementation can choose from multiple SCST target drivers.
scst_target_iqn_name = None	(String) Certain ISCSI targets have predefined target names, SCST target driver uses this name.

Table 3.47. Description of Sheepdog driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
sheepdog_store_address = 127.0.0.1	(String) IP address of sheep daemon.
sheepdog_store_port = 7000	(Port number) Port of sheep daemon.

Table 3.48. Description of Samba volume driver configuration options

Configuration option = Default value	Description
[DEFAULT]	

Configuration option = Default value	Description
smbfs_allocation_info_file_path = \$state_path/allocation_data	(String) The path of the automatically generated file containing information about volume disk space allocation.
smbfs_default_volume_format = qcow2	(String) Default format that will be used when creating volumes if no volume format is specified.
smbfs_mount_options = noperm,file_mode=0775,dir_mode=0775	(String) Mount options passed to the smbfs client. See mount.cifs man page for details.
smbfs_mount_point_base = \$state_path/mnt	(String) Base dir containing mount points for smbfs shares.
smbfs_oversub_ratio = 1.0	(Floating point) 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.
smbfs_shares_config = /etc/cinder/smbfs_shares	(String) File with the list of available smbfs shares.
smbfs_sparsed_volumes = True	(Boolean) Create volumes as sparsed files which take no space rather than regular files when using raw format, in which case volume creation takes lot of time.
smbfs_used_ratio = 0.95	(Floating point) Percent of ACTUAL usage of the underlying volume before no new volumes can be allocated to the volume destination.

Table 3.49. Description of SolidFire driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
sf_account_prefix = None	(String) Create SolidFire accounts with this prefix. Any string can be used here, but the string "hostname" is special and will create a prefix using the cinder node hostname (previous default behavior). The default is NO prefix.
sf_allow_template_caching = True	(Boolean) Create an internal cache of copy of images when a bootable volume is created to eliminate fetch from glance and qemu-conversion on subsequent calls.
sf_allow_tenant_qos = False	(Boolean) Allow tenants to specify QOS on create

Configuration option = Default value	Description
sf_api_port = 443	(Port number) SolidFire API port. Useful if the device api is behind a proxy on a different port.
sf_emulate_512 = True	(Boolean) Set 512 byte emulation on volume creation;
sf_enable_vag = False	(Boolean) Utilize volume access groups on a per-tenant basis.
sf_enable_volume_mapping = True	(Boolean) Create an internal mapping of volume IDs and account. Optimizes lookups and performance at the expense of memory, very large deployments may want to consider setting to False.
sf_svip = None	(String) Overrides default cluster SVIP with the one specified. This is required on deployments that have implemented the use of VLANs for iSCSI networks in their cloud.
sf_template_account_name = openstack-vtemplate	(String) Account name on the SolidFire Cluster to use as owner of template/cache volumes (created if does not exist).
sf_volume_prefix = UUID-	(String) Create SolidFire volumes with this prefix. Volume names are of the form <sf_volume_prefix> <cinder-volume-id>. The default is to use a prefix of 'UUID-'.

Table 3.50. Description of Ceph storage configuration options

Configuration option = Default value	Description
[DEFAULT]	
rados_connect_timeout = -1	(Integer) Timeout value (in seconds) used when connecting to ceph cluster. If value < 0, no timeout is set and default librados value is used.
rados_connection_interval = 5	(Integer) Interval value (in seconds) between connection retries to ceph cluster.
rados_connection_retries = 3	(Integer) Number of retries if connection to ceph cluster failed.
rbd_ceph_conf =	(String) Path to the ceph configuration file
rbd_cluster_name = ceph	(String) The name of ceph cluster

Configuration option = Default value	Description
rbd_flatten_volume_from_snapshot = False	(Boolean) Flatten volumes created from snapshots to remove dependency from volume to snapshot
rbd_max_clone_depth = 5	(Integer) Maximum number of nested volume clones that are taken before a flatten occurs. Set to 0 to disable cloning.
rbd_pool = rbd	(String) The RADOS pool where rbd volumes are stored
rbd_secret_uuid = None	(String) The libvirt uuid of the secret for the rbd_user volumes
rbd_store_chunk_size = 4	(Integer) Volumes will be chunked into objects of this size (in megabytes).
rbd_user = None	(String) The RADOS client name for accessing rbd volumes - only set when using cephx authentication
volume_tmp_dir = None	(String) 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.

Table 3.51. Description of GlusterFS storage configuration options

Configuration option = Default value	Description
[DEFAULT]	
glusterfs_mount_point_base = \$state_path/mnt	(String) Base dir containing mount points for gluster shares.
glusterfs_shares_config = /etc/cinder/glusterfs_shares	(String) File with the list of available gluster shares
nas_volume_prov_type = thin	(String) Provisioning type that will be used when creating volumes.

Table 3.52. Description of GPFS storage configuration options

Configuration option = Default value	Description
[DEFAULT]	
gpfs_images_dir = None	(String) Specifies the path of the Image service repository in GPFS. Leave undefined if not storing images in GPFS.
gpfs_images_share_mode = None	(String) 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.
gpfs_max_clone_depth = 0	(Integer) 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.
gpfs_mount_point_base = None	(String) Specifies the path of the GPFS directory where Block Storage volume and snapshot files are stored.
gpfs_sparse_volumes = True	(Boolean) 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.
gpfs_storage_pool = system	(String) Specifies the storage pool that volumes are assigned to. By default, the system storage pool is used.
nas_host =	(String) IP address or Hostname of NAS system.
nas_login = admin	(String) User name to connect to NAS system.
nas_password =	(String) Password to connect to NAS system.
nas_private_key =	(String) Filename of private key to use for SSH authentication.
nas_ssh_port = 22	(Port number) SSH port to use to connect to NAS system.

Table 3.53. Description of NFS storage configuration options

Configuration option = Default value	Description
[DEFAULT]	
nfs_mount_attempts = 3	(Integer) 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.
nfs_mount_options = None	(String) Mount options passed to the NFS client. See section of the NFS man page for details.
nfs_mount_point_base = \$state_path/mnt	(String) Base dir containing mount points for NFS shares.
nfs_shares_config = /etc/cinder/nfs_shares	(String) File with the list of available NFS shares
nfs_sparsed_volumes = True	(Boolean) 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.

Table 3.54. Description of storage configuration options

Configuration option = Default value	Description
[DEFAULT]	
allocated_capacity_weight_multiplier = -1.0	(Floating point) Multiplier used for weighing allocated capacity. Positive numbers mean to stack vs spread.
capacity_weight_multiplier = 1.0	(Floating point) Multiplier used for weighing free capacity. Negative numbers mean to stack vs spread.
enabled_backends = None	(List) A list of backend names to use. These backend names should be backed by a unique [CONFIG] group with its options
iscsi_helper = tgtadm	(String) iSCSI target user-land tool to use. tgtadm is default, use lioadm for LIO iSCSI support, scstadmin for SCST target support, ietadm for iSCSI Enterprise Target, iscsictl for Chelsio iSCSI Target or fake for testing.
iscsi_iotype = fileio	(String) Sets the behavior of the iSCSI target to either perform blockio or fileio optionally, auto can be set and Cinder will autodetect type of backing device

Configuration option = Default value	Description
iscsi_ip_address = \$my_ip	(String) The IP address that the iSCSI daemon is listening on
iscsi_port = 3260	(Port number) The port that the iSCSI daemon is listening on
iscsi_protocol = iscsi	(String) Determines the iSCSI protocol for new iSCSI volumes, created with tgtadm or lioadm target helpers. In order to enable RDMA, this parameter should be set with the value "iser". The supported iSCSI protocol values are "iscsi" and "iser".
iscsi_target_flags =	(String) Sets the target-specific flags for the iSCSI target. Only used for tgtadm to specify backing device flags using bsoflags option. The specified string is passed as is to the underlying tool.
iscsi_target_prefix = iqn.2010-10.org.openstack:	(String) Prefix for iSCSI volumes
iscsi_write_cache = on	(String) Sets the behavior of the iSCSI target to either perform write-back(on) or write-through(off). This parameter is valid if iscsi_helper is set to tgtadm.
iser_helper = tgtadm	(String) The name of the iSER target user-land tool to use
iser_ip_address = \$my_ip	(String) The IP address that the iSER daemon is listening on
iser_port = 3260	(Port number) The port that the iSER daemon is listening on
iser_target_prefix = iqn.2010-10.org.openstack:	(String) Prefix for iSER volumes
migration_create_volume_timeout_secs = 300	(Integer) Timeout for creating the volume to migrate to when performing volume migration (seconds)
num_iser_scan_tries = 3	(Integer) The maximum number of times to rescan iSER target to find volume
num_volume_device_scan_tries = 3	(Integer) The maximum number of times to rescan targets to find volume
volume_backend_name = None	(String) The backend name for a given driver implementation
volume_clear = zero	(String) Method used to wipe old volumes

Configuration option = Default value	Description
volume_clear_ionice = None	(String) The flag to pass to ionice to alter the i/o priority of the process used to zero a volume after deletion, for example "-c3" for idle only priority.
volume_clear_size = 0	(Integer) Size in MiB to wipe at start of old volumes. 1024 MiB at max. 0 => all
volume_copy_blkio_cgroup_name = cinder-volume-copy	(String) The blkio cgroup name to be used to limit bandwidth of volume copy
volume_copy_bps_limit = 0	(Integer) The upper limit of bandwidth of volume copy. 0 => unlimited
volume_dd_blocksize = 1M	(String) The default block size used when copying/clearing volumes
volume_driver = cinder.volume.drivers.lvm.LVMVolumeDriver	(String) Driver to use for volume creation
volume_manager = cinder.volume.manager.VolumeManager	(String) Full class name for the Manager for volume
volume_service_inithost_offload = False	(Boolean) Offload pending volume delete during volume service startup
volume_usage_audit_period = month	(String) Time period for which to generate volume usages. The options are hour, day, month, or year.
volumes_dir = \$state_path/volumes	(String) Volume configuration file storage directory

Table 3.55. Description of IBM Storwize driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
storwize_san_secondary_ip = None	(String) Specifies secondary management IP or hostname to be used if san_ip is invalid or becomes inaccessible.
storwize_svc_allow_tenant_qos = False	(Boolean) Allow tenants to specify QOS on create
storwize_svc_flashcopy_rate = 50	(Integer) Specifies the Storwize FlashCopy copy rate to be used when creating a full volume copy. The default is rate is 50, and the valid rates are 1-100.

Configuration option = Default value	Description
storwize_svc_flashcopy_timeout = 120	(Integer) Maximum number of seconds to wait for FlashCopy to be prepared.
storwize_svc_iscsi_chap_enabled = True	(Boolean) Configure CHAP authentication for iSCSI connections (Default: Enabled)
storwize_svc_multihostmap_enabled = True	(Boolean) DEPRECATED: This option no longer has any affect. It is deprecated and will be removed in the next release.
storwize_svc_multipath_enabled = False	(Boolean) Connect with multipath (FC only; iSCSI multipath is controlled by Nova)
storwize_svc_stretched_cluster_partner = None	(String) If operating in stretched cluster mode, specify the name of the pool in which mirrored copies are stored.Example: "pool2"
storwize_svc_vol_autoexpand = True	(Boolean) Storage system autoexpand parameter for volumes (True/False)
storwize_svc_vol_compression = False	(Boolean) Storage system compression option for volumes
storwize_svc_vol_easytier = True	(Boolean) Enable Easy Tier for volumes
storwize_svc_vol_grainsize = 256	(Integer) Storage system grain size parameter for volumes (32/64/128/256)
storwize_svc_vol_iogrp = 0	(Integer) The I/O group in which to allocate volumes
storwize_svc_vol_nofmtdisk = False	(Boolean) Specifies that the volume not be formatted during creation.
storwize_svc_vol_rsize = 2	(Integer) Storage system space-efficiency parameter for volumes (percentage)
storwize_svc_vol_warning = 0	(Integer) Storage system threshold for volume capacity warnings (percentage)
storwize_svc_volpool_name = volpool	(List) Comma separated list of storage system storage pools for volumes.

Table 3.56. Description of swift configuration options

Configuration option = Default value	Description
[DEFAULT]	

Configuration option = Default value	Description
backup_swift_auth_insecure = False	(Boolean) Bypass verification of server certificate when making SSL connection to Swift.
backup_swift_auth_url = None	(String) The URL of the Keystone endpoint

Table 3.57. Description of Tintri volume driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
tintri_api_version = v310	(String) API version for the storage system
tintri_image_cache_expiry_days = 30	(Integer) Delete unused image snapshots older than mentioned days
tintri_image_shares_config = None	(String) Path to image nfs shares file
tintri_server_hostname = None	(String) The hostname (or IP address) for the storage system
tintri_server_password = None	(String) Password for the storage system
tintri_server_username = None	(String) User name for the storage system

Table 3.58. Description of Violin volume driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
violin_dedup_capable_pools =	(List) Storage pools capable of dedup and other luns. (Comma separated list)
violin_dedup_only_pools =	(List) Storage pools to be used to setup dedup luns only. (Comma separated list)

Configuration option = Default value	Description
violin_iscsi_target_ips =	(List) Target iSCSI addresses to use.(Comma separated list)
violin_pool_allocation_method = random	(String) Method of choosing a storage pool for a lun.
violin_request_timeout = 300	(Integer) Global backend request timeout, in seconds.

Table 3.59. Description of Windows configuration options

Configuration option = Default value	Description
[DEFAULT]	
windows_iscsi_lun_path = C:\iSCSIVirtualDisks	(String) Path to store VHD backed volumes

Table 3.60. Description of ZFS Storage Appliance NFS driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
zfssa_cache_directory = os-cinder-cache	(String) Name of directory inside zfssa_nfs_share where cache volumes are stored.
zfssa_cache_project = os-cinder-cache	(String) Name of ZFSSA project where cache volumes are stored.
zfssa_data_ip = None	(String) Data path IP address
zfssa_enable_local_cache = True	(Boolean) Flag to enable local caching: True, False.
zfssa_https_port = 443	(String) HTTPS port number
zfssa_manage_policy = loose	(String) Driver policy for volume manage.
zfssa_nfs_mount_options =	(String) Options to be passed while mounting share over nfs
zfssa_nfs_pool =	(String) Storage pool name.

Configuration option = Default value	Description
zfssa_nfs_project = NFSProject	(String) Project name.
zfssa_nfs_share = nfs_share	(String) Share name.
zfssa_nfs_share_compression = off	(String) Data compression.
zfssa_nfs_share_logbias = latency	(String) Synchronous write bias-latency, throughput.
zfssa_rest_timeout = None	(Integer) REST connection timeout. (seconds)

Table 3.61. Description of zones configuration options

Configuration option = Default value	Description
[DEFAULT]	
cloned_volume_same_az = True	(Boolean) Ensure that the new volumes are the same AZ as snapshot or source volume

Table 3.62. Description of brocade zoning fabrics configuration options

Configuration option = Default value	Description
[BRCD_FABRIC_EXAMPLE]	
fc_fabric_address =	(String) Management IP of fabric.
fc_fabric_password =	(String) Password for user.
fc_fabric_port = 22	(Port number) Connecting port
fc_fabric_ssh_cert_path =	(String) Local SSH certificate Path.
fc_fabric_user =	(String) Fabric user ID.
fc_southbound_protocol = HTTP	(String) South bound connector for the fabric.
fc_virtual_fabric_id = None	(String) Virtual Fabric ID.

Configuration option = Default value	Description
principal_switch_wwn = None	(String) DEPRECATED: Principal switch WWN of the fabric. This option is not used anymore.
zone_activate = True	(Boolean) Overridden zoning activation state.
zone_name_prefix = openstack	(String) Overridden zone name prefix.
zoning_policy = initiator-target	(String) Overridden zoning policy.

Table 3.63. Description of cisco zoning fabrics configuration options

Configuration option = Default value	Description
[CISCO_FABRIC_EXAMPLE]	
cisco_fc_fabric_address =	(String) Management IP of fabric
cisco_fc_fabric_password =	(String) Password for user
cisco_fc_fabric_port = 22	(Port number) Connecting port
cisco_fc_fabric_user =	(String) Fabric user ID
cisco_zone_activate = True	(Boolean) overridden zoning activation state
cisco_zone_name_prefix = None	(String) overridden zone name prefix
cisco_zoning_policy = initiator-target	(String) overridden zoning policy
cisco_zoning_vsan = None	(String) VSAN of the Fabric

Table 3.64. Description of brocade zoning manager configuration options

Configuration option = Default value	Description
[fc-zone-manager]	
brcd_sb_connector = HTTP	(String) South bound connector for zoning operation

Table 3.65. Description of cisco zoning manager configuration options

Configuration option = Default value	Description
[fc-zone-manager]	
cisco_sb_connector = cinder.zonemanager.drivers.cisco.cisco_fc_z one_client_cli.CiscoFCZoneClientCLI	(String) Southbound connector for zoning operation

Table 3.66. Description of zoning configuration options

Configuration option = Default value	Description
[DEFAULT]	
zoning_mode = None	(String) FC Zoning mode configured
[fc-zone-manager]	
enable_unsupported_driver = False	(Boolean) Set this to True when you want to allow an unsupported zone manager driver to start. Drivers that haven't maintained a working CI system and testing are marked as unsupported until CI is working again. This also marks a driver as deprecated and may be removed in the next release.
fc_fabric_names = None	(String) 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
fc_san_lookup_service = cinder.zonemanager.drivers.brocade.brcd_fc _san_lookup_service.BrcdFCSanLookupSer vice	(String) FC SAN Lookup Service
zone_driver = cinder.zonemanager.drivers.brocade.brcd_fc _zone_driver.BrcdFCZoneDriver	(String) FC Zone Driver responsible for zone management
zoning_policy = initiator-target	(String) Zoning policy configured by user; valid values include "initiator-target" or "initiator"

3.1.2. New, Updated, and Deprecated Options in Newton for Block Storage

Table 3.67. New options

Option = default value	(Type) Help string
[DEFAULT] additional_retry_list =	(StrOpt) FSS additional retry list, separate by ;
[DEFAULT] backup_swift_project = None	(StrOpt) Swift project/account name. Required when connecting to an auth 3.0 system
[DEFAULT] backup_swift_project_domain = None	(StrOpt) Swift project domain name. Required when connecting to an auth 3.0 system
[DEFAULT] backup_swift_user_domain = None	(StrOpt) Swift user domain name. Required when connecting to an auth 3.0 system
[DEFAULT] backup_use_temp_snapshot = False	(BoolOpt) If this is set to True, the backup_use_temp_snapshot path will be used during the backup. Otherwise, it will use backup_use_temp_volume path.
[DEFAULT] chap = disabled	(StrOpt) CHAP authentication mode, effective only for iscsi (disabled enabled)
[DEFAULT] clone_volume_timeout = 680	(IntOpt) Create clone volume timeout.
[DEFAULT] cluster = None	(StrOpt) Name of this cluster. Used to group volume hosts that share the same backend configurations to work in HA Active-Active mode. Active-Active is not yet supported.
[DEFAULT] connection_type = iscsi	(StrOpt) Connection type to the IBM Storage Array
[DEFAULT] coprhd_emulate_snapshot = False	(BoolOpt) True False to indicate if the storage array in CoprHD is VMAX or VPLEX
[DEFAULT] coprhd_hostname = None	(StrOpt) Hostname for the CoprHD Instance
[DEFAULT] coprhd_password = None	(StrOpt) Password for accessing the CoprHD Instance
[DEFAULT] coprhd_port = 4443	(PortOpt) Port for the CoprHD Instance
[DEFAULT] coprhd_project = None	(StrOpt) Project to utilize within the CoprHD Instance
[DEFAULT] coprhd_scaleio_rest_gateway_host = None	(StrOpt) Rest Gateway IP or FQDN for Scaleio
[DEFAULT] coprhd_scaleio_rest_gateway_port = 4984	(PortOpt) Rest Gateway Port for Scaleio
[DEFAULT] coprhd_scaleio_rest_server_password = None	(StrOpt) Rest Gateway Password

Option = default value	(Type) Help string
[DEFAULT] coprhd_scaleio_rest_server_username = None	(StrOpt) Username for Rest Gateway
[DEFAULT] coprhd_tenant = None	(StrOpt) Tenant to utilize within the CoprHD Instance
[DEFAULT] coprhd_username = None	(StrOpt) Username for accessing the CoprHD Instance
[DEFAULT] coprhd_varray = None	(StrOpt) Virtual Array to utilize within the CoprHD Instance
[DEFAULT] datera_503_interval = 5	(IntOpt) Interval between 503 retries
[DEFAULT] datera_503_timeout = 120	(IntOpt) Timeout for HTTP 503 retry messages
[DEFAULT] datera_acl_allow_all = False	(BoolOpt) True to set acl 'allow_all' on volumes created
[DEFAULT] datera_debug = False	(BoolOpt) True to set function arg and return logging
[DEFAULT] datera_debug_replica_count_override = False	(BoolOpt) ONLY FOR DEBUG/TESTING PURPOSES True to set replica_count to 1
[DEFAULT] default_group_type = None	(StrOpt) Default group type to use
[DEFAULT] dell_server_os = Red Hat Linux 6.x	(StrOpt) Server OS type to use when creating a new server on the Storage Center.
[DEFAULT] drbdmanage_disk_options = {"c-min-rate": "4M"}	(StrOpt) Disk options to set on new resources. See http://www.drbd.org/en/doc/users-guide-90/re-drbdconf for all the details.
[DEFAULT] drbdmanage_net_options = {"connect-int": "4", "allow-two-primaries": "yes", "ko-count": "30", "max-buffers": "20000", "ping-timeout": "100"}	(StrOpt) Net options to set on new resources. See http://www.drbd.org/en/doc/users-guide-90/re-drbdconf for all the details.
[DEFAULT] drbdmanage_resource_options = {"auto-promote-timeout": "300"}	(StrOpt) Resource options to set on new resources. See http://www.drbd.org/en/doc/users-guide-90/re-drbdconf for all the details.
[DEFAULT] dsware_isthin = False	(BoolOpt) The flag of thin storage allocation.
[DEFAULT] dsware_manager =	(StrOpt) Fusionstorage manager ip addr for cinder-volume.

Option = default value	(Type) Help string
[DEFAULT] enable_unsupported_driver = False	(BoolOpt) Set this to True when you want to allow an unsupported driver to start. Drivers that haven't maintained a working CI system and testing are marked as unsupported until CI is working again. This also marks a driver as deprecated and may be removed in the next release.
[DEFAULT] fss_debug = False	(BoolOpt) Enable HTTP debugging to FSS
[DEFAULT] fss_pool =	(IntOpt) FSS pool id in which FalconStor volumes are stored.
[DEFAULT] fusionstorageagent =	(StrOpt) Fusionstorage agent ip addr range.
[DEFAULT] glance_catalog_info = image:glance:publicURL	(StrOpt) Info to match when looking for glance in the service catalog. Format is: separated values of the form: <service_type>:<service_name>:<endpoint_type> - Only used if glance_api_servers are not provided.
[DEFAULT] group_api_class = cinder.group.api.API	(StrOpt) The full class name of the group API class
[DEFAULT] hnas_chap_enabled = True	(BoolOpt) Whether the chap authentication is enabled in the iSCSI target or not.
[DEFAULT] hnas_cluster_admin_ip0 = None	(StrOpt) The IP of the HNAS cluster admin. Required only for HNAS multi-cluster setups.
[DEFAULT] hnas_mgmt_ip0 = None	(IPOpt) Management IP address of HNAS. This can be any IP in the admin address on HNAS or the SMU IP.
[DEFAULT] hnas_password = None	(StrOpt) HNAS password.
[DEFAULT] hnas_ssc_cmd = ssc	(StrOpt) Command to communicate to HNAS.
[DEFAULT] hnas_ssh_port = 22	(PortOpt) Port to be used for SSH authentication.
[DEFAULT] hnas_ssh_private_key = None	(StrOpt) Path to the SSH private key used to authenticate in HNAS SMU.
[DEFAULT] hnas_svc0_hdp = None	(StrOpt) Service 0 HDP
[DEFAULT] hnas_svc0_iscsi_ip = None	(IPOpt) Service 0 iSCSI IP
[DEFAULT] hnas_svc0_volume_type = None	(StrOpt) Service 0 volume type
[DEFAULT] hnas_svc1_hdp = None	(StrOpt) Service 1 HDP
[DEFAULT] hnas_svc1_iscsi_ip = None	(IPOpt) Service 1 iSCSI IP

Option = default value	(Type) Help string
[DEFAULT] hnas_svc1_volume_type = None	(StrOpt) Service 1 volume type
[DEFAULT] hnas_svc2_hdp = None	(StrOpt) Service 2 HDP
[DEFAULT] hnas_svc2_iscsi_ip = None	(IPOpt) Service 2 iSCSI IP
[DEFAULT] hnas_svc2_volume_type = None	(StrOpt) Service 2 volume type
[DEFAULT] hnas_svc3_hdp = None	(StrOpt) Service 3 HDP
[DEFAULT] hnas_svc3_iscsi_ip = None	(IPOpt) Service 3 iSCSI IP
[DEFAULT] hnas_svc3_volume_type = None	(StrOpt) Service 3 volume type
[DEFAULT] hnas_username = None	(StrOpt) HNAS username.
[DEFAULT] kaminario_nodedup_substring = K2-nodedup	(StrOpt) If volume-type name contains this substring nodedup volume will be created, otherwise dedup volume will be created.
[DEFAULT] lvm_suppress_fd_warnings = False	(BoolOpt) Suppress leaked file descriptor warnings in LVM commands.
[DEFAULT] message_ttl = 2592000	(IntOpt) message minimum life in seconds.
[DEFAULT] metro_domain_name = None	(StrOpt) The remote metro device domain name.
[DEFAULT] metro_san_address = None	(StrOpt) The remote metro device request url.
[DEFAULT] metro_san_password = None	(StrOpt) The remote metro device san password.
[DEFAULT] metro_san_user = None	(StrOpt) The remote metro device san user.
[DEFAULT] metro_storage_pools = None	(StrOpt) The remote metro device pool names.
[DEFAULT] nas_host =	(StrOpt) IP address or Hostname of NAS system.

Option = default value	(Type) Help string
[DEFAULT] netapp_replication_aggregate_map = None	(MultiOpt) Multi opt of dictionaries to represent the aggregate mapping between source and destination back ends when using whole back end replication. For every source aggregate associated with a cinder pool (NetApp FlexVol), you would need to specify the destination aggregate on the replication target device. A replication target device is configured with the configuration option replication_device. Specify this option as many times as you have replication devices. Each entry takes the standard dict config form: netapp_replication_aggregate_map = backend_id: <name_of_replication_device_section>,src_aggr_name1:dest_aggr_name1,src_aggr_name2:dest_aggr_name2,...
[DEFAULT] netapp_snapmirror_quiesce_timeout = 3600	(IntOpt) The maximum time in seconds to wait for existing SnapMirror transfers to complete before aborting during a failover.
[DEFAULT] nexenta_nbd_symlinks_dir = /dev/disk/by-path	(StrOpt) NexentaEdge logical path of directory to store symbolic links to NBDs
[DEFAULT] osapi_volume_use_ssl = False	(BoolOpt) Wraps the socket in a SSL context if True is set. A certificate file and key file must be specified.
[DEFAULT] pool_id_filter =	(ListOpt) Pool id permit to use.
[DEFAULT] pool_type = default	(StrOpt) Pool type, like sata-2copy.
[DEFAULT] proxy = storage.proxy.IBMStorageProxy	(StrOpt) Proxy driver that connects to the IBM Storage Array
[DEFAULT] quota_groups = 10	(IntOpt) Number of groups allowed per project
[DEFAULT] scaleio_server_certificate_path = None	(StrOpt) Server certificate path
[DEFAULT] scaleio_verify_server_certificate = False	(BoolOpt) verify server certificate
[DEFAULT] scheduler_weight_handler = cinder.scheduler.weights.OrderedHostWeightHandler	(StrOpt) Which handler to use for selecting the host/pool after weighing
[DEFAULT] secondary_san_ip =	(StrOpt) IP address of secondary DSM controller
[DEFAULT] secondary_san_login = Admin	(StrOpt) Secondary DSM user name
[DEFAULT] secondary_san_password =	(StrOpt) Secondary DSM user password name
[DEFAULT] secondary_sc_api_port = 3033	(PortOpt) Secondary Dell API port

Option = default value	(Type) Help string
[DEFAULT] sio_max_over_subscription_ratio = 10.0	(FloatOpt) max_over_subscription_ratio setting for the ScaleIO driver. This replaces the general max_over_subscription_ratio which has no effect in this driver. Maximum value allowed for ScaleIO is 10.0.
[DEFAULT] storage_protocol = iscsi	(StrOpt) Protocol for transferring data between host and storage back-end.
[DEFAULT] synology_admin_port = 5000	(PortOpt) Management port for Synology storage.
[DEFAULT] synology_device_id = None	(StrOpt) Device id for skip one time password check for logging in Synology storage if OTP is enabled.
[DEFAULT] synology_one_time_pass = None	(StrOpt) One time password of administrator for logging in Synology storage if OTP is enabled.
[DEFAULT] synology_password =	(StrOpt) Password of administrator for logging in Synology storage.
[DEFAULT] synology_pool_name =	(StrOpt) Volume on Synology storage to be used for creating lun.
[DEFAULT] synology_ssl_verify = True	(BoolOpt) Do certificate validation or not if \$driver_use_ssl is True
[DEFAULT] synology_username = admin	(StrOpt) Administrator of Synology storage.
[DEFAULT] violin_dedup_capable_pools =	(ListOpt) Storage pools capable of dedup and other luns.(Comma separated list)
[DEFAULT] violin_dedup_only_pools =	(ListOpt) Storage pools to be used to setup dedup luns only.(Comma separated list)
[DEFAULT] violin_iscsi_target_ips =	(ListOpt) Target iSCSI addresses to use.(Comma separated list)
[DEFAULT] violin_pool_allocation_method = random	(StrOpt) Method of choosing a storage pool for a lun.
[DEFAULT] vzstorage_default_volume_format = raw	(StrOpt) Default format that will be used when creating volumes if no volume format is specified.
[DEFAULT] zadara_default_snap_policy = False	(BoolOpt) VPSA - Attach snapshot policy for volumes
[DEFAULT] zadara_password = None	(StrOpt) VPSA - Password
[DEFAULT] zadara_use_iser = True	(BoolOpt) VPSA - Use ISER instead of iSCSI

Option = default value	(Type) Help string
[DEFAULT] zadara_user = None	(StrOpt) VPSA - Username
[DEFAULT] zadara_vol_encrypt = False	(BoolOpt) VPSA - Default encryption policy for volumes
[DEFAULT] zadara_vol_name_template = OS_%s	(StrOpt) VPSA - Default template for VPSA volume names
[DEFAULT] zadara_vpsa_host = None	(StrOpt) VPSA - Management Host name or IP address
[DEFAULT] zadara_vpsa_poolname = None	(StrOpt) VPSA - Storage Pool assigned for volumes
[DEFAULT] zadara_vpsa_port = None	(PortOpt) VPSA - Port number
[DEFAULT] zadara_vpsa_use_ssl = False	(BoolOpt) VPSA - Use SSL connection
[DEFAULT] zteAheadReadSize = 8	(IntOpt) Cache readahead size.
[DEFAULT] zteCachePolicy = 1	(IntOpt) Cache policy. 0, Write Back; 1, Write Through.
[DEFAULT] zteChunkSize = 4	(IntOpt) Virtual block size of pool. Unit : KB. Valid value : 4, 8, 16, 32, 64, 128, 256, 512.
[DEFAULT] zteControllerIP0 = None	(IPOpt) Main controller IP.
[DEFAULT] zteControllerIP1 = None	(IPOpt) Slave controller IP.
[DEFAULT] zteLocalIP = None	(IPOpt) Local IP.
[DEFAULT] ztePoolVoAllocatedPolicy = 0	(IntOpt) Pool volume allocated policy. 0, Auto; 1, High Performance Tier First; 2, Performance Tier First; 3, Capacity Tier First.
[DEFAULT] ztePoolVolAlarmStopAllocatedFlag = 0	(IntOpt) Pool volume alarm stop allocated flag.
[DEFAULT] ztePoolVolAlarmThreshold = 0	(IntOpt) Pool volume alarm threshold. [0, 100]
[DEFAULT] ztePoolVolInitAllocatedCapacity = 0	(IntOpt) Pool volume init allocated Capacity. Unit : KB.
[DEFAULT] ztePoolVollsThin = False	(IntOpt) Whether it is a thin volume.
[DEFAULT] ztePoolVolMovePolicy = 0	(IntOpt) Pool volume move policy. 0, Auto; 1, Highest Available; 2, Lowest Available; 3, No Relocation.

Option = default value	(Type) Help string
[DEFAULT] zteSSDCacheSwitch = 1	(IntOpt) SSD cache switch. 0, OFF; 1, ON.
[DEFAULT] zteStoragePool =	(ListOpt) Pool name list.
[DEFAULT] zteUserName = None	(StrOpt) User name.
[DEFAULT] zteUserPassword = None	(StrOpt) User password.
[barbican] auth_endpoint = http://localhost:5000/v3	(StrOpt) Use this endpoint to connect to Keystone
[barbican] barbican_api_version = None	(StrOpt) Version of the Barbican API, for example: "v1"
[barbican] barbican_endpoint = None	(StrOpt) Use this endpoint to connect to Barbican, for example: " http://localhost:9311/ "
[barbican] number_of_retries = 60	(IntOpt) Number of times to retry poll for key creation completion
[barbican] retry_delay = 1	(IntOpt) Number of seconds to wait before retrying poll for key creation completion
[fc-zone-manager] enable_unsupported_driver = False	(BoolOpt) Set this to True when you want to allow an unsupported zone manager driver to start. Drivers that haven't maintained a working CI system and testing are marked as unsupported until CI is working again. This also marks a driver as deprecated and may be removed in the next release.
[key_manager] api_class = castellan.key_manager.barbican_key_manager.BarbicanKeyManager	(StrOpt) The full class name of the key manager API class
[key_manager] fixed_key = None	(StrOpt) Fixed key returned by key manager, specified in hex

Table 3.68. New default values

Option	Previous default value	New default value
[DEFAULT] backup_service_inithost_offload	False	True
[DEFAULT] datera_num_replicas	1	3
[DEFAULT] default_timeout	525600	31536000

Option	Previous default value	New default value
[DEFAULT] glance_api_servers	\$glance_host:\$glance_port	None
[DEFAULT] io_port_list	*	None
[DEFAULT] iscsi_initiators		None
[DEFAULT] naviseccli_path		None
[DEFAULT] nexenta_chunksize	16384	32768
[DEFAULT] query_volume_filters	name, status, metadata, availability_zone, bootable	name, status, metadata, availability_zone, bootable, group_id
[DEFAULT] vmware_task_poll_interval	0.5	2.0

Table 3.69. Deprecated options

Deprecated option	New Option
[DEFAULT] enable_v1_api	None
[DEFAULT] enable_v2_api	None
[DEFAULT] eqix_chap_login	[DEFAULT] chap_username
[DEFAULT] eqix_chap_password	[DEFAULT] chap_password
[DEFAULT] eqix_use_chap	[DEFAULT] use_chap_auth
[DEFAULT] host	[DEFAULT] backend_host
[DEFAULT] nas_ip	[DEFAULT] nas_host
[DEFAULT] osapi_max_request_body_size	[oslo_middleware] max_request_body_size
[DEFAULT] use_syslog	None
[hyperv] force_volumeutils_v1	None

CHAPTER 4. IDENTITY SERVICE

This chapter details the OpenStack Identity service configuration options.

4.1. IDENTITY SERVICE CONFIGURATION OPTIONS

4.1.1. Identity Service Configuration File

The Identity service is configured in the `/etc/keystone/keystone.conf` file.

4.1.2. Description of Configuration Options

The following tables provide a comprehensive list of the Identity service options.

Table 4.1. Description of API configuration options

Configuration option = Default value	Description
[DEFAULT]	
admin_endpoint = None	(String) The base admin endpoint URL for Keystone that is advertised to clients (NOTE: this does NOT affect how Keystone listens for connections). Defaults to the base host URL of the request. For example, if keystone receives a request to <code>http://server:35357/v3/users</code> , then this will option will be automatically treated as <code>http://server:35357</code> . You should only need to set option if either the value of the base URL contains a path that keystone does not automatically infer (<code>/prefix/v3</code>), or if the endpoint should be found on a different host.
admin_token = None	(String) Using this feature is <i>NOT</i> recommended. Instead, use the <code>keystone-manage bootstrap</code> command. The value of this option is treated as a "shared secret" that can be used to bootstrap Keystone through the API. This "token" does not represent a user (it has no identity), and carries no explicit authorization (it effectively bypasses most authorization checks). If set to <i>None</i> , the value is ignored and the <code>admin_token</code> middleware is effectively disabled. However, to completely disable <code>admin_token</code> in production (highly recommended, as it presents a security risk), remove <code>AdminTokenAuthMiddleware</code> (the <code>admin_token_auth</code> filter) from your paste application pipelines (for example, in <code>keystone-paste.ini</code>).

Configuration option = Default value	Description
domain_id_immutable = True	(Boolean) DEPRECATED: Set this to false if you want to enable the ability for user, group and project entities to be moved between domains by updating their <i>domain_id</i> attribute. Allowing such movement is not recommended if the scope of a domain admin is being restricted by use of an appropriate policy file (see <i>etc/policy.v3cloudsample.json</i> as an example). This feature is deprecated and will be removed in a future release, in favor of strictly immutable domain IDs. The option to set <i>domain_id_immutable</i> to false has been deprecated in the M release and will be removed in the O release.
list_limit = None	(Integer) The maximum number of entities that will be returned in a collection. This global limit may be then overridden for a specific driver, by specifying a <i>list_limit</i> in the appropriate section (for example, <i>[assignment]</i>). No limit is set by default. In larger deployments, it is recommended that you set this to a reasonable number to prevent operations like listing all users and projects from placing an unnecessary load on the system.
max_param_size = 64	(Integer) Limit the sizes of user & project ID/names.
max_project_tree_depth = 5	(Integer) Maximum depth of the project hierarchy, excluding the project acting as a domain at the top of the hierarchy. WARNING: Setting it to a large value may adversely impact performance.
max_token_size = 8192	(Integer) Similar to <i>[DEFAULT] max_param_size</i> , but provides an exception for token values. With PKI / PKIZ tokens, this needs to be set close to 8192 (any higher, and other HTTP implementations may break), depending on the size of your service catalog and other factors. With Fernet tokens, this can be set as low as 255. With UUID tokens, this should be set to 32).
member_role_id = 9fe2ff9ee4384b1894a90878d3e92bab	(String) Similar to the <i>[DEFAULT] member_role_name</i> option, this represents the default role ID used to associate users with their default projects in the v2 API. This will be used as the explicit role where one is not specified by the v2 API. You do not need to set this value unless you want keystone to use an existing role with a different ID, other than the arbitrarily defined <i>_member_role</i> (in which case, you should set <i>[DEFAULT] member_role_name</i> as well).

Configuration option = Default value	Description
member_role_name = <code>_member_</code>	(String) This is the role name used in combination with the <code>[DEFAULT] member_role_id</code> option; see that option for more detail. You do not need to set this option unless you want keystone to use an existing role (in which case, you should set <code>[DEFAULT] member_role_id</code> as well).
public_endpoint = <code>None</code>	(String) The base public endpoint URL for Keystone that is advertised to clients (NOTE: this does NOT affect how Keystone listens for connections). Defaults to the base host URL of the request. For example, if keystone receives a request to <code>http://server:5000/v3/users</code> , then this option will be automatically treated as <code>http://server:5000</code> . You should only need to set option if either the value of the base URL contains a path that keystone does not automatically infer (<code>/prefix/v3</code>), or if the endpoint should be found on a different host.
secure_proxy_ssl_header = <code>HTTP_X_FORWARDED_PROTO</code>	(String) DEPRECATED: The HTTP header used to determine the scheme for the original request, even if it was removed by an SSL terminating proxy. This option has been deprecated in the N release and will be removed in the P release. Use <code>oslo.middleware.http_proxy_to_wsgi</code> configuration instead.
strict_password_check = <code>False</code>	(Boolean) If set to true, strict password length checking is performed for password manipulation. If a password exceeds the maximum length, the operation will fail with an HTTP 403 Forbidden error. If set to false, passwords are automatically truncated to the maximum length.
<code>[oslo_middleware]</code>	
enable_proxy_headers_parsing = <code>False</code>	(Boolean) Whether the application is behind a proxy or not. This determines if the middleware should parse the headers or not.
max_request_body_size = <code>114688</code>	(Integer) The maximum body size for each request, in bytes.
secure_proxy_ssl_header = <code>X-Forwarded-Proto</code>	(String) DEPRECATED: The HTTP Header that will be used to determine what the original request protocol scheme was, even if it was hidden by a SSL termination proxy.

Table 4.2. Description of assignment configuration options

Configuration option = Default value	Description
[assignment]	
driver = None	(String) Entrypoint for the assignment backend driver in the keystone.assignment namespace. Only an SQL driver is supplied. If an assignment driver is not specified, the identity driver will choose the assignment driver (driver selection based on <i>[identity]/driver</i> option is deprecated and will be removed in the "O" release).
prohibited_implied_role = admin	(List) A list of role names which are prohibited from being an implied role.

Table 4.3. Description of authorization configuration options

Configuration option = Default value	Description
[auth]	
external = None	(String) Entrypoint for the external (REMOTE_USER) auth plugin module in the keystone.auth.external namespace. Supplied drivers are DefaultDomain and Domain. The default driver is DefaultDomain.
methods = external, password, token, oauth1	(List) Allowed authentication methods.
oauth1 = None	(String) Entrypoint for the oAuth1.0 auth plugin module in the keystone.auth.oauth1 namespace.
password = None	(String) Entrypoint for the password auth plugin module in the keystone.auth.password namespace.
token = None	(String) Entrypoint for the token auth plugin module in the keystone.auth.token namespace.

Table 4.4. Description of cache configuration options

Configuration option = Default value	Description
[memcache]	

Configuration option = Default value	Description
dead_retry = 300	(Integer) 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).
pool_connection_get_timeout = 10	(Integer) 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).
pool_maxsize = 10	(Integer) 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).
pool_unused_timeout = 60	(Integer) 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).

Table 4.5. Description of catalog configuration options

Configuration option = Default value	Description
[catalog]	
cache_time = None	(Integer) Time to cache catalog data (in seconds). This has no effect unless global and catalog caching are enabled.
caching = True	(Boolean) Toggle for catalog caching. This has no effect unless global caching is enabled.
driver = sql	(String) Entrypoint for the catalog backend driver in the keystone.catalog namespace. Supplied drivers are kvs, sql, templated, and endpoint_filter.sql
list_limit = None	(Integer) Maximum number of entities that will be returned in a catalog collection.
template_file = default_catalog.templates	(String) Catalog template file name for use with the template catalog backend.

Table 4.6. Description of CA and SSL configuration options

Configuration option = Default value	Description
[eventlet_server_ssl]	
ca_certs = /etc/keystone/ssl/certs/ca.pem	(String) DEPRECATED: Path of the CA cert file for SSL.
cert_required = False	(Boolean) DEPRECATED: Require client certificate.
certfile = /etc/keystone/ssl/certs/keystone.pem	(String) DEPRECATED: Path of the certfile for SSL. For non-production environments, you may be interested in using <i>keystone-manage ssl_setup</i> to generate self-signed certificates.
enable = False	(Boolean) DEPRECATED: Toggle for SSL support on the Keystone eventlet servers.
keyfile = /etc/keystone/ssl/private/keystonekey.pem	(String) DEPRECATED: Path of the keyfile for SSL.
[signing]	
ca_certs = /etc/keystone/ssl/certs/ca.pem	(String) DEPRECATED: Path of the CA for token signing. PKI token support has been deprecated in the M release and will be removed in the O release. Fernet or UUID tokens are recommended.
ca_key = /etc/keystone/ssl/private/cakey.pem	(String) DEPRECATED: Path of the CA key for token signing. PKI token support has been deprecated in the M release and will be removed in the O release. Fernet or UUID tokens are recommended.
cert_subject = /C=US/ST=Unset/L=Unset/O=Unset/CN=www.example.com	(String) DEPRECATED: Certificate subject (auto generated certificate) for token signing. PKI token support has been deprecated in the M release and will be removed in the O release. Fernet or UUID tokens are recommended.
certfile = /etc/keystone/ssl/certs/signing_cert.pem	(String) DEPRECATED: Path of the certfile for token signing. For non-production environments, you may be interested in using <i>keystone-manage pki_setup</i> to generate self-signed certificates. PKI token support has been deprecated in the M release and will be removed in the O release. Fernet or UUID tokens are recommended.
key_size = 2048	(Integer) DEPRECATED: Key size (in bits) for token signing cert (auto generated certificate). PKI token support has been deprecated in the M release and will be removed in the O release. Fernet or UUID tokens are recommended.

Configuration option = Default value	Description
keyfile = /etc/keystone/ssl/private/signing_key.pem	(String) DEPRECATED: Path of the keyfile for token signing. PKI token support has been deprecated in the M release and will be removed in the O release. Fernet or UUID tokens are recommended.
valid_days = 3650	(Integer) DEPRECATED: Days the token signing cert is valid for (auto generated certificate). PKI token support has been deprecated in the M release and will be removed in the O release. Fernet or UUID tokens are recommended.
[ssl]	
ca_key = /etc/keystone/ssl/private/cakey.pem	(String) Path of the CA key file for SSL.
cert_subject = /C=US/ST=Unset/L=Unset/O=Unset/CN=local host	(String) SSL certificate subject (auto generated certificate).
key_size = 1024	(Integer) SSL key length (in bits) (auto generated certificate).
valid_days = 3650	(Integer) Days the certificate is valid for once signed (auto generated certificate).

Table 4.7. Description of common configuration options

Configuration option = Default value	Description
[DEFAULT]	
executor_thread_pool_size = 64	(Integer) Size of executor thread pool.
insecure_debug = False	(Boolean) If set to true, then the server will return information in HTTP responses that may allow an unauthenticated or authenticated user to get more information than normal, such as additional details about why authentication failed. This may be useful for debugging but is insecure.

Table 4.8. Description of credential configuration options

Configuration option = Default value	Description
[credential]	
driver = sql	(String) Entrypoint for the credential backend driver in the keystone.credential namespace.

Table 4.9. Description of logging configuration options

Configuration option = Default value	Description
[audit]	
namespace = openstack	(String) namespace prefix for generated id

Table 4.10. Description of domain configuration options

Configuration option = Default value	Description
[domain_config]	
cache_time = 300	(Integer) TTL (in seconds) to cache domain config data. This has no effect unless domain config caching is enabled.
caching = True	(Boolean) Toggle for domain config caching. This has no effect unless global caching is enabled.
driver = sql	(String) Entrypoint for the domain config backend driver in the keystone.resource.domain_config namespace.

Table 4.11. Description of federation configuration options

Configuration option = Default value	Description
[federation]	
assertion_prefix =	(String) Value to be used when filtering assertion parameters from the environment.

Configuration option = Default value	Description
driver = sql	(String) Entrypoint for the federation backend driver in the keystone.federation namespace.
federated_domain_name = Federated	(String) 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.
remote_id_attribute = None	(String) 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 <i>Shib-Identity-Provider</i>).
sso_callback_template = /etc/keystone/sso_callback_template.html	(String) Location of Single Sign-On callback handler, will return a token to a trusted dashboard host.
trusted_dashboard = []	(Multi-valued) A list of trusted dashboard hosts. Before accepting a Single Sign-On request to return a token, the origin host must be a member of the trusted_dashboard list. This configuration option may be repeated for multiple values. For example: trusted_dashboard=http://acme.com/auth/websso trusted_dashboard=http://beta.com/auth/websso

Table 4.12. Description of Fernet tokens configuration options

Configuration option = Default value	Description
[fernet_tokens]	
key_repository = /etc/keystone/fernet-keys/	(String) Directory containing Fernet token keys.
max_active_keys = 3	(Integer) 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.

Table 4.13. Description of identity configuration options

Configuration option = Default value	Description
[identity]	
cache_time = 600	(Integer) Time to cache identity data (in seconds). This has no effect unless global and identity caching are enabled.
caching = True	(Boolean) Toggle for identity caching. This has no effect unless global caching is enabled.
default_domain_id = default	(String) This references the domain to use for all Identity API v2 requests (which are not aware of domains). A domain with this ID will be created for you by keystone-manage db_sync in migration 008. The domain referenced by this ID cannot be deleted on the v3 API, to prevent accidentally breaking the v2 API. There is nothing special about this domain, other than the fact that it must exist in order to maintain support for your v2 clients.
domain_config_dir = /etc/keystone/domains	(String) Path for Keystone to locate the domain specific identity configuration files if domain_specific_drivers_enabled is set to true.
domain_configurations_from_database = False	(Boolean) 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.
domain_specific_drivers_enabled = False	(Boolean) A subset (or all) of domains can have their own identity driver, each with their own partial configuration options, stored in either the resource backend or in a file in a domain configuration directory (depending on the setting of domain_configurations_from_database). Only values specific to the domain need to be specified in this manner. This feature is disabled by default; set to true to enable.
driver = sql	(String) Entrypoint for the identity backend driver in the keystone.identity namespace. Supplied drivers are ldap and sql.
list_limit = None	(Integer) Maximum number of entities that will be returned in an identity collection.
max_password_length = 4096	(Integer) Maximum supported length for user passwords; decrease to improve performance.

Table 4.14. Description of KVS configuration options

Configuration option = Default value	Description
[kvs]	
backends =	(List) Extra dogpile.cache backend modules to register with the dogpile.cache library.
config_prefix = keystone.kvs	(String) Prefix for building the configuration dictionary for the KVS region. This should not need to be changed unless there is another dogpile.cache region with the same configuration name.
default_lock_timeout = 5	(Integer) Default lock timeout (in seconds) for distributed locking.
enable_key_mangler = True	(Boolean) Toggle to disable using a key-mangling function to ensure fixed length keys. This is toggle-able for debugging purposes, it is highly recommended to always leave this set to true.

Table 4.15. Description of LDAP configuration options

Configuration option = Default value	Description
[ldap]	
alias_dereferencing = default	(String) The LDAP dereferencing option for queries. The "default" option falls back to using default dereferencing configured by your ldap.conf.
allow_subtree_delete = False	(Boolean) Delete subtrees using the subtree delete control. Only enable this option if your LDAP server supports subtree deletion.
auth_pool_connection_lifetime = 60	(Integer) End user auth connection lifetime in seconds.
auth_pool_size = 100	(Integer) End user auth connection pool size.
chase_referrals = None	(Boolean) Override the system's default referral chasing behavior for queries.
debug_level = None	(Integer) 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.
dumb_member = cn=dumb,dc=nonexistent	(String) DN of the "dummy member" to use when "use_dumb_member" is enabled.

Configuration option = Default value	Description
group_additional_attribute_mapping =	(List) 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.
group_allow_create = True	(Boolean) DEPRECATED: Allow group creation in LDAP backend. Write support for Identity LDAP backends has been deprecated in the M release and will be removed in the O release.
group_allow_delete = True	(Boolean) DEPRECATED: Allow group deletion in LDAP backend. Write support for Identity LDAP backends has been deprecated in the M release and will be removed in the O release.
group_allow_update = True	(Boolean) DEPRECATED: Allow group update in LDAP backend. Write support for Identity LDAP backends has been deprecated in the M release and will be removed in the O release.
group_attribute_ignore =	(List) List of attributes stripped off the group on update.
group_desc_attribute = description	(String) LDAP attribute mapped to group description.
group_filter = None	(String) LDAP search filter for groups.
group_id_attribute = cn	(String) LDAP attribute mapped to group id.
group_member_attribute = member	(String) LDAP attribute mapped to show group membership.
group_members_are_ids = False	(Boolean) If the members of the group objectclass are user IDs rather than DNs, set this to true. This is the case when using posixGroup as the group objectclass and OpenDirectory.
group_name_attribute = ou	(String) LDAP attribute mapped to group name.
group_objectclass = groupOfNames	(String) LDAP objectclass for groups.
group_tree_dn = None	(String) Search base for groups. Defaults to the suffix value.
page_size = 0	(Integer) Maximum results per page; a value of zero ("0") disables paging.
password = None	(String) Password for the BindDN to query the LDAP server.

Configuration option = Default value	Description
pool_connection_lifetime = 600	(Integer) Connection lifetime in seconds.
pool_connection_timeout = -1	(Integer) Connector timeout in seconds. Value -1 indicates indefinite wait for response.
pool_retry_delay = 0.1	(Floating point) Time span in seconds to wait between two reconnect trials.
pool_retry_max = 3	(Integer) Maximum count of reconnect trials.
pool_size = 10	(Integer) Connection pool size.
query_scope = one	(String) The LDAP scope for queries, "one" represents oneLevel/singleLevel and "sub" represents subtree/wholeSubtree options.
suffix = cn=example,cn=com	(String) LDAP server suffix
tls_cacertdir = None	(String) CA certificate directory path for communicating with LDAP servers.
tls_cacertfile = None	(String) CA certificate file path for communicating with LDAP servers.
tls_req_cert = demand	(String) Specifies what checks to perform on client certificates in an incoming TLS session.
url = ldap://localhost	(String) URL(s) for connecting to the LDAP server. Multiple LDAP URLs may be specified as a comma separated string. The first URL to successfully bind is used for the connection.
use_auth_pool = True	(Boolean) Enable LDAP connection pooling for end user authentication. If use_pool is disabled, then this setting is meaningless and is not used at all.
use_dumb_member = False	(Boolean) If true, will add a dummy member to groups. This is required if the objectclass for groups requires the "member" attribute.
use_pool = True	(Boolean) Enable LDAP connection pooling.
use_tls = False	(Boolean) Enable TLS for communicating with LDAP servers.
user = None	(String) User BindDN to query the LDAP server.

Configuration option = Default value	Description
user_additional_attribute_mapping =	(List) List of additional LDAP attributes used for mapping additional attribute mappings for users. 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.
user_allow_create = True	(Boolean) DEPRECATED: Allow user creation in LDAP backend. Write support for Identity LDAP backends has been deprecated in the M release and will be removed in the O release.
user_allow_delete = True	(Boolean) DEPRECATED: Allow user deletion in LDAP backend. Write support for Identity LDAP backends has been deprecated in the M release and will be removed in the O release.
user_allow_update = True	(Boolean) DEPRECATED: Allow user updates in LDAP backend. Write support for Identity LDAP backends has been deprecated in the M release and will be removed in the O release.
user_attribute_ignore = default_project_id	(List) List of attributes stripped off the user on update.
user_default_project_id_attribute = None	(String) LDAP attribute mapped to default_project_id for users.
user_description_attribute = description	(String) LDAP attribute mapped to user description.
user_enabled_attribute = enabled	(String) LDAP attribute mapped to user enabled flag.
user_enabled_default = True	(String) 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".
user_enabled_emulation = False	(Boolean) If true, Keystone uses an alternative method to determine if a user is enabled or not by checking if they are a member of the "user_enabled_emulation_dn" group.
user_enabled_emulation_dn = None	(String) DN of the group entry to hold enabled users when using enabled emulation.
user_enabled_emulation_use_group_config = False	(Boolean) Use the "group_member_attribute" and "group_objectclass" settings to determine membership in the emulated enabled group.

Configuration option = Default value	Description
user_enabled_invert = False	(Boolean) 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.
user_enabled_mask = 0	(Integer) 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".
user_filter = None	(String) LDAP search filter for users.
user_id_attribute = cn	(String) LDAP attribute mapped to user id. WARNING: must not be a multivalued attribute.
user_mail_attribute = mail	(String) LDAP attribute mapped to user email.
user_name_attribute = sn	(String) LDAP attribute mapped to user name.
user_objectclass = inetOrgPerson	(String) LDAP objectclass for users.
user_pass_attribute = userPassword	(String) LDAP attribute mapped to password.
user_tree_dn = None	(String) Search base for users. Defaults to the suffix value.

Table 4.16. Description of mapping configuration options

Configuration option = Default value	Description
[identity_mapping]	

Configuration option = Default value	Description
backward_compatible_ids = True	(Boolean) 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.
driver = sql	(String) Entrypoint for the identity mapping backend driver in the keystone.identity.id_mapping namespace.
generator = sha256	(String) Entrypoint for the public ID generator for user and group entities in the keystone.identity.id_generator namespace. The Keystone identity mapper only supports generators that produce no more than 64 characters.

Table 4.17. Description of memcache configuration options

Configuration option = Default value	Description
[memcache]	
servers = localhost:11211	(List) Memcache servers in the format of "host:port".
socket_timeout = 3	(Integer) 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).

Table 4.18. Description of OAuth configuration options

Configuration option = Default value	Description
[oauth1]	
access_token_duration = 86400	(Integer) Duration (in seconds) for the OAuth Access Token.
driver = sql	(String) Entrypoint for the OAuth backend driver in the keystone.oauth1 namespace.
request_token_duration = 28800	(Integer) Duration (in seconds) for the OAuth Request Token.

Table 4.19. Description of os_inherit configuration options

Configuration option = Default value	Description
[os_inherit]	
enabled = True	(Boolean) DEPRECATED: role-assignment inheritance to projects from owning domain or from projects higher in the hierarchy can be optionally disabled. In the future, this option will be removed and the hierarchy will be always enabled. The option to enable the OS-INHERIT extension has been deprecated in the M release and will be removed in the O release. The OS-INHERIT extension will be enabled by default.

Table 4.20. Description of policy configuration options

Configuration option = Default value	Description
[policy]	
driver = sql	(String) Entrypoint for the policy backend driver in the keystone.policy namespace. Supplied drivers are rules and sql.
list_limit = None	(Integer) Maximum number of entities that will be returned in a policy collection.

Table 4.21. Description of Redis configuration options

Configuration option = Default value	Description
[matchmaker_redis]	
check_timeout = 20000	(Integer) Time in ms to wait before the transaction is killed.
host = 127.0.0.1	(String) DEPRECATED: Host to locate redis. Replaced by [DEFAULT]/transport_url
password =	(String) DEPRECATED: Password for Redis server (optional). Replaced by [DEFAULT]/transport_url
port = 6379	(Port number) DEPRECATED: Use this port to connect to redis host. Replaced by [DEFAULT]/transport_url
sentinel_hosts =	(List) DEPRECATED: List of Redis Sentinel hosts (fault tolerance mode) e.g. [host:port, host1:port ...] Replaced by [DEFAULT]/transport_url
socket_timeout = 10000	(Integer) Timeout in ms on blocking socket operations
wait_timeout = 2000	(Integer) Time in ms to wait between connection attempts.

Table 4.22. Description of revoke configuration options

Configuration option = Default value	Description
[revoke]	
cache_time = 3600	(Integer) Time to cache the revocation list and the revocation events (in seconds). This has no effect unless global and token caching are enabled.
caching = True	(Boolean) Toggle for revocation event caching. This has no effect unless global caching is enabled.
driver = sql	(String) Entrypoint for an implementation of the backend for persisting revocation events in the keystone.revoke namespace. Supplied drivers are kvs and sql.
expiration_buffer = 1800	(Integer) This value (calculated in seconds) is added to token expiration before a revocation event may be removed from the backend.

Table 4.23. Description of role configuration options

Configuration option = Default value	Description
[role]	
cache_time = None	(Integer) TTL (in seconds) to cache role data. This has no effect unless global caching is enabled.
caching = True	(Boolean) Toggle for role caching. This has no effect unless global caching is enabled.
driver = None	(String) Entrypoint for the role backend driver in the keystone.role namespace. Supplied drivers are ldap and sql.
list_limit = None	(Integer) Maximum number of entities that will be returned in a role collection.

Table 4.24. Description of SAML configuration options

Configuration option = Default value	Description
[saml]	
assertion_expiration_time = 3600	(Integer) Default TTL, in seconds, for any generated SAML assertion created by Keystone.
certfile = /etc/keystone/ssl/certs/signing_cert.pem	(String) Path of the certfile for SAML signing. For non-production environments, you may be interested in using <i>keystone-manage pki_setup</i> to generate self-signed certificates. Note, the path cannot contain a comma.
idp_contact_company = None	(String) Company of contact person.
idp_contact_email = None	(String) Email address of contact person.
idp_contact_name = None	(String) Given name of contact person
idp_contact_surname = None	(String) Surname of contact person.
idp_contact_telephone = None	(String) Telephone number of contact person.
idp_contact_type = other	(String) The contact type describing the main point of contact for the identity provider.

Configuration option = Default value	Description
idp_entity_id = None	(String) Entity ID value for unique Identity Provider identification. Usually FQDN is set with a suffix. A value is required to generate IDP Metadata. For example: https://keystone.example.com/v3/OS-FEDERATION/saml2/idp
idp_lang = en	(String) Language used by the organization.
idp_metadata_path = /etc/keystone/saml2_idp_metadata.xml	(String) Path to the Identity Provider Metadata file. This file should be generated with the keystone-manage saml_idp_metadata command.
idp_organization_display_name = None	(String) Organization name to be displayed.
idp_organization_name = None	(String) Organization name the installation belongs to.
idp_organization_url = None	(String) URL of the organization.
idp_sso_endpoint = None	(String) Identity Provider Single-Sign-On service value, required in the Identity Provider's metadata. A value is required to generate IDP Metadata. For example: https://keystone.example.com/v3/OS-FEDERATION/saml2/sso
keyfile = /etc/keystone/ssl/private/signing_key.pem	(String) Path of the keyfile for SAML signing. Note, the path cannot contain a comma.
relay_state_prefix = ss:mem:	(String) The prefix to use for the RelayState SAML attribute, used when generating ECP wrapped assertions.
xmlsec1_binary = xmlsec1	(String) Binary to be called for XML signing. Install the appropriate package, specify absolute path or adjust your PATH environment variable if the binary cannot be found.

Table 4.25. Description of security configuration options

Configuration option = Default value	Description
[DEFAULT]	

Configuration option = Default value	Description
crypt_strength = 10000	(Integer) The value passed as the keyword "rounds" to passlib's encrypt method. This option represents a trade off between security and performance. Higher values lead to slower performance, but higher security. Changing this option will only affect newly created passwords as existing password hashes already have a fixed number of rounds applied, so it is safe to tune this option in a running cluster. For more information, see https://pythonhosted.org/passlib/password_hash_api.html#choosing-the-right-rounds-value

Table 4.26. Description of Tokenless Authorization configuration options

Configuration option = Default value	Description
[tokenless_auth]	
issuer_attribute = SSL_CLIENT_I_DN	(String) The issuer attribute that is served as an IdP ID for the X.509 tokenless authorization along with the protocol to look up its corresponding mapping. It is the environment variable in the WSGI environment that references to the issuer of the client certificate.
protocol = x509	(String) The protocol name for the X.509 tokenless authorization along with the option issuer_attribute below can look up its corresponding mapping.
trusted_issuer = []	(Multi-valued) The list of trusted issuers to further filter the certificates that are allowed to participate in the X.509 tokenless authorization. If the option is absent then no certificates will be allowed. The naming format for the attributes of a Distinguished Name(DN) must be separated by a comma and contain no spaces. This configuration option may be repeated for multiple values. For example: trusted_issuer=CN=john,OU=keystone,O=openstack trusted_issuer=CN=mary,OU=eng,O=abc

Table 4.27. Description of token configuration options

Configuration option = Default value	Description
[token]	

Configuration option = Default value	Description
allow_rescope_scoped_token = True	(Boolean) Allow rescopeing of scoped token. Setting <code>allow_rescope_scoped_token</code> to false prevents a user from exchanging a scoped token for any other token.
bind =	(List) External auth mechanisms that should add bind information to token, e.g., <code>kerberos,x509</code> .
cache_time = None	(Integer) Time to cache tokens (in seconds). This has no effect unless global and token caching are enabled.
caching = True	(Boolean) Toggle for token system caching. This has no effect unless global caching is enabled.
driver = sql	(String) Entrypoint for the token persistence backend driver in the <code>keystone.token.persistence</code> namespace. Supplied drivers are <code>kvs</code> , <code>memcache</code> , <code>memcache_pool</code> , and <code>sql</code> .
enforce_token_bind = permissive	(String) Enforcement policy on tokens presented to Keystone with bind information. One of <code>disabled</code> , <code>permissive</code> , <code>strict</code> , <code>required</code> or a specifically required bind mode, e.g., <code>kerberos</code> or <code>x509</code> to require binding to that authentication.
expiration = 3600	(Integer) Amount of time a token should remain valid (in seconds).
hash_algorithm = md5	(String) DEPRECATED: The hash algorithm to use for PKI tokens. This can be set to any algorithm that <code>hashlib</code> supports. WARNING: Before changing this value, the <code>auth_token</code> middleware must be configured with the <code>hash_algorithms</code> , otherwise token revocation will not be processed correctly. PKI token support has been deprecated in the M release and will be removed in the O release. Fernet or UUID tokens are recommended.
infer_roles = True	(Boolean) Add roles to token that are not explicitly added, but that are linked implicitly to other roles.
provider = uuid	(String) Controls the token construction, validation, and revocation operations. Entrypoint in the <code>keystone.token.provider</code> namespace. Core providers are <code>[fernet pkiz pki uuid]</code> .

Configuration option = Default value	Description
revoke_by_id = True	(Boolean) Revoke token by token identifier. Setting <code>revoke_by_id</code> to <code>true</code> enables various forms of enumerating tokens, e.g. <i>list tokens for user</i> . These enumerations are processed to determine the list of tokens to revoke. Only disable if you are switching to using the Revoke extension with a backend other than KVS, which stores events in memory.

Table 4.28. Description of trust configuration options

Configuration option = Default value	Description
[trust]	
allow_redelegation = False	(Boolean) Enable redelegation feature.
driver = sql	(String) Entrypoint for the trust backend driver in the <code>keystone.trust</code> namespace.
enabled = True	(Boolean) Delegation and impersonation features can be optionally disabled.
max_redelegation_count = 3	(Integer) Maximum depth of trust redelegation.

4.1.3. New, Updated, and Deprecated Options in Newton for Identity Service

Table 4.29. Deprecated options

Deprecated option	New Option
[DEFAULT] use_syslog	None

CHAPTER 5. NETWORKING

This chapter explains the OpenStack Networking configuration options.

5.1. NETWORKING CONFIGURATION OPTIONS

5.1.1. Description of Configuration Options

The following tables provide a comprehensive list of the Networking configuration options.

Table 5.1. Description of agent configuration options

Configuration option = Default value	Description
[DEFAULT]	
external_pids = \$state_path/external/pids	(String) Location to store child pid files
[AGENT]	
agent_type = Open vSwitch agent	(String) DEPRECATED: Selects the Agent Type reported
availability_zone = nova	(String) Availability zone of this node

Table 5.2. Description of API configuration options

Configuration option = Default value	Description
[DEFAULT]	
allow_bulk = True	(Boolean) Allow the usage of the bulk API
allow_pagination = True	(Boolean) DEPRECATED: Allow the usage of the pagination. This option has been deprecated and will now be enabled unconditionally.
allow_sorting = True	(Boolean) DEPRECATED: Allow the usage of the sorting. This option has been deprecated and will now be enabled unconditionally.

Configuration option = Default value	Description
api_extensions_path =	(String) The path for API extensions. Note that this can be a colon-separated list of paths. For example: <code>api_extensions_path = extensions:/path/to/more/exts:/even/more/exts</code> . The <code>__path__</code> of <code>neutron.extensions</code> is appended to this, so if your extensions are in there you don't need to specify them here.
api_paste_config = api-paste.ini	(String) File name for the <code>paste.deploy</code> config for api service
backlog = 4096	(Integer) Number of backlog requests to configure the socket with
client_socket_timeout = 900	(Integer) 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.
max_header_line = 16384	(Integer) Maximum line size of message headers to be accepted. <code>max_header_line</code> may need to be increased when using large tokens (typically those generated when keystone is configured to use PKI tokens with big service catalogs).
pagination_max_limit = -1	(String) The maximum number of items returned in a single response, value was 'infinite' or negative integer means no limit
retry_until_window = 30	(Integer) Number of seconds to keep retrying to listen
service_plugins =	(List) The service plugins Neutron will use
tcp_keepidle = 600	(Integer) Sets the value of <code>TCP_KEEPIDLE</code> in seconds for each server socket. Not supported on OS X.
use_ssl = False	(Boolean) Enable SSL on the API server
wsgi_default_pool_size = 100	(Integer) Size of the pool of greenthreads used by wsgi
wsgi_keep_alive = True	(Boolean) If False, closes the client socket connection explicitly.
wsgi_log_format = %(client_ip)s "%(request_line)s" status: %(status_code)s len: %(body_length)s time: %(wall_seconds).7f	(String) A python format string that is used as the template to generate log lines. The following values can be formatted into it: <code>client_ip</code> , <code>date_time</code> , <code>request_line</code> , <code>status_code</code> , <code>body_length</code> , <code>wall_seconds</code> .

Configuration option = Default value	Description
[oslo_middleware]	
enable_proxy_headers_parsing = False	(Boolean) Whether the application is behind a proxy or not. This determines if the middleware should parse the headers or not.
max_request_body_size = 114688	(Integer) The maximum body size for each request, in bytes.
secure_proxy_ssl_header = X-Forwarded-Proto	(String) DEPRECATED: The HTTP Header that will be used to determine what the original request protocol scheme was, even if it was hidden by a SSL termination proxy.
[oslo_versionedobjects]	
fatal_exception_format_errors = False	(Boolean) Make exception message format errors fatal

Table 5.3. Description of common configuration options

Configuration option = Default value	Description
[DEFAULT]	
agent_down_time = 75	(Integer) Seconds to regard the agent is down; should be at least twice report_interval, to be sure the agent is down for good.
allow_automatic_dhcp_failover = True	(Boolean) Automatically remove networks from offline DHCP agents.
allow_automatic_l3agent_failover = False	(Boolean) Automatically reschedule routers from offline L3 agents to online L3 agents.
api_workers = None	(Integer) Number of separate API worker processes for service. If not specified, the default is equal to the number of CPUs available for best performance.
auth_ca_cert = None	(String) Certificate Authority public key (CA cert) file for ssl
auth_strategy = keystone	(String) The type of authentication to use

Configuration option = Default value	Description
base_mac = fa:16:3e:00:00:00	(String) The base MAC address Neutron will use for VIFs. The first 3 octets will remain unchanged. If the 4th octet is not 00, it will also be used. The others will be randomly generated.
bind_host = 0.0.0.0	(String) The host IP to bind to
bind_port = 9696	(Port number) The port to bind to
cache_url =	(String) DEPRECATED: URL to connect to the cache back end. This option is deprecated in the Newton release and will be removed. Please add a [cache] group for oslo.cache in your neutron.conf and add "enable" and "backend" options in this section.
core_plugin = None	(String) The core plugin Neutron will use
default_availability_zones =	(List) Default value of availability zone hints. The availability zone aware schedulers use this when the resources availability_zone_hints is empty. Multiple availability zones can be specified by a comma separated string. This value can be empty. In this case, even if availability_zone_hints for a resource is empty, availability zone is considered for high availability while scheduling the resource.
dhcp_agent_notification = True	(Boolean) Allow sending resource operation notification to DHCP agent
dhcp_agents_per_network = 1	(Integer) Number of DHCP agents scheduled to host a tenant network. If this number is greater than 1, the scheduler automatically assigns multiple DHCP agents for a given tenant network, providing high availability for DHCP service.
dhcp_broadcast_reply = False	(Boolean) Use broadcast in DHCP replies.
dhcp_confs = \$state_path/dhcp	(String) Location to store DHCP server config files.
dhcp_domain = openstacklocal	(String) DEPRECATED: Domain to use for building the hostnames. This option is deprecated. It has been moved to neutron.conf as dns_domain. It will be removed in a future release.
dhcp_lease_duration = 86400	(Integer) DHCP lease duration (in seconds). Use -1 to tell dnsmasq to use infinite lease times.

Configuration option = Default value	Description
dhcp_load_type = networks	(String) 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
dns_domain = openstacklocal	(String) Domain to use for building the hostnames
enable_new_agents = True	(Boolean) Agent starts with admin_state_up=False when enable_new_agents=False. In the case, user's resources will not be scheduled automatically to the agent until admin changes admin_state_up to True.
enable_services_on_agents_with_admin_state_down = False	(Boolean) 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.
executor_thread_pool_size = 64	(Integer) Size of executor thread pool.
external_dns_driver = None	(String) Driver for external DNS integration.
global_physnet_mtu = 1500	(Integer) MTU of the underlying physical network. Neutron uses this value to calculate MTU for all virtual network components. For flat and VLAN networks, neutron uses this value without modification. For overlay networks such as VXLAN, neutron automatically subtracts the overlay protocol overhead from this value. Defaults to 1500, the standard value for Ethernet.
ip_lib_force_root = False	(Boolean) Force ip_lib calls to use the root helper
ipam_driver = internal	(String) Neutron IPAM (IP address management) driver to use. By default, the reference implementation of the Neutron IPAM driver is used.
mac_generation_retries = 16	(Integer) DEPRECATED: How many times Neutron will retry MAC generation. This option is now obsolete and so is deprecated to be removed in the Ocata release.

Configuration option = Default value	Description
max_allowed_address_pair = 10	(Integer) Maximum number of allowed address pairs
max_dns_nameservers = 5	(Integer) Maximum number of DNS nameservers per subnet
max_fixed_ips_per_port = 5	(Integer) DEPRECATED: Maximum number of fixed ips per port. This option is deprecated and will be removed in the Ocata release.
max_rtr_adv_interval = 100	(Integer) MaxRtrAdvInterval setting for radvd.conf
max_subnet_host_routes = 20	(Integer) Maximum number of host routes per subnet
min_rtr_adv_interval = 30	(Integer) MinRtrAdvInterval setting for radvd.conf
periodic_fuzzy_delay = 5	(Integer) Range of seconds to randomly delay when starting the periodic task scheduler to reduce stampeding. (Disable by setting to 0)
periodic_interval = 40	(Integer) Seconds between running periodic tasks.
report_interval = 300	(Integer) Interval between two metering reports
state_path = /var/lib/neutron	(String) Where to store Neutron state files. This directory must be writable by the agent.
vlan_transparent = False	(Boolean) If True, then allow plugins that support it to create VLAN transparent networks.
web_framework = legacy	(String) This will choose the web framework in which to run the Neutron API server. 'pecan' is a new experimental rewrite of the API server.
[AGENT]	
check_child_processes_action = respawn	(String) Action to be executed when a child process dies
check_child_processes_interval = 60	(Integer) Interval between checks of child process liveness (seconds), use 0 to disable
debug_iptables_rules = False	(Boolean) Duplicate every iptables difference calculation to ensure the format being generated matches the format of iptables-save. This option should not be turned on for production systems because it imposes a performance penalty.
log_agent_heartbeats = False	(Boolean) Log agent heartbeats

Configuration option = Default value	Description
polling_interval = 2	(Integer) The number of seconds the agent will wait between polling for local device changes.
root_helper = sudo	(String) Root helper application. 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_daemon = None	(String) Root helper daemon application to use when possible.
[profiler]	
connection_string = messaging://	<p>(String) Connection string for a notifier backend. Default value is messaging:// which sets the notifier to oslo_messaging.</p> <p>Examples of possible values:</p> <ul style="list-style-type: none"> ● messaging://: use oslo_messaging driver for sending notifications.
enabled = False	<p>(Boolean) Enables the profiling for all services on this node. Default value is False (fully disable the profiling feature).</p> <p>Possible values:</p> <ul style="list-style-type: none"> ● True: Enables the feature ● False: Disables the feature. The profiling cannot be started via this project operations. If the profiling is triggered by another project, this project part will be empty.
hmac_keys = SECRET_KEY	<p>(String) Secret key(s) to use for encrypting context data for performance profiling. This string value should have the following format: <key1>[,<key2>,...<keyn>], where each key is some random string. A user who triggers the profiling via the REST API has to set one of these keys in the headers of the REST API call to include profiling results of this node for this particular project.</p> <p>Both "enabled" flag and "hmac_keys" config options should be set to enable profiling. Also, to generate correct profiling information across all services at least one key needs to be consistent between OpenStack projects. This ensures it can be used from client side to generate the trace, containing information from all possible resources.</p>

Configuration option = Default value	Description
trace_sqlalchemy = False	(Boolean) Enables SQL requests profiling in services. Default value is False (SQL requests won't be traced). Possible values: <ul style="list-style-type: none"> • True: Enables SQL requests profiling. Each SQL query will be part of the trace and can be analyzed by how much time was spent for that. • False: Disables SQL requests profiling. The spent time is only shown on a higher level of operations. Single SQL queries cannot be analyzed this way.
[qos]	
notification_drivers = message_queue	(List) Drivers list to use to send the update notification
[service_providers]	
service_provider = []	(Multi-valued) Defines providers for advanced services using the format: <service_type><name>:<driver>[:default]

Table 5.4. Description of Compute configuration options

Configuration option = Default value	Description
[DEFAULT]	
notify_nova_on_port_data_changes = True	(Boolean) Send notification to nova when port data (fixed_ips/floatingip) changes so nova can update its cache.
notify_nova_on_port_status_changes = True	(Boolean) Send notification to nova when port status changes
nova_client_cert =	(String) Client certificate for nova metadata api server.
nova_client_priv_key =	(String) Private key of client certificate.
send_events_interval = 2	(Integer) Number of seconds between sending events to nova if there are any events to send.

Table 5.5. Description of DHCP agent configuration options

Configuration option = Default value	Description
[DEFAULT]	
advertise_mtu = True	(Boolean) DEPRECATED: If True, advertise network MTU values if core plugin calculates them. MTU is advertised to running instances via DHCP and RA MTU options.
dhcp_driver = neutron.agent.linux.dhcp.Dnsmasq	(String) The driver used to manage the DHCP server.
dnsmasq_base_log_dir = None	(String) Base log dir for dnsmasq logging. The log contains DHCP and DNS log information and is useful for debugging issues with either DHCP or DNS. If this section is null, disable dnsmasq log.
dnsmasq_config_file =	(String) Override the default dnsmasq settings with this file.
dnsmasq_dns_servers =	(List) Comma-separated list of the DNS servers which will be used as forwarders.
dnsmasq_lease_max = 16777216	(Integer) Limit number of leases to prevent a denial-of-service.
dnsmasq_local_resolv = False	(Boolean) Enables the dnsmasq service to provide name resolution for instances via DNS resolvers on the host running the DHCP agent. Effectively removes the '--no-resolv' option from the dnsmasq process arguments. Adding custom DNS resolvers to the 'dnsmasq_dns_servers' option disables this feature.
enable_isolated_metadata = False	(Boolean) The DHCP server can assist with providing metadata support on isolated networks. Setting this value to True will cause the DHCP server to append specific host routes to the DHCP request. The metadata service will only be activated when the subnet does not contain any router port. The guest instance must be configured to request host routes via DHCP (Option 121). This option doesn't have any effect when force_metadata is set to True.
enable_metadata_network = False	(Boolean) Allows for serving metadata requests coming from a dedicated metadata access network whose CIDR is 169.254.169.254/16 (or larger prefix), and is connected to a Neutron router from which the VMs send metadata:1 request. In this case DHCP Option 121 will not be injected in VMs, as they will be able to reach 169.254.169.254 through a router. This option requires enable_isolated_metadata = True.

Configuration option = Default value	Description
force_metadata = False	(Boolean) In some cases the Neutron router is not present to provide the metadata IP but the DHCP server can be used to provide this info. Setting this value will force the DHCP server to append specific host routes to the DHCP request. If this option is set, then the metadata service will be activated for all the networks.
host = example.domain	(String) 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.
interface_driver = None	(String) The driver used to manage the virtual interface.
num_sync_threads = 4	(Integer) Number of threads to use during sync process. Should not exceed connection pool size configured on server.
resync_interval = 5	(Integer) The DHCP agent will resync its state with Neutron to recover from any transient notification or RPC errors. The interval is number of seconds between attempts.

Table 5.6. Description of DVR configuration options

Configuration option = Default value	Description
[DEFAULT]	
dvr_base_mac = fa:16:3f:00:00:00	(String) The base mac address used for unique DVR instances by Neutron. The first 3 octets will remain unchanged. If the 4th octet is not 00, it will also be used. The others will be randomly generated. The 'dvr_base_mac' <i>must</i> 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
router_distributed = False	(Boolean) System-wide flag to determine the type of router that tenants can create. Only admin can override.

Table 5.7. Description of FDB agent configuration options

Configuration option = Default value	Description
[FDB]	
shared_physical_device_mappings =	(List) Comma-separated list of <physical_network>: <network_device> tuples mapping physical network names to the agent's node-specific shared physical network device between SR-IOV and OVS or SR-IOV and linux bridge

Table 5.8. Description of IPv6 router advertisement configuration options

Configuration option = Default value	Description
[DEFAULT]	
ra_confs = \$state_path/ra	(String) Location to store IPv6 RA config files

Table 5.9. Description of L2 agent extension configuration options

Configuration option = Default value	Description
[agent]	
extensions =	(List) Extensions list to use

Table 5.10. Description of L3 agent configuration options

Configuration option = Default value	Description
[DEFAULT]	
enable_snat_by_default = True	(Boolean) Define the default value of enable_snat if not provided in external_gateway_info.

Configuration option = Default value	Description
external_network_bridge =	(String) DEPRECATED: Name of bridge used for external network traffic. When this parameter is set, the L3 agent will plug an interface directly into an external bridge which will not allow any wiring by the L2 agent. Using this will result in incorrect port statuses. This option is deprecated and will be removed in Ocata.
ha_confs_path = \$state_path/ha_confs	(String) Location to store keepalived/contrackd config files
ha_vrrp_advert_int = 2	(Integer) The advertisement interval in seconds
ha_vrrp_auth_password = None	(String) VRRP authentication password
ha_vrrp_auth_type = PASS	(String) VRRP authentication type
host = example.domain	(String) 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.
interface_driver = None	(String) The driver used to manage the virtual interface.
ipv6_pd_enabled = False	(Boolean) Enables IPv6 Prefix Delegation for automatic subnet CIDR allocation. Set to True to enable IPv6 Prefix Delegation for subnet allocation in a PD-capable environment. Users making subnet creation requests for IPv6 subnets without providing a CIDR or subnetpool ID will be given a CIDR via the Prefix Delegation mechanism. Note that enabling PD will override the behavior of the default IPv6 subnetpool.
l3_ha = False	(Boolean) Enable HA mode for virtual routers.
l3_ha_net_cidr = 169.254.192.0/18	(String) Subnet used for the l3 HA admin network.
l3_ha_network_physical_name =	(String) The physical network name with which the HA network can be created.
l3_ha_network_type =	(String) The network type to use when creating the HA network for an HA router. By default or if empty, the first 'tenant_network_types' is used. This is helpful when the VRRP traffic should use a specific network which is not the default one.
max_l3_agents_per_router = 3	(Integer) Maximum number of L3 agents which a HA router will be scheduled on. If it is set to 0 then the router will be scheduled on every agent.

Configuration option = Default value	Description
min_l3_agents_per_router = 2	(Integer) DEPRECATED: Minimum number of L3 agents that have to be available in order to allow a new HA router to be scheduled. This option is deprecated in the Newton release and will be removed for the Ocata release where the scheduling of new HA routers will always be allowed.
[AGENT]	
comment_iptables_rules = True	(Boolean) Add comments to iptables rules. Set to false to disallow the addition of comments to generated iptables rules that describe each rule's purpose. System must support the iptables comments module for addition of comments.
use_helper_for_ns_read = True	(Boolean) 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.

Table 5.11. Description of LBaaS agent configuration options

Configuration option = Default value	Description
[DEFAULT]	
debug = False	(Boolean) If set to true, the logging level will be set to DEBUG instead of the default INFO level. Mutable This option can be changed without restarting.
device_driver = ['neutron_lbaas.drivers.haproxy.namespace_driver.HaproxyNSDriver']	(Multi-valued) Drivers used to manage loadbalancing devices
interface_driver = None	(String) The driver used to manage the virtual interface.
periodic_interval = 40	(Integer) Seconds between running periodic tasks.
[haproxy]	
loadbalancer_state_path = \$state_path/lbaas	(String) Location to store config and state files

Configuration option = Default value	Description
send_gratuitous_arp = 3	(Integer) When delete and re-add the same vip, send this many gratuitous ARPs to flush the ARP cache in the Router. Set it below or equal to 0 to disable this feature.
user_group = nogroup	(String) The user group

Table 5.12. Description of LBaaS Embrane, Radware, NetScaler, HAproxy, plug-in configuration options

Configuration option = Default value	Description
[DEFAULT]	
loadbalancer_scheduler_driver = neutron_lbaas.agent_scheduler.ChanceScheduler	(String) Driver to use for scheduling to a default loadbalancer agent
[haproxy]	
jinja_config_template = /usr/lib/python/site-packages/neutron-lbaas/neutron_lbaas/drivers/haproxy/templates/haproxy.loadbalancer.j2	(String) Jinja template file for haproxy configuration
[radwarev2]	
child_workflow_template_names = manage_l3	(List) Name of child workflow templates used. Default: manage_l3
ha_secondary_address = None	(String) IP address of secondary vDirect server.
service_adc_type = VA	(String) Service ADC type. Default: VA.
service_adc_version =	(String) Service ADC version.
service_cache = 20	(Integer) Size of service cache. Default: 20.
service_compression_throughput = 100	(Integer) Service compression throughput. Default: 100.
service_ha_pair = False	(Boolean) Enables or disables the Service HA pair. Default: False.
service_isl_vlan = -1	(Integer) A required VLAN for the interswitch link to use.

Configuration option = Default value	Description
service_resource_pool_ids =	(List) Resource pool IDs.
service_session_mirroring_enabled = False	(Boolean) Enable or disable Alteon interswitch link for stateful session failover. Default: False.
service_ssl_throughput = 100	(Integer) Service SSL throughput. Default: 100.
service_throughput = 1000	(Integer) Service throughput. Default: 1000.
stats_action_name = stats	(String) Name of the workflow action for statistics. Default: stats.
vdirect_address = None	(String) IP address of vDirect server.
vdirect_password = radware	(String) vDirect user password.
vdirect_user = vDirect	(String) vDirect user name.
workflow_action_name = apply	(String) Name of the workflow action. Default: apply.
workflow_params = {'data_ip_address': '192.168.200.99', 'ha_network_name': 'HA-Network', 'ha_port': 2, 'allocate_ha_ips': True, 'ha_ip_pool_name': 'default', 'allocate_ha_vrrp': True, 'data_port': 1, 'gateway': '192.168.200.1', 'twoleg_enabled': '_REPLACE_', 'data_ip_mask': '255.255.255.0'}	(Dict) Parameter for I2_I3 workflow constructor.
workflow_template_name = os_lb_v2	(String) Name of the workflow template. Default: os_lb_v2.
[radwarev2_debug]	
configure_l3 = True	(Boolean) Configure ADC with L3 parameters?
configure_l4 = True	(Boolean) Configure ADC with L4 parameters?
provision_service = True	(Boolean) Provision ADC service?

Table 5.13. Description of Load-Balancer-as-a-Service configuration options

Configuration option = Default value	Description
[certificates]	

Configuration option = Default value	Description
barbican_auth = barbican_acl_auth	(String) Name of the Barbican authentication method to use
cert_manager_type = barbican	(String) Certificate Manager plugin. Defaults to barbican.
storage_path = /var/lib/neutron-lbaas/certificates/	(String) Absolute path to the certificate storage directory. Defaults to env[OS_LBAAS_TLS_STORAGE].

Table 5.14. Description of metadata configuration options

Configuration option = Default value	Description
[DEFAULT]	
metadata_backlog = 4096	(Integer) Number of backlog requests to configure the metadata server socket with
metadata_proxy_group =	(String) Group (gid or name) running metadata proxy after its initialization (if empty: agent effective group).
metadata_proxy_shared_secret =	(String) When proxying metadata requests, Neutron signs the Instance-ID header with a shared secret to prevent spoofing. You may select any string for a secret, but it must match here and in the configuration used by the Nova Metadata Server. NOTE: Nova uses the same config key, but in [neutron] section.
metadata_proxy_socket = \$state_path/metadata_proxy	(String) Location of Metadata Proxy UNIX domain socket
metadata_proxy_socket_mode = deduce	(String) Metadata Proxy UNIX domain socket mode, 4 values allowed: 'deduce': deduce mode from metadata_proxy_user/group values, 'user': set metadata proxy socket mode to 0o644, to use when metadata_proxy_user is agent effective user or root, 'group': set metadata proxy socket mode to 0o664, to use when metadata_proxy_group is agent effective group or root, 'all': set metadata proxy socket mode to 0o666, to use otherwise.
metadata_proxy_user =	(String) User (uid or name) running metadata proxy after its initialization (if empty: agent effective user).

Configuration option = Default value	Description
metadata_proxy_watch_log = None	(Boolean) Enable/Disable log watch by metadata proxy. It should be disabled when metadata_proxy_user/group is not allowed to read/write its log file and copytruncate logrotate option must be used if logrotate is enabled on metadata proxy log files. Option default value is deduced from metadata_proxy_user: watch log is enabled if metadata_proxy_user is agent effective user id/name.
metadata_workers = 0	(Integer) Number of separate worker processes for metadata server (defaults to half of the number of CPUs)
nova_metadata_insecure = False	(Boolean) Allow to perform insecure SSL (https) requests to nova metadata
nova_metadata_ip = 127.0.0.1	(String) IP address used by Nova metadata server.
nova_metadata_port = 8775	(Port number) TCP Port used by Nova metadata server.
nova_metadata_protocol = http	(String) Protocol to access nova metadata, http or https

Table 5.15. Description of ML2 Flat mechanism driver configuration options

Configuration option = Default value	Description
[ml2_type_flat]	
flat_networks = *	(List) List of physical_network names with which flat networks can be created. Use default '*' to allow flat networks with arbitrary physical_network names. Use an empty list to disable flat networks.

Table 5.16. Description of ML2 Geneve type driver configuration options

Configuration option = Default value	Description
[ml2_type_geneve]	

Configuration option = Default value	Description
max_header_size = 30	(Integer) Geneve encapsulation header size is dynamic, this value is used to calculate the maximum MTU for the driver. This is the sum of the sizes of the outer ETH + IP + UDP + GENEVE header sizes. The default size for this field is 50, which is the size of the Geneve header without any additional option headers.
vni_ranges =	(List) Comma-separated list of <vni_min>:<vni_max> tuples enumerating ranges of Geneve VNI IDs that are available for tenant network allocation

Table 5.17. Description of ML2 GRE configuration options

Configuration option = Default value	Description
[ml2_type_gre]	
tunnel_id_ranges =	(List) Comma-separated list of <tun_min>:<tun_max> tuples enumerating ranges of GRE tunnel IDs that are available for tenant network allocation

Table 5.18. Description of ML2 L2 population configuration options

Configuration option = Default value	Description
[l2pop]	
agent_boot_time = 180	(Integer) Delay within which agent is expected to update existing ports when it restarts

Table 5.19. Description of ML2 ML2 SR-IOV driver configuration options

Configuration option = Default value	Description
[ml2_sriov]	

Configuration option = Default value	Description
supported_pci_vendor_devs = None	(List) DEPRECATED: Comma-separated list of supported PCI vendor devices, as defined by vendor_id:product_id according to the PCI ID Repository. Default None accept all PCI vendor devices. DEPRECATED: This option is deprecated in the Newton release and will be removed in the Ocata release. Starting from Ocata the mechanism driver will accept all PCI vendor devices.

Table 5.20. Description of ML2 VLAN configuration options

Configuration option = Default value	Description
[ml2_type_vlan]	
network_vlan_ranges =	(List) 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.

Table 5.21. Description of ML2 VXLAN configuration options

Configuration option = Default value	Description
[ml2_type_vxlan]	
vni_ranges =	(List) 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	(String) Multicast group for VXLAN. When configured, will enable sending all broadcast traffic to this multicast group. When left unconfigured, will disable multicast VXLAN mode.

Table 5.22. Description of ML2 configuration options

Configuration option = Default value	Description
[ml2]	
extension_drivers =	(List) An ordered list of extension driver endpoints to be loaded from the neutron.ml2.extension_drivers namespace. For example: extension_drivers = port_security,qos
external_network_type = None	(String) Default network type for external networks when no provider attributes are specified. By default it is None, which means that if provider attributes are not specified while creating external networks then they will have the same type as tenant networks. Allowed values for external_network_type config option depend on the network type values configured in type_drivers config option.
mechanism_drivers =	(List) An ordered list of networking mechanism driver endpoints to be loaded from the neutron.ml2.mechanism_drivers namespace.
overlay_ip_version = 4	(Integer) IP version of all overlay (tunnel) network endpoints. Use a value of 4 for IPv4 or 6 for IPv6.
path_mtu = 0	(Integer) Maximum size of an IP packet (MTU) that can traverse the underlying physical network infrastructure without fragmentation when using an overlay/tunnel protocol. This option allows specifying a physical network MTU value that differs from the default global_physnet_mtu value.
physical_network_mtus =	(List) A list of mappings of physical networks to MTU values. The format of the mapping is <physnet>:<mtu val>. This mapping allows specifying a physical network MTU value that differs from the default global_physnet_mtu value.
tenant_network_types = local	(List) Ordered list of network_types to allocate as tenant networks. The default value 'local' is useful for single-box testing but provides no connectivity between hosts.
type_drivers = local, flat, vlan, gre, vxlan, geneve	(List) List of network type driver endpoints to be loaded from the neutron.ml2.type_drivers namespace.

Table 5.23. Description of nova configuration options

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

Configuration option = Default value	Description
[nova]	
auth_section = None	(Unknown) Config Section from which to load plugin specific options
auth_type = None	(Unknown) Authentication type to load
cafile = None	(String) PEM encoded Certificate Authority to use when verifying HTTPs connections.
certfile = None	(String) PEM encoded client certificate cert file
endpoint_type = public	(String) Type of the nova endpoint to use. This endpoint will be looked up in the keystone catalog and should be one of public, internal or admin.
insecure = False	(Boolean) Verify HTTPS connections.
keyfile = None	(String) PEM encoded client certificate key file
region_name = None	(String) Name of nova region to use. Useful if keystone manages more than one region.
timeout = None	(Integer) Timeout value for http requests

Table 5.24. Description of Open vSwitch agent configuration options

Configuration option = Default value	Description
[DEFAULT]	
ovs_integration_bridge = br-int	(String) Name of Open vSwitch bridge to use
ovs_use_veth = False	(Boolean) Uses 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_vsctl_timeout = 10	(Integer) Timeout in seconds for ovs-vsctl commands. If the timeout expires, ovs commands will fail with ALARMCLOCK error.
[AGENT]	

Configuration option = Default value	Description
arp_responder = False	(Boolean) Enable local ARP responder if it is supported. Requires OVS 2.1 and ML2 I2population driver. Allows the switch (when supporting an overlay) to respond to an ARP request locally without performing a costly ARP broadcast into the overlay.
dont_fragment = True	(Boolean) Set or un-set the don't fragment (DF) bit on outgoing IP packet carrying GRE/VXLAN tunnel.
drop_flows_on_start = False	(Boolean) Reset flow table on start. Setting this to True will cause brief traffic interruption.
enable_distributed_routing = False	(Boolean) Make the I2 agent run in DVR mode.
I2_population = False	(Boolean) Use ML2 I2population mechanism driver to learn remote MAC and IPs and improve tunnel scalability.
minimize_polling = True	(Boolean) Minimize polling by monitoring ovssdb for interface changes.
ovssdb_monitor_respawn_interval = 30	(Integer) The number of seconds to wait before respawning the ovssdb monitor after losing communication with it.
prevent_arp_spoofing = True	(Boolean) DEPRECATED: Enable suppression of ARP responses that don't match an IP address that belongs to the port from which they originate. Note: This prevents the VMs attached to this agent from spoofing, it doesn't protect them from other devices which have the capability to spoof (e.g. bare metal or VMs attached to agents without this flag set to True). Spoofing rules will not be added to any ports that have port security disabled. For LinuxBridge, this requires ebtables. For OVS, it requires a version that supports matching ARP headers. This option will be removed in Ocata so the only way to disable protection will be via the port security extension.
quitting_rpc_timeout = 10	(Integer) Set new timeout in seconds for new rpc calls after agent receives SIGTERM. If value is set to 0, rpc timeout won't be changed
tunnel_csum = False	(Boolean) Set or un-set the tunnel header checksum on outgoing IP packet carrying GRE/VXLAN tunnel.
tunnel_types =	(List) Network types supported by the agent (gre and/or vxlan).
veth_mtu = 9000	(Integer) MTU size of veth interfaces
vxlan_udp_port = 4789	(Port number) The UDP port to use for VXLAN tunnels.

Configuration option = Default value	Description
[OVS]	
bridge_mappings =	(List) Comma-separated list of <physical_network>: <bridge> tuples mapping physical network names to the agent's node-specific Open vSwitch bridge names to be used for flat and VLAN networks. The length of bridge names should be no more than 11. Each bridge must exist, and should have a physical network interface configured as a port. All physical networks configured on the server should have mappings to appropriate bridges on each agent. Note: If you remove a bridge from this mapping, make sure to disconnect it from the integration bridge as it won't be managed by the agent anymore.
datapath_type = system	(String) OVS datapath to use. 'system' is the default value and corresponds to the kernel datapath. To enable the userspace datapath set this value to 'netdev'.
int_peer_patch_port = patch-tun	(String) Peer patch port in integration bridge for tunnel bridge.
integration_bridge = br-int	(String) Integration bridge to use. Do not change this parameter unless you have a good reason to. This is the name of the OVS integration bridge. There is one per hypervisor. The integration bridge acts as a virtual 'patch bay'. All VM VIFs are attached to this bridge and then 'patched' according to their network connectivity.
local_ip = None	(IP) IP address of local overlay (tunnel) network endpoint. Use either an IPv4 or IPv6 address that resides on one of the host network interfaces. The IP version of this value must match the value of the 'overlay_ip_version' option in the ML2 plug-in configuration file on the neutron server node(s).
of_connect_timeout = 30	(Integer) Timeout in seconds to wait for the local switch connecting the controller. Used only for 'native' driver.
of_interface = native	(String) OpenFlow interface to use.
of_listen_address = 127.0.0.1	(IP) Address to listen on for OpenFlow connections. Used only for 'native' driver.
of_listen_port = 6633	(Port number) Port to listen on for OpenFlow connections. Used only for 'native' driver.
of_request_timeout = 10	(Integer) Timeout in seconds to wait for a single OpenFlow request. Used only for 'native' driver.

Configuration option = Default value	Description
ovsdb_connection = tcp:127.0.0.1:6640	(String) The connection string for the native OVSDb backend. Requires the native ovsdb_interface to be enabled.
ovsdb_interface = native	(String) The interface for interacting with the OVSDb
tun_peer_patch_port = patch-int	(String) Peer patch port in tunnel bridge for integration bridge.
tunnel_bridge = br-tun	(String) Tunnel bridge to use.
use_veth_interconnection = False	(Boolean) Use veths instead of patch ports to interconnect the integration bridge to physical networks. Support kernel without Open vSwitch patch port support so long as it is set to True.
vhostuser_socket_dir = /var/run/openvswitch	(String) OVS vhost-user socket directory.

Table 5.25. Description of IPv6 Prefix Delegation driver configuration options

Configuration option = Default value	Description
[DEFAULT]	
pd_confs = \$state_path/pd	(String) Location to store IPv6 PD files.
pd_dhcp_driver = dibbler	(String) Service to handle DHCPv6 Prefix delegation.
vendor_pen = 8888	(String) A decimal value as Vendor's Registered Private Enterprise Number as required by RFC3315 DUID-EN.

Table 5.26. Description of policy configuration options

Configuration option = Default value	Description
[DEFAULT]	
allow_overlapping_ips = False	(Boolean) Allow overlapping IP support in Neutron. Attention: the following parameter MUST be set to False if Neutron is being used in conjunction with Nova security groups.

Table 5.27. Description of QoS configuration options

Configuration option = Default value	Description
[QOS]	
kernel_hz = 250	(Integer) Value of host kernel tick rate (hz) for calculating minimum burst value in bandwidth limit rules for a port with QoS. See kernel configuration file for HZ value and tc-tbf manual for more information.
tbf_latency = 50	(Integer) Value of latency (ms) for calculating size of queue for a port with QoS. See tc-tbf manual for more information.

Table 5.28. Description of quotas configuration options

Configuration option = Default value	Description
[DEFAULT]	
max_routes = 30	(Integer) Maximum number of routes per router
[QUOTAS]	
default_quota = -1	(Integer) Default number of resource allowed per tenant. A negative value means unlimited.
quota_driver = neutron.db.quota.driver.DbQuotaDriver	(String) Default driver to use for quota checks.
quota_firewall = 10	(Integer) Number of firewalls allowed per tenant. A negative value means unlimited.
quota_firewall_policy = 10	(Integer) Number of firewall policies allowed per tenant. A negative value means unlimited.
quota_firewall_rule = 100	(Integer) Number of firewall rules allowed per tenant. A negative value means unlimited.
quota_floatingip = 50	(Integer) Number of floating IPs allowed per tenant. A negative value means unlimited.
quota_healthmonitor = -1	(Integer) Number of health monitors allowed per tenant. A negative value means unlimited.
quota_listener = -1	(Integer) Number of Loadbalancer Listeners allowed per tenant. A negative value means unlimited.

Configuration option = Default value	Description
quota_loadbalancer = 10	(Integer) Number of LoadBalancers allowed per tenant. A negative value means unlimited.
quota_member = -1	(Integer) Number of pool members allowed per tenant. A negative value means unlimited.
quota_network = 10	(Integer) Number of networks allowed per tenant. A negative value means unlimited.
quota_pool = 10	(Integer) Number of pools allowed per tenant. A negative value means unlimited.
quota_port = 50	(Integer) Number of ports allowed per tenant. A negative value means unlimited.
quota_rbac_policy = 10	(Integer) Default number of RBAC entries allowed per tenant. A negative value means unlimited.
quota_router = 10	(Integer) Number of routers allowed per tenant. A negative value means unlimited.
quota_security_group = 10	(Integer) Number of security groups allowed per tenant. A negative value means unlimited.
quota_security_group_rule = 100	(Integer) Number of security rules allowed per tenant. A negative value means unlimited.
quota_subnet = 10	(Integer) Number of subnets allowed per tenant, A negative value means unlimited.
track_quota_usage = True	(Boolean) Keep in track in the database of current resource quota usage. Plugins which do not leverage the neutron database should set this flag to False.

Table 5.29. Description of Redis configuration options

Configuration option = Default value	Description
[matchmaker_redis]	
check_timeout = 20000	(Integer) Time in ms to wait before the transaction is killed.
host = 127.0.0.1	(String) DEPRECATED: Host to locate redis. Replaced by [DEFAULT]/transport_url

Configuration option = Default value	Description
password =	(String) DEPRECATED: Password for Redis server (optional). Replaced by [DEFAULT]/transport_url
port = 6379	(Port number) DEPRECATED: Use this port to connect to redis host. Replaced by [DEFAULT]/transport_url
sentinel_hosts =	(List) DEPRECATED: List of Redis Sentinel hosts (fault tolerance mode) e.g. [host:port, host1:port ...] Replaced by [DEFAULT]/transport_url
socket_timeout = 10000	(Integer) Timeout in ms on blocking socket operations
wait_timeout = 2000	(Integer) Time in ms to wait between connection attempts.

Table 5.30. Description of scheduler configuration options

Configuration option = Default value	Description
[DEFAULT]	
network_auto_schedule = True	(Boolean) Allow auto scheduling networks to DHCP agent.
network_scheduler_driver = neutron.scheduler.dhcp_agent_scheduler.WeightScheduler	(String) Driver to use for scheduling network to DHCP agent
router_auto_schedule = True	(Boolean) Allow auto scheduling of routers to L3 agent.
router_scheduler_driver = neutron.scheduler.l3_agent_scheduler.LeastRoutersScheduler	(String) Driver to use for scheduling router to a default L3 agent

Table 5.31. Description of security groups configuration options

Configuration option = Default value	Description
[SECURITYGROUP]	

Configuration option = Default value	Description
enable_ipset = True	(Boolean) Use ipset to speed-up the iptables based security groups. Enabling ipset support requires that ipset is installed on L2 agent node.
enable_security_group = True	(Boolean) 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.
firewall_driver = None	(String) Driver for security groups firewall in the L2 agent

Table 5.32. Description of LBaaS service authentication configuration options

Configuration option = Default value	Description
[service_auth]	
admin_password = password	(String) The service admin password
admin_project_domain = admin	(String) The admin project domain name
admin_tenant_name = admin	(String) The service admin tenant name
admin_user = admin	(String) The service admin user name
admin_user_domain = admin	(String) The admin user domain name
auth_url = http://127.0.0.1:5000/v2.0	(String) Authentication endpoint
auth_version = 2	(String) The auth version used to authenticate
endpoint_type = public	(String) The endpoint_type to be used
insecure = False	(Boolean) Disable server certificate verification
region = RegionOne	(String) The deployment region
service_name = lbaas	(String) The name of the service

Table 5.33. Description of SR-IOV agent configuration options

Configuration option = Default value	Description
[SRIOV_NIC]	
exclude_devices =	(List) Comma-separated list of <network_device>: <vfs_to_exclude> tuples, mapping network_device to the agent's node-specific list of virtual functions that should not be used for virtual networking. vfs_to_exclude is a semicolon-separated list of virtual functions to exclude from network_device. The network_device in the mapping should appear in the physical_device_mappings list.
physical_device_mappings =	(List) Comma-separated list of <physical_network>: <network_device> tuples mapping physical network names to the agent's node-specific physical network device interfaces of SR-IOV physical function to be used for VLAN networks. All physical networks listed in network_vlan_ranges on the server should have mappings to appropriate interfaces on each agent.

5.1.2. New, Updated, and Deprecated Options in Newton for Networking

Table 5.34. New options

Option = default value	(Type) Help string
[DEFAULT] cache_url =	(StrOpt) URL to connect to the cache back end. This option is deprecated in the Newton release and will be removed. Please add a [cache] group for oslo.cache in your neutron.conf and add "enable" and "backend" options in this section.
[AGENT] debug_iptables_rules = False	(BoolOpt) Duplicate every iptables difference calculation to ensure the format being generated matches the format of iptables-save. This option should not be turned on for production systems because it imposes a performance penalty.
[FDB] shared_physical_device_mappings =	(ListOpt) Comma-separated list of <physical_network>:<network_device> tuples mapping physical network names to the agent's node-specific shared physical network device between SR-IOV and OVS or SR-IOV and linux bridge

Option = default value	(Type) Help string
[cache] backend = dogpile.cache.null	(StrOpt) Dogpile.cache backend module. It is recommended that Memcache or Redis (dogpile.cache.redis) be used in production deployments. For eventlet-based or highly threaded servers, Memcache with pooling (oslo_cache.memcache_pool) is recommended. For low thread servers, dogpile.cache.memcached is recommended. Test environments with a single instance of the server can use the dogpile.cache.memory backend.
[cache] backend_argument = []	(MultiStrOpt) Arguments supplied to the backend module. Specify this option once per argument to be passed to the dogpile.cache backend. Example format: "<argname><value>".
[cache] config_prefix = cache.oslo	(StrOpt) 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.
[cache] debug_cache_backend = False	(BoolOpt) 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.
[cache] enabled = False	(BoolOpt) Global toggle for caching.
[cache] expiration_time = 600	(IntOpt) 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.
[cache] memcache_dead_retry = 300	(IntOpt) Number of seconds memcached server is considered dead before it is tried again. (dogpile.cache.memcache and oslo_cache.memcache_pool backends only).
[cache] memcache_pool_connection_get_timeout = 10	(IntOpt) Number of seconds that an operation will wait to get a memcache client connection.
[cache] memcache_pool_maxsize = 10	(IntOpt) Max total number of open connections to every memcached server. (oslo_cache.memcache_pool backend only).
[cache] memcache_pool_unused_timeout = 60	(IntOpt) Number of seconds a connection to memcached is held unused in the pool before it is closed. (oslo_cache.memcache_pool backend only).
[cache] memcache_servers = localhost:11211	(ListOpt) Memcache servers in the format of "host:port". (dogpile.cache.memcache and oslo_cache.memcache_pool backends only).

Option = default value	(Type) Help string
[cache] memcache_socket_timeout = 3	(IntOpt) Timeout in seconds for every call to a server. (dogpile.cache.memcache and oslo_cache.memcache_pool backends only).
[cache] proxies =	(ListOpt) Proxy classes to import that will affect the way the dogpile.cache backend functions. See the dogpile.cache documentation on changing-backend-behavior.
[ml2] overlay_ip_version = 4	(IntOpt) IP version of all overlay (tunnel) network endpoints. Use a value of 4 for IPv4 or 6 for IPv6.
[profiler] connection_string = messaging://	(StrOpt) Connection string for a notifier backend. Default value is messaging:// which sets the notifier to oslo_messaging. Examples of possible values: * messaging://: use oslo_messaging driver for sending notifications.
[profiler] enabled = False	(BoolOpt) Enables the profiling for all services on this node. Default value is False (fully disable the profiling feature). Possible values: * True: Enables the feature * False: Disables the feature. The profiling cannot be started via this project operations. If the profiling is triggered by another project, this project part will be empty.
[profiler] hmac_keys = SECRET_KEY	(StrOpt) Secret key(s) to use for encrypting context data for performance profiling. This string value should have the following format: <key1>[,<key2>,...<keyn>], where each key is some random string. A user who triggers the profiling via the REST API has to set one of these keys in the headers of the REST API call to include profiling results of this node for this particular project. Both "enabled" flag and "hmac_keys" config options should be set to enable profiling. Also, to generate correct profiling information across all services at least one key needs to be consistent between OpenStack projects. This ensures it can be used from client side to generate the trace, containing information from all possible resources.
[profiler] trace_sqlalchemy = False	(BoolOpt) Enables SQL requests profiling in services. Default value is False (SQL requests won't be traced). Possible values: * True: Enables SQL requests profiling. Each SQL query will be part of the trace and can be analyzed by how much time was spent for that. * False: Disables SQL requests profiling. The spent time is only shown on a higher level of operations. Single SQL queries cannot be analyzed this way.

Table 5.35. New default values

Option	Previous default value	New default value
[DEFAULT] allow_pagination	False	True
[DEFAULT] allow_sorting	False	True
[DEFAULT] dnsmasq_dns_servers	None	
[DEFAULT] external_network_bridge	br-ex	
[DEFAULT] ipam_driver	None	internal
[OVS] of_interface	ovs-ofctl	native
[OVS] ovsdb_interface	vsctl	native
[ml2] path_mtu	1500	0
[ml2_sriov] supported_pci_vendor_devs	15b3:1004, 8086:10ca	None
[ml2_type_geneve] max_header_size	50	30

Table 5.36. Deprecated options

Deprecated option	New Option
[DEFAULT] min_l3_agents_per_router	None
[DEFAULT] use_syslog	None
[ml2_sriov] supported_pci_vendor_devs	None

CHAPTER 6. OBJECT STORAGE

6.1. OBJECT STORAGE CONFIGURATION OPTIONS

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.

6.1.1. Description of Configuration Options

The following tables provide a comprehensive list of the Object Storage configuration options.

Table 6.1. Description of configuration options for **[account-auditor]** in **account-server.conf**

Configuration option = Default value	Description
accounts_per_second = 200	Maximum accounts audited per second. Should be tuned according to individual system specs. 0 is unlimited.
interval = 1800	Minimum time for a pass to take
log_address = /dev/log	Location where syslog sends the logs to
log_facility = LOG_LOCAL0	Syslog log facility
log_level = INFO	Logging level
log_name = account-auditor	Label used when logging
recon_cache_path = /var/cache/swift	Directory where stats for a few items will be stored

Table 6.2. Description of configuration options for **[account-reaper]** in **account-server.conf**

Configuration option = Default value	Description
concurrency = 25	Number of replication workers to spawn
conn_timeout = 0.5	Connection timeout to external services

Configuration option = Default value	Description
delay_reaping = 0	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. bind to giving up worker can 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.
interval = 3600	Minimum time for a pass to take
log_address = /dev/log	Location where syslog sends the logs to
log_facility = LOG_LOCAL0	Syslog log facility
log_level = INFO	Logging level
log_name = account-reaper	Label used when logging
node_timeout = 10	Request timeout to external services
reap_warn_after = 2592000	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). This is in addition to any time requested by delay_reaping.

Table 6.3. Description of configuration options for[**account-replicator**] in **account-server.conf**

Configuration option = Default value	Description
concurrency = 8	Number of replication workers to spawn
conn_timeout = 0.5	Connection timeout to external services
interval = 30	Minimum time for a pass to take
log_address = /dev/log	Location where syslog sends the logs to

Configuration option = Default value	Description
log_facility = LOG_LOCAL0	Syslog log facility
log_level = INFO	Logging level
log_name = account-replicator	Label used when logging
max_diffs = 100	Caps how long the replicator spends trying to sync a database per pass
node_timeout = 10	Request timeout to external services
per_diff = 1000	Limit number of items to get per diff
reclaim_age = 604800	Time elapsed in seconds before an object can be reclaimed
recon_cache_path = /var/cache/swift	Directory where stats for a few items will be stored
rsync_compress = no	Allow rsync to compress data which is transmitted to destination node during sync. However, this is applicable only when destination node is in a different region than the local one.
rsync_module = {replication_ip}::account	Format of the rsync module where the replicator will send data. The configuration value can include some variables that will be extracted from the ring. Variables must follow the format {NAME} where NAME is one of: ip, port, replication_ip, replication_port, region, zone, device, meta. See etc/rsyncd.conf-sample for some examples. uses what's set here, or what's set in the DEFAULT section, or 10 (though other sections use 3 as the final default).
run_pause = 30	Time in seconds to wait between replication passes

Table 6.4. Description of configuration options for `[app-account-server]` in `account-server.conf`

Configuration option = Default value	Description
auto_create_account_prefix = .	Prefix to use when automatically creating accounts
replication_server = false	If defined, tells server how to handle replication verbs in requests. When set to True (or 1), only replication verbs will be accepted. When set to False, replication verbs will be rejected. When undefined, server will accept any verb in the request.

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_level = INFO	Log level
set log_name = account-server	Label to use when logging
set log_requests = true	Whether or not to log requests
use = egg:swift#account	Entry point of paste.deploy in the server

Table 6.5. Description of configuration options for [DEFAULT] in `account-server.conf`

Configuration option = Default value	Description
backlog = 4096	Maximum number of allowed pending TCP connections
bind_ip = 0.0.0.0	IP Address for server to bind to
bind_port = 6002	Port for server to bind to
bind_timeout = 30	Seconds to attempt bind before giving up
db_preallocation = off	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. 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. server. For most cases, this should be
devices = /srv/node	Parent directory of where devices are mounted
disable_fallocate = false	Disable "fast fail" fallocate checks if the underlying filesystem does not support it.
eventlet_debug = false	If true, turn on debug logging for eventlet

Configuration option = Default value	Description
fallocate_reserve = 0	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
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_LOCAL0	Syslog log facility
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 = swift	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 <code>log_statsd_default_sample_rate</code> 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.
max_clients = 1024	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.
mount_check = true	Whether or not check if the devices are mounted to prevent accidentally writing to the root device

Configuration option = Default value	Description
swift_dir = /etc/swift	Swift configuration directory
user = swift	User to run as
workers = auto	a much higher value, one can reduce the impact of slow file system operations in one request from negatively impacting other requests.

Table 6.6. Description of configuration options for **[filter-healthcheck]** in **account-server.conf**

Configuration option = Default value	Description
disable_path =	An optional filesystem path, which if present, will cause the healthcheck URL to return "503 Service Unavailable" with a body of "DISABLED BY FILE"
use = egg:swift#healthcheck	Entry point of paste.deploy in the server

Table 6.7. Description of configuration options for **[filter-recon]** in **account-server.conf**

Configuration option = Default value	Description
recon_cache_path = /var/cache/swift	Directory where stats for a few items will be stored
use = egg:swift#recon	Entry point of paste.deploy in the server

Table 6.8. Description of configuration options for **[filter-xprofile]** in **account-server.conf**

Configuration option = Default value	Description
dump_interval = 5.0	the profile data will be dumped to local disk based on above naming rule in this interval (seconds).
dump_timestamp = false	Be careful, this option will enable the profiler to dump data into the file with a time stamp which means that there will be lots of files piled up in the directory.
flush_at_shutdown = false	Clears the data when the wsgi server shutdowns.

Configuration option = Default value	Description
log_filename_prefix = /tmp/log/swift/profile/default.profile	This prefix is used to combine the process ID and timestamp to name the profile data file. Make sure the executing user has permission to write into this path. Any missing path segments will be created, if necessary. When you enable profiling in more than one type of daemon, you must override it with a unique value like: /var/log/swift/profile/accoutn.profile
path = /__profile__	This is the path of the URL to access the mini web UI.
profile_module = eventlet.green.profile	This option enables you to switch profilers which inherit from the Python standard profiler. Currently, the supported value can be 'cProfile', 'eventlet.green.profile', etc.
unwind = false	unwind the iterator of applications
use = egg:swift#xprofile	Entry point of paste.deploy in the server

Table 6.9. Description of configuration options for **[pipeline-main]** in **account-server.conf**

Configuration option = Default value	Description
pipeline = healthcheck recon account-server	Pipeline to use for processing operations.

Table 6.10. Description of configuration options for **[app-proxy-server]** in **container-reconciler.conf**

Configuration option = Default value	Description
use = egg:swift#proxy	Entry point of paste.deploy in the server

Table 6.11. Description of configuration options for **[container-reconciler]** in **container-reconciler.conf**

Configuration option = Default value	Description
interval = 30	Minimum time for a pass to take

Configuration option = Default value	Description
reclaim_age = 604800	Time elapsed in seconds before an object can be reclaimed
request_tries = 3	Server errors from requests will be retried by default

Table 6.12. Description of configuration options for **[DEFAULT]** in **container-reconciler.conf**

Configuration option = Default value	Description
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_LOCAL0	Syslog log facility
log_level = INFO	Logging level
log_name = swift	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 <code>log_statsd_default_sample_rate</code> 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.
swift_dir = /etc/swift	Swift configuration directory
user = swift	User to run as

Table 6.13. Description of configuration options for `[filter-cache]` in `container-reconciler.conf`

Configuration option = Default value	Description
<code>use = egg:swift#memcache</code>	Entry point of paste.deploy in the server

Table 6.14. Description of configuration options for `[filter-catch_errors]` in `container-reconciler.conf`

Configuration option = Default value	Description
<code>use = egg:swift#catch_errors</code>	Entry point of paste.deploy in the server

Table 6.15. Description of configuration options for `[filter-proxy-logging]` in `container-reconciler.conf`

Configuration option = Default value	Description
<code>use = egg:swift#proxy_logging</code>	Entry point of paste.deploy in the server

Table 6.16. Description of configuration options for `[pipeline-main]` in `container-reconciler.conf`

Configuration option = Default value	Description
<code>pipeline = catch_errors proxy-logging cache proxy-server</code>	Pipeline to use for processing operations.

Table 6.17. Description of configuration options for `[app-container-server]` in `container-server.conf`

Configuration option = Default value	Description
<code>allow_versions = false</code>	Enable/Disable object versioning feature
<code>auto_create_account_prefix = .</code>	Prefix to use when automatically creating accounts
<code>conn_timeout = 0.5</code>	Connection timeout to external services

Configuration option = Default value	Description
node_timeout = 3	Request timeout to external services
replication_server = false	If defined, tells server how to handle replication verbs in requests. When set to True (or 1), only replication verbs will be accepted. When set to False, replication verbs will be rejected. When undefined, server will accept any verb in the request.
set log_address = /dev/log	Location where syslog sends the logs to
set log_facility = LOG_LOCAL0	Syslog log facility
set log_level = INFO	Log level
set log_name = container-server	Label to use when logging
set log_requests = true	Whether or not to log requests
use = egg:swift#container	Entry point of paste.deploy in the server

Table 6.18. Description of configuration options for `[container-auditor]` in `container-server.conf`

Configuration option = Default value	Description
containers_per_second = 200	Maximum containers 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.
interval = 1800	Minimum time for a pass to take
log_address = /dev/log	Location where syslog sends the logs to
log_facility = LOG_LOCAL0	Syslog log facility
log_level = INFO	Logging level
log_name = container-auditor	Label used when logging
recon_cache_path = /var/cache/swift	Directory where stats for a few items will be stored

Table 6.19. Description of configuration options for[**container-replicator**] in **container-server.conf**

Configuration option = Default value	Description
concurrency = 8	Number of replication workers to spawn
conn_timeout = 0.5	Connection timeout to external services
interval = 30	Minimum time for a pass to take
log_address = /dev/log	Location where syslog sends the logs to
log_facility = LOG_LOCAL0	Syslog log facility
log_level = INFO	Logging level
log_name = container-replicator	Label used when logging
max_diffs = 100	Caps how long the replicator spends trying to sync a database per pass
node_timeout = 10	Request timeout to external services
per_diff = 1000	Limit number of items to get per diff
reclaim_age = 604800	Time elapsed in seconds before an object can be reclaimed
recon_cache_path = /var/cache/swift	Directory where stats for a few items will be stored
rsync_compress = no	Allow rsync to compress data which is transmitted to destination node during sync. However, this is applicable only when destination node is in a different region than the local one.
rsync_module = {replication_ip}::container	Format of the rsync module where the replicator will send data. The configuration value can include some variables that will be extracted from the ring. Variables must follow the format {NAME} where NAME is one of: ip, port, replication_ip, replication_port, region, zone, device, meta. See etc/rsyncd.conf-sample for some examples. uses what's set here, or what's set in the DEFAULT section, or 10 (though other sections use 3 as the final default).
run_pause = 30	Time in seconds to wait between replication passes

Table 6.20. Description of configuration options for[**container-sync**] in **container-server.conf**

Configuration option = Default value	Description
conn_timeout = 5	Connection timeout to external services
container_time = 60	Maximum amount of time to spend syncing each container
internal_client_conf_path = /etc/swift/internal-client.conf	Internal client config file path
interval = 300	Minimum time for a pass to take
log_address = /dev/log	Location where syslog sends the logs to
log_facility = LOG_LOCAL0	Syslog log facility
log_level = INFO	Logging level
log_name = container-sync	Label used when logging
request_tries = 3	Server errors from requests will be retried by default
sync_proxy = http://10.1.1.1:8888,http://10.1.1.2:8888	If you need to use an HTTP proxy, set it here. Defaults to no proxy.

Table 6.21. Description of configuration options for[**container-updater**] in **container-server.conf**

Configuration option = Default value	Description
account_suppression_time = 60	Seconds to suppress updating an account that has generated an error (timeout, not yet found, etc.)
concurrency = 4	Number of replication workers to spawn
conn_timeout = 0.5	Connection timeout to external services
interval = 300	Minimum time for a pass to take
log_address = /dev/log	Location where syslog sends the logs to
log_facility = LOG_LOCAL0	Syslog log facility
log_level = INFO	Logging level
log_name = container-updater	Label used when logging
node_timeout = 3	Request timeout to external services

Configuration option = Default value	Description
recon_cache_path = /var/cache/swift	Directory where stats for a few items will be stored
slowdown = 0.01	Time in seconds to wait between objects

Table 6.22. Description of configuration options for [DEFAULT] in `container-server.conf`

Configuration option = Default value	Description
allowed_sync_hosts = 127.0.0.1	The list of hosts that are allowed to send syncs to.
backlog = 4096	Maximum number of allowed pending TCP connections
bind_ip = 0.0.0.0	IP Address for server to bind to
bind_port = 6001	Port for server to bind to
bind_timeout = 30	Seconds to attempt bind before giving up
db_preallocation = off	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. 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. server. For most cases, this should be
devices = /srv/node	Parent directory of where devices are mounted
disable_fallocate = false	Disable "fast fail" fallocate checks if the underlying filesystem does not support it.
eventlet_debug = false	If true, turn on debug logging for eventlet
fallocate_reserve = 0	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. 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

Configuration option = Default value	Description
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_LOCAL0	Syslog log facility
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 = swift	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 <code>log_statsd_default_sample_rate</code> 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.
max_clients = 1024	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.
mount_check = true	Whether or not check if the devices are mounted to prevent accidentally writing to the root device
swift_dir = /etc/swift	Swift configuration directory
user = swift	User to run as

Configuration option = Default value	Description
workers = auto	a much higher value, one can reduce the impact of slow file system operations in one request from negatively impacting other requests.

Table 6.23. Description of configuration options for `[filter-healthcheck]` in `container-server.conf`

Configuration option = Default value	Description
disable_path =	An optional filesystem path, which if present, will cause the healthcheck URL to return "503 Service Unavailable" with a body of "DISABLED BY FILE"
use = egg:swift#healthcheck	Entry point of paste.deploy in the server

Table 6.24. Description of configuration options for `[filter-recon]` in `container-server.conf`

Configuration option = Default value	Description
recon_cache_path = /var/cache/swift	Directory where stats for a few items will be stored
use = egg:swift#recon	Entry point of paste.deploy in the server

Table 6.25. Description of configuration options for `[filter-xprofile]` in `container-server.conf`

Configuration option = Default value	Description
dump_interval = 5.0	the profile data will be dumped to local disk based on above naming rule in this interval (seconds).
dump_timestamp = false	Be careful, this option will enable the profiler to dump data into the file with a time stamp which means that there will be lots of files piled up in the directory.
flush_at_shutdown = false	Clears the data when the wsgi server shutdowns.

Configuration option = Default value	Description
log_filename_prefix = /tmp/log/swift/profile/default.profile	This prefix is used to combine the process ID and timestamp to name the profile data file. Make sure the executing user has permission to write into this path. Any missing path segments will be created, if necessary. When you enable profiling in more than one type of daemon, you must override it with a unique value like: /var/log/swift/profile/object.profile
path = /__profile__	This is the path of the URL to access the mini web UI.
profile_module = eventlet.green.profile	This option enables you to switch profilers which inherit from the Python standard profiler. Currently, the supported value can be 'cProfile', 'eventlet.green.profile', etc.
unwind = false	unwind the iterator of applications
use = egg:swift#xprofile	Entry point of paste.deploy in the server

Table 6.26. Description of configuration options for[**pipeline-main**] in **container-server.conf**

Configuration option = Default value	Description
pipeline = healthcheck recon container-server	Pipeline to use for processing operations.

Table 6.27. Description of configuration options for[**DEFAULT**] in **container-sync-realms.conf**

Configuration option = Default value	Description
mtime_check_interval = 300	The number of seconds between checking the modified time of this config file for changes and therefore reloading it.

Table 6.28. Description of configuration options for[**realm1**] in **container-sync-realms.conf**

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

Configuration option = Default value	Description
cluster_clustername1 = https://host1/v1/	Any values in the realm section whose names begin with cluster_ will indicate the name and endpoint of a cluster and will be used by external users in their containers' X-Container-Sync-To metadata header values with the format "realm_name/cluster_name/container_name". Realm and cluster names are considered case insensitive.
cluster_clustername2 = https://host2/v1/	Any values in the realm section whose names begin with cluster_ will indicate the name and endpoint of a cluster and will be used by external users in their containers' X-Container-Sync-To metadata header values with the format "realm_name/cluster_name/container_name". Realm and cluster names are considered case insensitive.
key = realm1key	The key is the overall cluster-to-cluster key used in combination with the external users' key that they set on their containers' X-Container-Sync-Key metadata header values. These keys will be used to sign each request the container sync daemon makes and used to validate each incoming container sync request.
key2 = realm1key2	The key2 is optional and is an additional key incoming requests will be checked against. This is so you can rotate keys if you wish; you move the existing key to key2 and make a new key value.

Table 6.29. Description of configuration options for [realm2] in container-sync-realms.conf

Configuration option = Default value	Description
cluster_clustername3 = https://host3/v1/	Any values in the realm section whose names begin with cluster_ will indicate the name and endpoint of a cluster and will be used by external users in their containers' X-Container-Sync-To metadata header values with the format "realm_name/cluster_name/container_name". Realm and cluster names are considered case insensitive.
cluster_clustername4 = https://host4/v1/	Any values in the realm section whose names begin with cluster_ will indicate the name and endpoint of a cluster and will be used by external users in their containers' X-Container-Sync-To metadata header values with the format "realm_name/cluster_name/container_name". Realm and cluster names are considered case insensitive.

Configuration option = Default value	Description
key = realm2key	The key is the overall cluster-to-cluster key used in combination with the external users' key that they set on their containers' X-Container-Sync-Key metadata header values. These keys will be used to sign each request the container sync daemon makes and used to validate each incoming container sync request.
key2 = realm2key2	The key2 is optional and is an additional key incoming requests will be checked against. This is so you can rotate keys if you wish; you move the existing key to key2 and make a new key value.

Table 6.30. Description of configuration options for **[dispersion]** in **dispersion.conf**

Configuration option = Default value	Description
auth_key = testing	No help text available for this option.
auth_url = http://localhost:8080/auth/v1.0	Endpoint for auth server, such as keystone
auth_user = test:tester	Default user for dispersion in this context
auth_version = 1.0	Indicates which version of auth
concurrency = 25	Number of replication workers to spawn
container_populate = yes	No help text available for this option.
container_report = yes	No help text available for this option.
dispersion_coverage = 1.0	No help text available for this option.
dump_json = no	No help text available for this option.
endpoint_type = publicURL	Indicates whether endpoint for auth is public or internal
keystone_api_insecure = no	Allow accessing insecure keystone server. The keystone's certificate will not be verified.
object_populate = yes	No help text available for this option.
object_report = yes	No help text available for this option.

Configuration option = Default value	Description
project_domain_name = project_domain	No help text available for this option.
project_name = project	No help text available for this option.
retries = 5	No help text available for this option.
swift_dir = /etc/swift	Swift configuration directory
user_domain_name = user_domain	No help text available for this option.

Table 6.31. Description of configuration options for `[drive-audit]` in `drive-audit.conf`

Configuration option = Default value	Description
device_dir = /srv/node	Directory devices are mounted under
error_limit = 1	Number of errors to find before a device is unmounted
log_address = /dev/log	Location where syslog sends the logs to
log_facility = LOG_LOCAL0	Syslog log facility
log_file_pattern = /var/log/kern.*[!][!g][!z]	Location of the log file with globbing pattern to check against device errors locate device blocks with errors in the log file
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 = drive-audit	Label used when logging
log_to_console = False	No help text available for this option.
minutes = 60	Number of minutes to look back in
recon_cache_path = /var/cache/swift	Directory where stats for a few items will be stored
regex_pattern_1 = \berror\b.*\b(dm-[0-9]{1,2}\d?)\b	No help text available for this option.
unmount_failed_device = True	No help text available for this option.

Table 6.32. Description of configuration options for `[app-proxy-server]` in `internal-client.conf`

Configuration option = Default value	Description
<code>use = egg:swift#proxy</code>	Entry point of paste.deploy in the server

Table 6.33. Description of configuration options for `[DEFAULT]` in `internal-client.conf`

Configuration option = Default value	Description
<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_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>swift_dir = /etc/swift</code>	Swift configuration directory
<code>user = swift</code>	User to run as

Table 6.34. Description of configuration options for `[filter-cache]` in `internal-client.conf`

Configuration option = Default value	Description
use = egg:swift#memcache	Entry point of paste.deploy in the server

Table 6.35. Description of configuration options for `[filter-catch_errors]` in `internal-client.conf`

Configuration option = Default value	Description
use = egg:swift#catch_errors	Entry point of paste.deploy in the server

Table 6.36. Description of configuration options for `[filter-proxy-logging]` in `internal-client.conf`

Configuration option = Default value	Description
use = egg:swift#proxy_logging	Entry point of paste.deploy in the server

Table 6.37. Description of configuration options for `[pipeline-main]` in `internal-client.conf`

Configuration option = Default value	Description
pipeline = catch_errors proxy-logging cache proxy-server	No help text available for this option.

Table 6.38. Description of configuration options for `[memcache]` in `memcache.conf`

Configuration option = Default value	Description
connect_timeout = 0.3	Timeout in seconds (float) for connection.
io_timeout = 2.0	Timeout in seconds (float) for read and write.
memcache_max_connections = 2	Max number of connections to each memcached server per worker services.
memcache_serialization_support = 2	Sets how memcache values are serialized and deserialized.

Configuration option = Default value	Description
memcache_servers = 127.0.0.1:11211	Comma-separated list of memcached servers ip:port services.
pool_timeout = 1.0	Timeout in seconds (float) for pooled connection.
tries = 3	Number of servers to retry on failures getting a pooled connection.

Table 6.39. Description of configuration options for **[app-proxy-server]** in **object-expirer.conf**

Configuration option = Default value	Description
use = egg:swift#proxy	Entry point of paste.deploy in the server

Table 6.40. Description of configuration options for **[DEFAULT]** in **object-expirer.conf**

Configuration option = Default value	Description
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_LOCAL0	Syslog log facility
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 = swift	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.

Configuration option = Default value	Description
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 <code>log_statsd_default_sample_rate</code> 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.
swift_dir = /etc/swift	Swift configuration directory
user = swift	User to run as

Table 6.41. Description of configuration options for `[filter-cache]` in `object-expirer.conf`

Configuration option = Default value	Description
use = egg:swift#memcache	Entry point of paste.deploy in the server

Table 6.42. Description of configuration options for `[filter-catch_errors]` in `object-expirer.conf`

Configuration option = Default value	Description
use = egg:swift#catch_errors	Entry point of paste.deploy in the server

Table 6.43. Description of configuration options for `[filter-proxy-logging]` in `object-expirer.conf`

Configuration option = Default value	Description
access_log_address = /dev/log	Location where syslog sends the logs to. If not set, logging directives from [DEFAULT] without "access_" will be used.
access_log_facility = LOG_LOCAL0	Syslog facility to receive log lines. If not set, logging directives from [DEFAULT] without "access_" will be used.

Configuration option = Default value	Description
access_log_headers = false	Header to receive log lines. If not set, logging directives from [DEFAULT] without "access_" will be used.
access_log_headers_only =	If access_log_headers is True and access_log_headers_only is set only these headers are logged. Multiple headers can be defined as comma separated list like this: access_log_headers_only = Host, X-Object-Meta-Mtime
access_log_level = INFO	Syslog logging level to receive log lines. If not set, logging directives from [DEFAULT] without "access_" will be used.
access_log_name = swift	Label used when logging. If not set, logging directives from [DEFAULT] without "access_" will be used.
access_log_statsd_default_sample_rate = 1.0	Defines the probability of sending a sample for any given event or timing measurement. If not set, logging directives from [DEFAULT] without "access_" will be used.
access_log_statsd_host = localhost	You can use log_statsd_* from [DEFAULT], or override them here. StatsD server. IPv4/IPv6 addresses and hostnames are supported. If a hostname resolves to an IPv4 and IPv6 address, the IPv4 address will be used.
access_log_statsd_metric_prefix =	Value will be prepended to every metric sent to the StatsD server. If not set, logging directives from [DEFAULT] without "access_" will be used.
access_log_statsd_port = 8125	Port value for the StatsD server. If not set, logging directives from [DEFAULT] without "access_" will be used.
access_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. If not set, logging directives from [DEFAULT] without "access_" will be used.
access_log_udp_host =	If not set, the UDP receiver for syslog is disabled. If not set, logging directives from [DEFAULT] without "access_" will be used.
access_log_udp_port = 514	Port value for UDP receiver, if enabled. If not set, logging directives from [DEFAULT] without "access_" will be used.


Configuration option = Default value	Description
log_statsd_valid_http_methods = GET,HEAD,POST,PUT,DELETE,COPY,OPTIONS	What HTTP methods are allowed for StatsD logging (comma-sep). request methods not in this list will have "BAD_METHOD" for the <verb> portion of the metric.
reveal_sensitive_prefix = 16	<p>By default, the X-Auth-Token is logged. To obscure the value, set reveal_sensitive_prefix to the number of characters to log. For example, if set to 12, only the first 12 characters of the token appear in the log. An unauthorized access of the log file won't allow unauthorized usage of the token. However, the first 12 or so characters is unique enough that you can trace/debug token usage. Set to 0 to suppress the token completely (replaced by '...' in the log).</p> <div style="display: flex; align-items: flex-start;">  <div> <p>NOTE</p> <p>reveal_sensitive_prefix will not affect the value logged with access_log_headers=True.</p> </div> </div>
use = egg:swift#proxy_logging	Entry point of paste.deploy in the server

Table 6.44. Description of configuration options for **[object-expirer]** in **object-expirer.conf**

Configuration option = Default value	Description
auto_create_account_prefix = .	Prefix to use when automatically creating accounts
concurrency = 1	Number of replication workers to spawn
expiring_objects_account_name = expiring_objects	Account name for expiring objects.
interval = 300	Minimum time for a pass to take
process = 0	(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.
processes = 0	for each port (disk) in the ring. If you have 24 disks per server, and this setting is 4, then each storage node will have 1 + (24 * 4) = 97 total object-server processes running. This gives complete I/O isolation, drastically reducing the impact of slow disks on storage node performance. The object-replicator and object-reconstructor need to see this setting too, so it must be in the [DEFAULT] section.

Configuration option = Default value	Description
reclaim_age = 604800	Time elapsed in seconds before an object can be reclaimed
recon_cache_path = /var/cache/swift	Directory where stats for a few items will be stored
report_interval = 300	Interval in seconds between reports.

Table 6.45. Description of configuration options for **[pipeline-main]** in **object-expirer.conf**

Configuration option = Default value	Description
pipeline = catch_errors proxy-logging cache proxy-server	Pipeline to use for processing operations.

Table 6.46. Description of configuration options for **[app-object-server]** in **object-server.conf**

Configuration option = Default value	Description
allowed_headers = Content-Disposition, Content-Encoding, X-Delete-At, X-Object-Manifest, X-Static-Large-Object	Comma-separated list of headers that can be set in metadata of an object
auto_create_account_prefix = .	Prefix to use when automatically creating accounts
keep_cache_private = false	Allow non-public objects to stay in kernel's buffer cache
keep_cache_size = 5242880	Largest object size to keep in buffer cache
max_upload_time = 86400	Maximum time allowed to upload an object
mb_per_sync = 512	On PUT requests, sync file every n MB
replication_concurrency = 4	Set to restrict the number of concurrent incoming REPLICATION requests; set to 0 for unlimited
replication_failure_ratio = 1.0	If the value of failures / successes of REPLICATION subrequests exceeds this ratio, the overall REPLICATION request will be aborted
replication_failure_threshold = 100	The number of subrequest failures before the replication_failure_ratio is checked

Configuration option = Default value	Description
replication_lock_timeout = 15	Number of seconds to wait for an existing replication device lock before giving up.
replication_one_per_device = True	Restricts incoming REPLICATION requests to one per device, replication_currency above allowing. This can help control I/O to each device, but you may wish to set this to False to allow multiple REPLICATION requests (up to the above replication_concurrency setting) per device.
replication_server = false	If defined, tells server how to handle replication verbs in requests. When set to True (or 1), only replication verbs will be accepted. When set to False, replication verbs will be rejected. When undefined, server will accept any verb in the request.
set log_address = /dev/log	Location where syslog sends the logs to
set log_facility = LOG_LOCAL0	Syslog log facility
set log_level = INFO	Log level
set log_name = object-server	Label to use when logging
set log_requests = true	Whether or not to log requests
slow = 0	If > 0, Minimum time in seconds for a PUT or DELETE request to complete
splice = no	Use splice() for zero-copy object GETs. This requires Linux kernel version 3.0 or greater. When you set "splice = yes" but the kernel does not support it, error messages will appear in the object server logs at startup, but your object servers should continue to function.
threads_per_disk = 0	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.
use = egg:swift#object	Entry point of paste.deploy in the server

Table 6.47. Description of configuration options for[DEFAULT] in object-server.conf

Configuration option = Default value	Description
backlog = 4096	Maximum number of allowed pending TCP connections
bind_ip = 0.0.0.0	IP Address for server to bind to
bind_port = 6000	Port for server to bind to
bind_timeout = 30	Seconds to attempt bind before giving up
client_timeout = 60	Time to wait while receiving each chunk of data from a client or another backend node
conn_timeout = 0.5	Connection timeout to external services
container_update_timeout = 1.0	Time to wait while sending a container update on object update. object server. For most cases, this should be
devices = /srv/node	Parent directory of where devices are mounted
disable_fallocate = false	Disable "fast fail" fallocate checks if the underlying filesystem does not support it.
disk_chunk_size = 65536	Size of chunks to read/write to disk
eventlet_debug = false	If true, turn on debug logging for eventlet
expiring_objects_account_name = expiring_objects	Account name for the expiring objects
expiring_objects_container_divisor = 86400	Divisor for the expiring objects container
fallocate_reserve = 0	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. 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
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_LOCAL0	Syslog log facility
log_level = INFO	Logging level

Configuration option = Default value	Description
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 = swift	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 <code>log_statsd_default_sample_rate</code> 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.
max_clients = 1024	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.
mount_check = true	Whether or not check if the devices are mounted to prevent accidentally writing to the root device
network_chunk_size = 65536	Size of chunks to read/write over the network
node_timeout = 3	Request timeout to external services
servers_per_port = 0	If each disk in each storage policy ring has unique port numbers for its "ip" value, you can use this setting to have each object-server worker only service requests for the single disk matching the port in the ring. The value of this setting determines how many worker processes run for each port (disk) in the
swift_dir = /etc/swift	Swift configuration directory
user = swift	User to run as

Configuration option = Default value	Description
workers = auto	a much higher value, one can reduce the impact of slow file system operations in one request from negatively impacting other requests.

Table 6.48. Description of configuration options for **[filter-healthcheck]** in **object-server.conf**

Configuration option = Default value	Description
disable_path =	An optional filesystem path, which if present, will cause the healthcheck URL to return "503 Service Unavailable" with a body of "DISABLED BY FILE"
use = egg:swift#healthcheck	Entry point of paste.deploy in the server

Table 6.49. Description of configuration options for **[filter-recon]** in **object-server.conf**

Configuration option = Default value	Description
recon_cache_path = /var/cache/swift	Directory where stats for a few items will be stored
recon_lock_path = /var/lock	Directory where lock files will be stored
use = egg:swift#recon	Entry point of paste.deploy in the server

Table 6.50. Description of configuration options for **[filter-xprofile]** in **object-server.conf**

Configuration option = Default value	Description
dump_interval = 5.0	the profile data will be dumped to local disk based on above naming rule in this interval (seconds).
dump_timestamp = false	Be careful, this option will enable the profiler to dump data into the file with a time stamp which means that there will be lots of files piled up in the directory.
flush_at_shutdown = false	Clears the data when the wsgi server shutdowns.

Configuration option = Default value	Description
log_filename_prefix = /tmp/log/swift/profile/default.profile	This prefix is used to combine the process ID and timestamp to name the profile data file. Make sure the executing user has permission to write into this path. Any missing path segments will be created, if necessary. When you enable profiling in more than one type of daemon, you must override it with a unique value like: /var/log/swift/profile/object.profile
path = /__profile__	This is the path of the URL to access the mini web UI.
profile_module = eventlet.green.profile	This option enables you to switch profilers which inherit from the Python standard profiler. Currently, the supported value can be 'cProfile', 'eventlet.green.profile', etc.
unwind = false	unwind the iterator of applications
use = egg:swift#xprofile	Entry point of paste.deploy in the server

Table 6.51. Description of configuration options for[**object-auditor**] in **object-server.conf**

Configuration option = Default value	Description
bytes_per_second = 10000000	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
concurrency = 1	Number of replication workers to spawn
disk_chunk_size = 65536	Size of chunks to read/write to disk
files_per_second = 20	Maximum files audited per second. Should be tuned according to individual system specs. 0 is unlimited.

Configuration option = Default value	Description
log_address = /dev/log	Location where syslog sends the logs to
log_facility = LOG_LOCAL0	Syslog log facility
log_level = INFO	Logging level
log_name = object-auditor	Label used when logging
log_time = 3600	Frequency of status logs in seconds.
object_size_stats =	Takes a comma-separated list of ints. When set, the object auditor will increment a counter for every object whose size is greater or equal to the given breaking points and reports the result after a full scan.
recon_cache_path = /var/cache/swift	Directory where stats for a few items will be stored
zero_byte_files_per_second = 50	Maximum zero byte files audited per second.


Table 6.52. Description of configuration options for `[object-reconstructor]` in `object-server.conf`

Configuration option = Default value	Description
concurrency = 1	Number of replication workers to spawn
daemonize = on	Whether or not to run replication as a daemon
handoffs_first = False	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.
http_timeout = 60	Maximum duration for an HTTP request
interval = 30	Minimum time for a pass to take
lockup_timeout = 1800	Attempts to kill all workers if nothing replications for lockup_timeout seconds
log_address = /dev/log	Location where syslog sends the logs to
log_facility = LOG_LOCAL0	Syslog log facility
log_level = INFO	Logging level

Configuration option = Default value	Description
log_name = object-reconstructor	Label used when logging
node_timeout = 10	Request timeout to external services
reclaim_age = 604800	Time elapsed in seconds before an object can be reclaimed
recon_cache_path = /var/cache/swift	Directory where stats for a few items will be stored
ring_check_interval = 15	How often (in seconds) to check the ring
run_pause = 30	Time in seconds to wait between replication passes
stats_interval = 300	Interval in seconds between logging replication statistics

Table 6.53. Description of configuration options for **[object-replicator]** in **object-server.conf**

Configuration option = Default value	Description
concurrency = 1	Number of replication workers to spawn
daemonize = on	Whether or not to run replication as a daemon
handoff_delete = auto	By default handoff partitions will be removed when it has successfully replicated to all the canonical nodes. If set to an integer n, it will remove the partition if it is successfully replicated to n nodes. The default setting should not be changed, except for extremem 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).
handoffs_first = False	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.
http_timeout = 60	Maximum duration for an HTTP request
interval = 30	Minimum time for a pass to take
lockup_timeout = 1800	Attempts to kill all workers if nothing replications for lockup_timeout seconds
log_address = /dev/log	Location where syslog sends the logs to

Configuration option = Default value	Description
log_facility = LOG_LOCAL0	Syslog log facility
log_level = INFO	Logging level
log_name = object-replicator	Label used when logging
node_timeout = <whatever's in the DEFAULT section or 10>	Request timeout to external services
reclaim_age = 604800	Time elapsed in seconds before an object can be reclaimed
recon_cache_path = /var/cache/swift	Directory where stats for a few items will be stored
ring_check_interval = 15	How often (in seconds) to check the ring
rsync_bwlimit = 0	bandwidth limit for rsync in kB/s. 0 means unlimited
rsync_compress = no	<p>Allows rsync to compress data which is transmitted to the destination node during sync. However, this applies only when the destination node is in a different region than the local one.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p>NOTE</p> <p>Objects that are already compressed (for example: .tar.gz, .mp3) might slow down the syncing process.</p> </div> </div>
rsync_error_log_line_length = 0	Limits the length of the rsync error log lines. 0 will log the entire line.
rsync_io_timeout = 30	Passed to rsync for a max duration (seconds) of an I/O op
rsync_module = {replication_ip}::object	Format of the rsync module where the replicator will send data. The configuration value can include some variables that will be extracted from the ring. Variables must follow the format {NAME} where NAME is one of: ip, port, replication_ip, replication_port, region, zone, device, meta. See etc/rsyncd.conf-sample for some examples. uses what's set here, or what's set in the DEFAULT section, or 10 (though other sections use 3 as the final default).
rsync_timeout = 900	Max duration (seconds) of a partition rsync
run_pause = 30	Time in seconds to wait between replication passes

Configuration option = Default value	Description
stats_interval = 300	Interval in seconds between logging replication statistics
sync_method = rsync	default is rsync, alternative is ssync

Table 6.54. Description of configuration options for **[object-updater]** in **object-server.conf**

Configuration option = Default value	Description
concurrency = 1	Number of replication workers to spawn
interval = 300	Minimum time for a pass to take
log_address = /dev/log	Location where syslog sends the logs to
log_facility = LOG_LOCAL0	Syslog log facility
log_level = INFO	Logging level
log_name = object-updater	Label used when logging
node_timeout = <whatever's in the DEFAULT section or 10>	Request timeout to external services
recon_cache_path = /var/cache/swift	Directory where stats for a few items will be stored
slowdown = 0.01	Time in seconds to wait between objects

Table 6.55. Description of configuration options for **[pipeline-main]** in **object-server.conf**

Configuration option = Default value	Description
pipeline = healthcheck recon object-server	Pipeline to use for processing operations.

Table 6.56. Description of configuration options for **[app-proxy-server]** in **proxy-server.conf**

Configuration option = Default value	Description
account_autocreate = false	If set to 'true' authorized accounts that do not yet exist within the Swift cluster will be automatically created.
allow_account_management = false	Whether account PUTs and DELETEs are even callable.
auto_create_account_prefix = .	Prefix to use when automatically creating accounts.
client_chunk_size = 65536	Chunk size to read from clients.
conn_timeout = 0.5	Connection timeout to external services.
deny_host_headers =	Comma separated list of Host headers to which the proxy will deny requests.
error_suppression_interval = 60	Time in seconds that must elapse since the last error for a node to be considered no longer error limited.
error_suppression_limit = 10	Error count to consider a node error limited.
log_handoffs = true	<p>Log handoff requests if handoff logging is enabled and the handoff was not expected.</p> <p>We only log handoffs when we've pushed the handoff count further than we would normally have expected under normal circumstances, that is (request_node_count - num primaries), when handoffs goes higher than that it means one of the primaries must have been skipped because of error limiting before we consumed all of our nodes_left.</p>
max_containers_per_account = 0	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_whitelist =	is a comma separated list of account names that ignore the max_containers_per_account cap.
node_timeout = 10	Request timeout to external services.
object_chunk_size = 65536	Chunk size to read from object servers.
object_post_as_copy = true	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.

Configuration option = Default value	Description
post_quorum_timeout = 0.5	How long to wait for requests to finish after a quorum has been established.
put_queue_depth = 10	Depth of the proxy put queue.
read_affinity = r1z1=100, r1z2=200, r2=300	<p>Which backend servers to prefer on reads. Format is r<N> for region N or r<N>z<M> for region N, zone M. The value after the equals is the priority; lower numbers are higher priority.</p> <p>Example: first read from region 1 zone 1, then region 1 zone 2, then anything in region 2, then everything else: read_affinity = r1z1=100, r1z2=200, r2=300</p> <p>Default is empty, meaning no preference.</p>
recheck_account_existence = 60	Cache timeout in seconds to send memcached for account existence.
recheck_container_existence = 60	Cache timeout in seconds to send memcached for container existence.
recoverable_node_timeout = node_timeout	Request timeout to external services for requests that, on failure, can be recovered from. For example, object GET. from a client external services.
request_node_count = 2 * replicas	<p>replicas Set to the number of nodes to 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. conf file for values will only be shown to the list of swift_owners. The exact default definition of a swift_owner is headers> up to the auth system in use, but usually indicates administrative responsibilities. paste.deploy to use for auth. To use tempauth set to:</p>
set log_address = /dev/log	Location where syslog sends the logs to.
set log_facility = LOG_LOCAL0	Syslog log facility.
set log_level = INFO	Log level.
set log_name = proxy-server	Label to use when logging.

Configuration option = Default value	Description
sorting_method = shuffle	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", or "timing".
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-container-meta-temp-url-key, x-container-meta-temp-url-key-2, x-account-access-control	These are the headers whose conf file for values will only be shown to the list of swift_owners. The exact default definition of a swift_owner is headers> up to the auth system in use, but usually indicates administrative responsibilities. paste.deploy to use for auth. To use tempauth set to:
timing_expiry = 300	If the "timing" sorting_method is used, the timings will only be valid for the number of seconds configured by timing_expiry.
use = egg:swift#proxy	Entry point of paste.deploy in the server.
write_affinity = r1, r2	This setting lets you trade data distribution for throughput. It makes the proxy server prefer local back-end servers for object PUT requests over non-local ones. Note that only object PUT requests are affected by the write_affinity setting; POST, GET, HEAD, DELETE, OPTIONS, and account/container PUT requests are not affected. The format is r<N> for region N or r<N>z<M> for region N, zone M. If this is set, then when handling an object PUT request, some number (see the write_affinity_node_count setting) of local backend servers will be tried before any nonlocal ones. Example: try to write to regions 1 and 2 before writing to any other nodes: write_affinity = r1, r2
write_affinity_node_count = 2 * replicas	This setting is only useful in conjunction with write_affinity; it governs how many local object servers will be tried before falling back to non-local ones. 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: write_affinity_node_count = 2 * replicas

Table 6.57. Description of configuration options for [DEFAULT] in proxy-server.conf

Configuration option = Default value	Description
admin_key = secret_admin_key	To use for admin calls that are HMAC signed. Default is empty, which will disable admin calls to /info.
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	Time to wait while receiving each chunk of data from a client or another backend node.
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.
disallowed_sections = swift.valid_api_versions, container_quotas, tempurl	Allows the ability to withhold sections from showing up in the public calls to /info. You can withhold subsections by separating the dict level with a ".". The following would cause the sections 'container_quotas' and 'tempurl' to not be listed, and the key max_failed_deletes would be removed from bulk_delete. Default value is 'swift.valid_api_versions' which allows all registered features to be listed via HTTP GET /info except swift.valid_api_versions information.
eventlet_debug = false	If true, turn on debug logging for eventlet.
expiring_objects_account_name = expiring_objects	Account name for the expiring objects.
expiring_objects_container_divisor = 86400	Divisor for the expiring objects container.
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.

Configuration option = Default value	Description
log_custom_handlers =	Comma-separated list of functions to call to setup custom log handlers.
log_facility = LOG_LOCAL0	Syslog log facility.
log_headers = false	Enables the ability to log request headers.
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 = swift	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.
max_clients = 1024	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.
strict_cors_mode = True	Enforce CORS.
swift_dir = /etc/swift	Swift configuration directory.
trans_id_suffix =	This optional suffix (default is empty) that would be appended to the swift transaction id allows one to easily figure out from which cluster that X-Trans-Id belongs to. This is very useful when one is managing more than one swift cluster.

Configuration option = Default value	Description
user = swift	User to run as.
workers = auto	a much higher value, one can reduce the impact of slow file system operations in one request from negatively impacting other requests.

Table 6.58. Description of configuration options for `[filter-account-quotas]` in `proxy-server.conf`

Configuration option = Default value	Description
use = egg:swift#account_quotas	Entry point of paste.deploy in the server

Table 6.59. Description of configuration options for `[filter-authtoken]` in `proxy-server.conf`

Configuration option = Default value	Description
auth_plugin = password	Authentication module to use.
auth_uri = http://keystonehost:5000	auth_uri should point to a Keystone service from which users may retrieve tokens. This value is used in the WWW-Authenticate header that auth_token sends with any denial response.
auth_url = http://keystonehost:35357	auth_url points to the Keystone Admin service. This information is used by the middleware to actually query Keystone about the validity of the authentication tokens. It is not necessary to append any Keystone API version number to this URI.
cache = swift.cache	cache is set to swift.cache . This means that the middleware will get the Swift memcache from the request environment.
delay_auth_decision = False	delay_auth_decision defaults to False, but leaving it as false will prevent other auth systems, staticweb, tempurl, formpost, and ACLs from working. This value must be explicitly set to True.
include_service_catalog = False	include_service_catalog defaults to True if not set. This means that when validating a token, the service catalog is retrieved and stored in the X-Service-Catalog header. Since Swift does not use the X-Service-Catalog header, there is no point in getting the service catalog. We recommend you set include_service_catalog to False.

Configuration option = Default value	Description
password = password	Password for service user.
paste.filter_factory = keystonemiddleware.auth_token:filter_factory	Entry point of paste.filter_factory in the server.
project_domain_id = default	Service project domain.
project_name = service	Service project name.
user_domain_id = default	Service user domain.
username = swift	Service user name.

Table 6.60. Description of configuration options for **[filter-bulk]** in **proxy-server.conf**

Configuration option = Default value	Description
delete_container_retry_count = 0	The parameter is used during a bulk delete of objects and their container. This would frequently fail because it is very likely that all replicated objects have not been deleted by the time the middleware got a successful response. It can be configured the number of retries. And the number of seconds to wait between each retry will be $1.5 \times \text{retry}$.
max_containers_per_extraction = 10000	The maximum numbers of containers per extraction.
max_deletes_per_request = 10000	The maximum numbers of deletion per request.
max_failed_deletes = 1000	The maximum number of tries to delete before failure.
max_failed_extractions = 1000	The maximum number of tries to extract before failure.
use = egg:swift#bulk	Entry point of paste.deploy in the server.
yield_frequency = 10	In order to keep a connection active during a potentially long bulk request, Swift may return whitespace prepended to the actual response body. This whitespace will be yielded no more than every <code>yield_frequency</code> seconds.

Table 6.61. Description of configuration options for **[filter-cache]** in **proxy-server.conf**

Configuration option = Default value	Description
memcache_max_connections = 2	Max number of connections to each memcached server per worker services
memcache_serialization_support = 2	Sets how memcache values are serialized and deserialized
memcache_servers = 127.0.0.1:11211	Comma-separated list of memcached servers ip:port services
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 = cache	Label to use when logging
use = egg:swift#memcache	Entry point of paste.deploy in the server

Table 6.62. Description of configuration options for `[filter-catch_errors]` 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 = catch_errors	Label to use when logging
use = egg:swift#catch_errors	Entry point of paste.deploy in the server

Table 6.63. Description of configuration options for `[filter-cname_lookup]` in `proxy-server.conf`

Configuration option = Default value	Description
lookup_depth = 1	Because CNAMEs can be recursive, specifies the number of levels through which to search.
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 = cname_lookup	Label to use when logging
storage_domain = example.com	Domain that matches your cloud. Multiple domains can be specified using a comma-separated list.
use = egg:swift#cname_lookup	Entry point of paste.deploy in the server

Table 6.64. Description of configuration options for **[filter-container-quotas]** in **proxy-server.conf**

Configuration option = Default value	Description
use = egg:swift#container_quotas	Entry point of paste.deploy in the server

Table 6.65. Description of configuration options for **[filter-container_sync]** in **proxy-server.conf**

Configuration option = Default value	Description
allow_full_urls = true	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.
current = //REALM/CLUSTER	Set this to specify this cluster //realm/cluster as "current" in /info.
use = egg:swift#container_sync	Entry point of paste.deploy in the server.

Table 6.66. Description of configuration options for `[filter-dlo]` in `proxy-server.conf`

Configuration option = Default value	Description
<code>max_get_time = 86400</code>	Time limit on GET requests (seconds).
<code>rate_limit_after_segment = 10</code>	Rate limit the download of large object segments after this segment is downloaded.
<code>rate_limit_segments_per_sec = 1</code>	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:
<code>use = egg:swift#dlo</code>	Entry point of paste.deploy in the server.

Table 6.67. Description of configuration options for `[filter-domain_remap]` in `proxy-server.conf`

Configuration option = Default value	Description
<code>default_reseller_prefix =</code>	If the reseller prefixes do not match, the default reseller prefix is used. When no default reseller prefix is configured, any request with an account prefix not in that list will be ignored by this middleware.
<code>path_root = v1</code>	Root path.
<code>reseller_prefixes = AUTH</code>	Browsers can convert a host header to lowercase, so check that reseller prefix on the account is the correct case. This is done by comparing the items in the <code>reseller_prefixes</code> config option to the found prefix. If they match except for case, the item from <code>reseller_prefixes</code> will be used instead of the found reseller prefix.
<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_headers = false</code>	If True, log headers in each request.
<code>set log_level = INFO</code>	Log level.
<code>set log_name = domain_remap</code>	Label to use when logging.
<code>storage_domain = example.com</code>	Domain that matches your cloud. Multiple domains can be specified using a comma-separated list.

Configuration option = Default value	Description
use = egg:swift#domain_remap	Entry point of paste.deploy in the server.

Table 6.68. Description of configuration options for **[filter-formpost]** in **proxy-server.conf**

Configuration option = Default value	Description
use = egg:swift#formpost	Entry point of paste.deploy in the server

Table 6.69. 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 6.70. Description of configuration options for **[filter-healthcheck]** in **proxy-server.conf**

Configuration option = Default value	Description
disable_path =	An optional filesystem path, which if present, will cause the healthcheck URL to return "503 Service Unavailable" with a body of "DISABLED BY FILE".
use = egg:swift#healthcheck	Entry point of paste.deploy in the server.

Table 6.71. 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".
is_admin = false	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.
operator_roles = admin, swiftoperator	Operator role defines the user which is allowed to manage a tenant and create containers or give ACL to others. This parameter may be prefixed with an appropriate prefix.
reseller_admin_role = ResellerAdmin	The reseller admin role gives the ability to create and delete accounts.
reseller_prefix = AUTH	The naming scope for the auth service.
service_roles =	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.
use = egg:swift#keystoneauth	Entry point of paste.deploy in the server.

Table 6.72. Description of configuration options for `[filter-list-endpoints]` in `proxy-server.conf`

Configuration option = Default value	Description
list_endpoints_path = /endpoints/	Path to list endpoints for an object, account or container.
use = egg:swift#list_endpoints	Entry point of paste.deploy in the server.

Configuration option = Default value	Description
access_log_statsd_host = localhost	You can use log_statsd_* from [DEFAULT], or override them here. StatsD server. IPv4/IPv6 addresses and hostnames are supported. If a hostname resolves to an IPv4 and IPv6 address, the IPv4 address will be used.
access_log_statsd_metric_prefix =	Value will be prepended to every metric sent to the StatsD server. If not set, logging directives from [DEFAULT] without "access_" will be used.
access_log_statsd_port = 8125	Port value for the StatsD server. If not set, logging directives from [DEFAULT] without "access_" will be used.
access_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. If not set, logging directives from [DEFAULT] without "access_" will be used.
access_log_udp_host =	If not set, the UDP receiver for syslog is disabled. If not set, logging directives from [DEFAULT] without "access_" will be used.
access_log_udp_port = 514	Port value for UDP receiver, if enabled. If not set, logging directives from [DEFAULT] without "access_" will be used.
log_statsd_valid_http_methods = GET,HEAD,POST,PUT,DELETE,COPY,OPTIONS	What HTTP methods are allowed for StatsD logging (comma-sep). request methods not in this list will have "BAD_METHOD" for the <verb> portion of the metric.
reveal_sensitive_prefix = 16	<p>The X-Auth-Token is sensitive data. If revealed to an unauthorised person, they can now make requests against an account until the token expires. Set reveal_sensitive_prefix to the number of characters of the token that are logged. For example reveal_sensitive_prefix = 12 so only first 12 characters of the token are logged. Or, set to 0 to completely remove the token.</p> <div style="display: flex; align-items: flex-start;"> <div style="flex: 1;">  </div> <div style="flex: 2;"> <p>NOTE</p> <p>reveal_sensitive_prefix will not affect the value logged with access_log_headers=True.</p> </div> </div>
use = egg:swift#proxy_logging	Entry point of paste.deploy in the server.

Table 6.75. Description of configuration options for [filter-ratelimit] in proxy-server.conf

Configuration option = Default value	Description
account_blacklist = c,d	Comma separated lists of account names that will not be allowed. Returns a 497 response. r: for containers of size x, limit requests per second to r. Will limit PUT, DELETE, and POST requests to /a/c/o. container_listing_ratelimit_x = r: for containers of size x, limit listing requests per second to r. Will limit GET requests to /a/c.
account_ratelimit = 0	If set, will limit PUT and DELETE requests to /account_name/container_name. Number is in requests per second.
account_whitelist = a,b	Comma separated lists of account names that will not be rate limited.
clock_accuracy = 1000	Represents 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.
container_listing_ratelimit_0 = 100	with container_listing_ratelimit_x = r, for containers of size x, limit container GET (listing) requests per second to r. The container rate will be linearly interpolated from the values given. With the default values, a container of size 5 will get a rate of 75.
container_listing_ratelimit_10 = 50	with container_listing_ratelimit_x = r, for containers of size x, limit container GET (listing) requests per second to r. The container rate will be linearly interpolated from the values given. With the default values, a container of size 5 will get a rate of 75.
container_listing_ratelimit_50 = 20	with container_listing_ratelimit_x = r, for containers of size x, limit container GET (listing) requests per second to r. The container rate will be linearly interpolated from the values given. With the default values, a container of size 5 will get a rate of 75.
container_ratelimit_0 = 100	with container_ratelimit_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 default values, a container of size 5 will get a rate of 75.
container_ratelimit_10 = 50	with container_ratelimit_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 default values, a container of size 5 will get a rate of 75.

Configuration option = Default value	Description
container_ratelimit_50 = 20	with container_ratelimit_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 default values, a container of size 5 will get a rate of 75.
log_sleep_time_seconds = 0	To allow visibility into rate limiting set this value > 0 and all sleeps greater than the number will be logged.
max_sleep_time_seconds = 60	App will immediately return a 498 response if the necessary sleep time ever exceeds the given max_sleep_time_seconds.
rate_buffer_seconds = 5	Number of seconds the rate counter can drop and be allowed to catch up (at a faster than listed rate). A larger number will result in larger spikes in rate but better average accuracy.
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 = ratelimit	Label to use when logging.
use = egg:swift#ratelimit	Entry point of paste.deploy in the server.

Table 6.76. Description of configuration options for `[filter-slo]` in `proxy-server.conf`

Configuration option = Default value	Description
max_get_time = 86400	Time limit on GET requests (seconds)
max_manifest_segments = 1000	Maximum number of segments.
max_manifest_size = 2097152	Maximum size of segments.
min_segment_size = 1048576	Minimum size of segments.
rate_limit_after_segment = 10	Rate limit the download of large object segments after this segment is downloaded.

Configuration option = Default value	Description
rate_limit_segments_per_sec = 0	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:
use = egg:swift#slo	Entry point of paste.deploy in the server.

Table 6.77. Description of configuration options for **[filter-staticweb]** in **proxy-server.conf**

Configuration option = Default value	Description
use = egg:swift#staticweb	Entry point of paste.deploy in the server

Table 6.78. Description of configuration options for **[filter-tempauth]** in **proxy-server.conf**

Configuration option = Default value	Description
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.
auth_prefix = /auth/	The HTTP request path prefix for the auth service. Swift itself reserves anything beginning with the letter.
require_group =	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 (for example, SERVICE_require_group=blah). By default, no group is needed. Do not use .admin.
reseller_prefix = AUTH	The naming scope for the auth service.
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.

Configuration option = Default value	Description
set log_level = INFO	Log level.
set log_name = tempauth	Label to use when logging.
storage_url_scheme = default	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.
token_life = 86400	The number of seconds a token is valid.
use = egg:swift#tempauth	Entry point of paste.deploy in the server.
user_<account>_<user> = <key> [group] [group] [...] [storage_url]	<p>List of all the accounts and user you want.</p> <p>The following are example entries required for running the tests:</p> <ul style="list-style-type: none"> • user_admin_admin = admin .admin .reseller_admin • user_test2_tester2 = testing2 .admin • user_test5_tester5 = testing5 service • user_test_tester = testing .admin • user_test_tester3 = testing3

Table 6.79. Description of configuration options for [filter-tempurl] in proxy-server.conf

Configuration option = Default value	Description
incoming_allow_headers =	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_remove_headers = x-timestamp	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.
methods = GET HEAD PUT POST DELETE	HTTP methods allowed with Temporary URLs.

Configuration option = Default value	Description
outgoing_allow_headers = x-object-meta-public-*	Headers allowed as exceptions to outgoing_allow_headers. Simply a whitespace delimited list of header names and names can optionally end with '*' to indicate a prefix match.
outgoing_remove_headers = x-object-meta-*	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.
use = egg:swift#tempurl	Entry point of paste.deploy in the server.

Table 6.80. Description of configuration options for `[filter-versioned_writes]` in `proxy-server.conf`


Configuration option = Default value	Description
allow_versioned_writes = false	<p>Enables using versioned writes middleware and exposing configuration settings via HTTP GET /info.</p> <div style="background-color: #fff9c4; padding: 10px; border: 1px solid #ccc;">  <p>WARNING</p> <p>Setting this option bypasses the allow_versions option in the container configuration file, which will be eventually deprecated. For more details, see Object Versioning.</p> </div>
use = egg:swift#versioned_writes	Entry point of paste.deploy in the server.

Table 6.81. Description of configuration options for `[filter-xprofile]` in `proxy-server.conf`

Configuration option = Default value	Description
dump_interval = 5.0	The profile data will be dumped to local disk based on above naming rule in this interval (seconds).
dump_timestamp = false	Be careful, this option will enable the profiler to dump data into the file with a time stamp which means that there will be lots of files piled up in the directory.

Configuration option = Default value	Description
flush_at_shutdown = false	Clears the data when the wsgi server shutdowns.
log_filename_prefix = /tmp/log/swift/profile/default.profile	This prefix is used to combine the process ID and timestamp to name the profile data file. Make sure the executing user has permission to write into this path. Any missing path segments will be created, if necessary. When you enable profiling in more than one type of daemon, you must override it with a unique value like: /var/log/swift/profile/accountn.profile.
path = /__profile__	This is the path of the URL to access the mini web UI.
profile_module = eventlet.green.profile	This option enables you to switch profilers which inherit from the Python standard profiler. Currently, the supported value can be 'cProfile', 'eventlet.green.profile', etc.
unwind = false	Unwind the iterator of applications.
use = egg:swift#xprofile	Entry point of paste.deploy in the server.

Table 6.82. Description of configuration options for **[pipeline-main]** in **proxy-server.conf**

Configuration option = Default value	Description
pipeline = catch_errors gatekeeper healthcheck proxy-logging cache container_sync bulk tempurl ratelimit tempauth container-quotas account-quotas slo dlo versioned_writes proxy-logging proxy-server	Pipeline to use for processing operations.

Table 6.83. Description of configuration options for **[account]** in **rsyncd.conf**

Configuration option = Default value	Description
lock file = /var/lock/account.lock	No help text available for this option.
max connections = 2	No help text available for this option.
path = /srv/node	No help text available for this option.
read only = false	No help text available for this option.

Table 6.84. Description of configuration options in `rsyncd.conf`

Configuration option = Default value	Description
gid = swift	Group ID for rsyncd.
log file = /var/log/rsyncd.log	Log file for rsyncd.
pid file = /var/run/rsyncd.pid	PID file for rsyncd.
uid = swift	User ID for rsyncd.
max connections =	Maximum number of connections for rsyncd. This option should be set for each account, container, or object.
path = /srv/node	Working directory for rsyncd to use. This option should be set for each account, container, or object.
read only = false	Set read only. This option should be set for each account, container, or object.
lock file =	Lock file for rsyncd. This option should be set for each account, container, or object.

Table 6.85. Description of configuration options for `[storage-policy-0]` in `swift.conf`

Configuration option = Default value	Description
default = yes	If no policies are defined a policy with index 0 will be automatically created for backwards compatibility and given the name Policy-0. A default policy is used when creating new containers when no policy is specified in the request. If no other policies are defined the policy with index 0 will be declared the default. If multiple policies are defined you must define a policy with index 0 and you must specify a default. It is recommended you always define a section for <code>storage-policy:0</code> . Aliases are not required when defining a storage policy.
name = Policy-0	No help text available for this option.
policy_type = replication	No help text available for this option.

Table 6.86. Description of configuration options for `[swift-constraints]` in `swift.conf`

Configuration option = Default value	Description
account_listing_limit = 10000	The default (and maximum) number of items returned for an account listing request.
container_listing_limit = 10000	The default (and maximum) number of items returned for a container listing request.
extra_header_count = 0	By default the maximum number of allowed headers depends on the number of max allowed metadata settings plus a default value of 32 for regular http headers. If for some reason this is not enough (custom middleware for example) it can be increased with the extra_header_count constraint.
max_account_name_length = 256	The maximum number of bytes in the utf8 encoding of an account name.
max_container_name_length = 256	The maximum number of bytes in the utf8 encoding of a container name.
max_file_size = 5368709122	The largest normal object that can be saved in the cluster. This is also the limit on the size of each segment of a large object when using the large object manifest support. This value is set in bytes. Setting it to lower than 1MiB will cause some tests to fail. It is STRONGLY recommended to leave this value at the default ($5 * 2^{30} + 2$).
max_header_size = 8192	The max number of bytes in the utf8 encoding of each header. Using 8192 as default because eventlet use 8192 as maximum size of header line. You may need to increase this value when using identity v3 API tokens including more than 7 catalog entries. See also include_service_catalog in proxy-server.conf-sample (documented in overview_auth.rst).
max_meta_count = 90	The max number of metadata keys that can be stored on a single account, container, or object.
max_meta_name_length = 128	The max number of bytes in the utf8 encoding of the name portion of a metadata header.
max_meta_overall_size = 4096	The max number of bytes in the utf8 encoding of the metadata (keys + values).
max_meta_value_length = 256	The max number of bytes in the utf8 encoding of a metadata value.
max_object_name_length = 1024	The max number of bytes in the utf8 encoding of an object name.
valid_api_versions = v0,v1,v2	No help text available for this option.

Table 6.87. Description of configuration options for `[swift-hash]` in `swift.conf`

Configuration option = Default value	Description
swift_hash_path_prefix = changeme	A prefix used by <code>hash_path</code> 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 = changeme	A suffix used by <code>hash_path</code> 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.

6.1.2. New, Updated, and Deprecated Options in Newton for OpenStack Object Storage

There are no new, updated, and deprecated options in Newton for OpenStack Object Storage.

CHAPTER 7. TELEMETRY

7.1. TELEMETRY CONFIGURATION OPTIONS

The Telemetry service collects measurements within OpenStack. Its various agents and services are configured in the `/etc/ceilometer/ceilometer.conf` file.

7.1.1. Description of Configuration Options

The following tables provide a comprehensive list of the Telemetry configuration options.

Table 7.1. Description of API configuration options

Configuration option = Default value	Description
[DEFAULT]	
api_paste_config = api_paste.ini	(String) Configuration file for WSGI definition of API.
event_pipeline_cfg_file = event_pipeline.yaml	(String) Configuration file for event pipeline definition.
pipeline_cfg_file = pipeline.yaml	(String) Configuration file for pipeline definition.
pipeline_polling_interval = 20	(Integer) Polling interval for pipeline file configuration in seconds.
refresh_event_pipeline_cfg = False	(Boolean) Refresh Event Pipeline configuration on-the-fly.
refresh_pipeline_cfg = False	(Boolean) Refresh Pipeline configuration on-the-fly.
reserved_metadata_keys =	(List) List of metadata keys reserved for metering use. And these keys are additional to the ones included in the namespace.
reserved_metadata_length = 256	(Integer) Limit on length of reserved metadata values.
reserved_metadata_namespace = metering.	(List) List of metadata prefixes reserved for metering use.
[api]	
aodh_is_enabled = None	(Boolean) Set True to redirect alarms URLs to aodh. Default autodetection by querying keystone.
aodh_url = None	(String) The endpoint of Aodh to redirect alarms URLs to Aodh API. Default autodetection by querying keystone.

Configuration option = Default value	Description
default_api_return_limit = 100	(Integer) Default maximum number of items returned by API request.
gnocchi_is_enabled = None	(Boolean) Set True to disable resource/meter/sample URLs. Default autodetection by querying keystone.
panko_is_enabled = None	(Boolean) Set True to redirect events URLs to Panko. Default autodetection by querying keystone.
panko_url = None	(String) The endpoint of Panko to redirect events URLs to Panko API. Default autodetection by querying keystone.
pecan_debug = False	(Boolean) Toggle Pecan Debug Middleware.
[oslo_middleware]	
enable_proxy_headers_parsing = False	(Boolean) Whether the application is behind a proxy or not. This determines if the middleware should parse the headers or not.
max_request_body_size = 114688	(Integer) The maximum body size for each request, in bytes.
secure_proxy_ssl_header = X-Forwarded-Proto	(String) DEPRECATED: The HTTP Header that will be used to determine what the original request protocol scheme was, even if it was hidden by a SSL termination proxy.

Table 7.2. Description of authorization configuration options

Configuration option = Default value	Description
[service_credentials]	
auth_section = None	(Unknown) Config Section from which to load plugin specific options
auth_type = None	(Unknown) Authentication type to load
cafile = None	(String) PEM encoded Certificate Authority to use when verifying HTTPs connections.
certfile = None	(String) PEM encoded client certificate cert file
insecure = False	(Boolean) Verify HTTPS connections.

Configuration option = Default value	Description
interface = public	(String) Type of endpoint in Identity service catalog to use for communication with OpenStack services.
keyfile = None	(String) PEM encoded client certificate key file
region_name = None	(String) Region name to use for OpenStack service endpoints.
timeout = None	(Integer) Timeout value for http requests

Table 7.3. Description of collector configuration options

Configuration option = Default value	Description
[collector]	
batch_size = 1	(Integer) Number of notification messages to wait before dispatching them
batch_timeout = None	(Integer) Number of seconds to wait before dispatching samples when batch_size is not reached (None means indefinitely)
udp_address = 0.0.0.0	(String) Address to which the UDP socket is bound. Set to an empty string to disable.
udp_port = 4952	(Port number) Port to which the UDP socket is bound.
workers = 1	(Integer) Number of workers for collector service. default value is 1.
[dispatcher_file]	
backup_count = 0	(Integer) The max number of the files to keep.
file_path = None	(String) Name and the location of the file to record meters.
max_bytes = 0	(Integer) The max size of the file.

Table 7.4. Description of common configuration options

Configuration option = Default value	Description
[DEFAULT]	
batch_polled_samples = True	(Boolean) To reduce polling agent load, samples are sent to the notification agent in a batch. To gain higher throughput at the cost of load set this to False.
executor_thread_pool_size = 64	(Integer) Size of executor thread pool.
host = <your_hostname>	(String) Name of this node, which must be valid in an AMQP key. Can be an opaque identifier.
http_timeout = 600	(Integer) Timeout seconds for HTTP requests. Set it to None to disable timeout.
polling_namespaces = ['compute', 'central']	(Unknown) Polling namespace(s) to be used while resource polling
pollster_list = []	(Unknown) List of pollsters (or wildcard templates) to be used while polling
rootwrap_config = /etc/ceilometer/rootwrap.conf	(String) Path to the rootwrap configuration file touse for running commands as root
shuffle_time_before_polling_task = 0	(Integer) To reduce large requests at same time to Nova or other components from different compute agents, shuffle start time of polling task.
[compute]	
resource_update_interval = 0	(Integer) New instances will be discovered periodically based on this option (in seconds). By default, the agent discovers instances according to pipeline polling interval. If option is greater than 0, the instance list to poll will be updated based on this option's interval. Measurements relating to the instances will match intervals defined in pipeline.
workload_partitioning = False	(Boolean) Enable work-load partitioning, allowing multiple compute agents to be run simultaneously.
[coordination]	
backend_url = None	(String) The backend URL to use for distributed coordination. If left empty, per-deployment central agent and per-host compute agent won't do workload partitioning and will only function correctly if a single instance of that service is running.
check_watchers = 10.0	(Floating point) Number of seconds between checks to see if group membership has changed

Configuration option = Default value	Description
heartbeat = 1.0	(Floating point) Number of seconds between heartbeats for distributed coordination.
max_retry_interval = 30	(Integer) Maximum number of seconds between retry to join partitioning group
retry_backoff = 1	(Integer) Retry backoff factor when retrying to connect with coordination backend
[database]	
event_connection = None	(String) The connection string used to connect to the event database. (if unset, connection is used)
event_time_to_live = -1	(Integer) Number of seconds that events are kept in the database for (<= 0 means forever).
metering_connection = None	(String) The connection string used to connect to the metering database. (if unset, connection is used)
metering_time_to_live = -1	(Integer) Number of seconds that samples are kept in the database for (<= 0 means forever).
sql_expire_samples_only = False	(Boolean) Indicates if exirer expires only samples. If set true, expired samples will be deleted, but residual resource and meter definition data will remain.
[meter]	
meter_definitions_cfg_file = meters.yaml	(String) Configuration file for defining meter notifications.
[polling]	
partitioning_group_prefix = None	(String) 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	(String) Secret value for signing messages. Set value empty if signing is not required to avoid computational overhead.
[publisher_notifier]	

Configuration option = Default value	Description
event_topic = event	(String) The topic that ceilometer uses for event notifications.
metering_topic = metering	(String) The topic that ceilometer uses for metering notifications.
telemetry_driver = messagingv2	(String) The driver that ceilometer uses for metering notifications.

Table 7.5. Description of logging configuration options

Configuration option = Default value	Description
[DEFAULT]	
nova_http_log_debug = False	(Boolean) DEPRECATED: Allow novaclient's debug log output. (Use default_log_levels instead)

Table 7.6. Description of HTTP dispatcher configuration options

Configuration option = Default value	Description
[dispatcher_http]	
event_target = None	(String) 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.
target =	(String) 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
timeout = 5	(Integer) The max time in seconds to wait for a request to timeout.
verify_ssl = None	(String) The path to a server certificate or directory if the system CAs are not used or if a self-signed certificate is used. Set to False to ignore SSL cert verification.

Table 7.7. Description of Dispatchers configuration options

Configuration option = Default value	Description
[DEFAULT]	
event_dispatchers = []	(Multi-valued) Dispatchers to process event data.
meter_dispatchers = ['database']	(Multi-valued) Dispatchers to process metering data.

Table 7.8. Description of events configuration options

Configuration option = Default value	Description
[event]	
definitions_cfg_file = event_definitions.yaml	(String) Configuration file for event definitions.
drop_unmatched_notifications = False	(Boolean) Drop notifications if no event definition matches. (Otherwise, we convert them with just the default traits)
store_raw = []	(Multi-valued) Store the raw notification for select priority levels (info and/or error). By default, raw details are not captured.
[notification]	
ack_on_event_error = True	(Boolean) Acknowledge message when event persistence fails.
workers = 1	(Integer) Number of workers for notification service, default value is 1.
workload_partitioning = False	(Boolean) Enable workload partitioning, allowing multiple notification agents to be run simultaneously.

Table 7.9. Description of exchange configuration options

Configuration option = Default value	Description
[DEFAULT]	
ceilometer_control_exchange = ceilometer	(String) Exchange name for ceilometer notifications.

Configuration option = Default value	Description
cinder_control_exchange = cinder	(String) Exchange name for Cinder notifications.
dns_control_exchange = central	(String) Exchange name for DNS service notifications.
glance_control_exchange = glance	(String) Exchange name for Glance notifications.
heat_control_exchange = heat	(String) Exchange name for Heat notifications
http_control_exchanges = ['nova', 'glance', 'neutron', 'cinder']	(Multi-valued) Exchanges name to listen for notifications.
ironic_exchange = ironic	(String) Exchange name for Ironic notifications.
keystone_control_exchange = keystone	(String) Exchange name for Keystone notifications.
magnum_control_exchange = magnum	(String) Exchange name for Magnum notifications.
neutron_control_exchange = neutron	(String) Exchange name for Neutron notifications.
nova_control_exchange = nova	(String) Exchange name for Nova notifications.
sahara_control_exchange = sahara	(String) Exchange name for Data Processing notifications.
sample_source = openstack	(String) Source for samples emitted on this instance.
swift_control_exchange = swift	(String) Exchange name for Swift notifications.

Table 7.10. Description of Hyper-V configuration options

Configuration option = Default value	Description
[hyperv]	
force_volumeutils_v1 = False	(Boolean) DEPRECATED: Force V1 volume utility class

Table 7.11. Description of inspector configuration options

Configuration option = Default value	Description
[DEFAULT]	
hypervisor_inspector = libvirt	(String) Inspector to use for inspecting the hypervisor layer. Known inspectors are libvirt, hyperv, vmware, xenapi and powervm.
libvirt_type = kvm	(String) Libvirt domain type.
libvirt_uri =	(String) Override the default libvirt URI (which is dependent on libvirt_type).

Table 7.12. Description of IPMI configuration options

Configuration option = Default value	Description
[ipmi]	
node_manager_init_retry = 3	(Integer) Number of retries upon Intel Node Manager initialization failure
polling_retry = 3	(Integer) Tolerance of IPMI/NM polling failures before disable this pollster. Negative indicates retrying forever.

Table 7.13. Description of notification configuration options

Configuration option = Default value	Description
[notification]	
batch_size = 100	(Integer) Number of notification messages to wait before publishing them. Batching is advised when transformations are applied in pipeline.
batch_timeout = 5	(Integer) Number of seconds to wait before publishing samples when batch_size is not reached (None means indefinitely)
disable_non_metric_meters = True	(Boolean) 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.

Configuration option = Default value	Description
messaging_urls = []	(Multi-valued) Messaging URLs to listen for notifications. Example: rabbit://user:pass@host1:port1[,user:pass@hostN:port N]/virtual_host (DEFAULT/transport_url is used if empty). This is useful when you have dedicate messaging nodes for each service, for example, all nova notifications go to rabbit-nova:5672, while all cinder notifications go to rabbit-cinder:5672.
pipeline_processing_queues = 10	(Integer) Number of queues to parallelize workload across. This value should be larger than the number of active notification agents for optimal results. WARNING: Once set, lowering this value may result in lost data.

Table 7.14. Description of Redis configuration options

Configuration option = Default value	Description
[matchmaker_redis]	
check_timeout = 20000	(Integer) Time in ms to wait before the transaction is killed.
host = 127.0.0.1	(String) DEPRECATED: Host to locate redis. Replaced by [DEFAULT]/transport_url
password =	(String) DEPRECATED: Password for Redis server (optional). Replaced by [DEFAULT]/transport_url
port = 6379	(Port number) DEPRECATED: Use this port to connect to redis host. Replaced by [DEFAULT]/transport_url
sentinel_hosts =	(List) DEPRECATED: List of Redis Sentinel hosts (fault tolerance mode) e.g. [host:port, host1:port ...] Replaced by [DEFAULT]/transport_url
socket_timeout = 10000	(Integer) Timeout in ms on blocking socket operations
wait_timeout = 2000	(Integer) Time in ms to wait between connection attempts.

Table 7.15. Description of RADOS gateway configuration options

Configuration option = Default value	Description
[rgw_admin_credentials]	
access_key = None	(String) Access key for Radosgw Admin.
secret_key = None	(String) Secret key for Radosgw Admin.

Table 7.16. Description of service types configuration options

Configuration option = Default value	Description
[service_types]	
glance = image	(String) Glance service type.
kwapi = energy	(String) Kwapi service type.
neutron = network	(String) Neutron service type.
neutron_lbaas_version = v2	(String) Neutron load balancer version.
nova = compute	(String) Nova service type.
radosgw = object-store	(String) Radosgw service type.
swift = object-store	(String) Swift service type.

Table 7.17. Description of storage configuration options

Configuration option = Default value	Description
[storage]	
max_retries = 10	(Integer) Maximum number of connection retries during startup. Set to -1 to specify an infinite retry count.
retry_interval = 10	(Integer) Interval (in seconds) between retries of connection.

Table 7.18. Description of swift configuration options

Configuration option = Default value	Description
[DEFAULT]	
reseller_prefix = AUTH_	(String) Swift reseller prefix. Must be on par with reseller_prefix in proxy-server.conf.

Table 7.19. Description of TripleO configuration options

Configuration option = Default value	Description
[hardware]	
meter_definitions_file = snmp.yaml	(String) Configuration file for defining hardware snmp meters.
readonly_user_auth_proto = None	(String) SNMPd v3 authentication algorithm of all the nodes running in the cloud
readonly_user_name = ro_snmp_user	(String) SNMPd user name of all nodes running in the cloud.
readonly_user_password = password	(String) SNMPd v3 authentication password of all the nodes running in the cloud.
readonly_user_priv_password = None	(String) SNMPd v3 encryption password of all the nodes running in the cloud.
readonly_user_priv_proto = None	(String) SNMPd v3 encryption algorithm of all the nodes running in the cloud
url_scheme = snmp://	(String) URL scheme to use for hardware nodes.

Table 7.20. Description of VMware configuration options

Configuration option = Default value	Description
[vmware]	
api_retry_count = 10	(Integer) Number of times a VMware vSphere API may be retried.

Configuration option = Default value	Description
ca_file = None	(String) CA bundle file to use in verifying the vCenter server certificate.
host_ip =	(String) IP address of the VMware vSphere host.
host_password =	(String) Password of VMware vSphere.
host_port = 443	(Port number) Port of the VMware vSphere host.
host_username =	(String) Username of VMware vSphere.
insecure = False	(Boolean) If true, the vCenter server certificate is not verified. If false, then the default CA truststore is used for verification. This option is ignored if "ca_file" is set.
task_poll_interval = 0.5	(Floating point) Sleep time in seconds for polling an ongoing async task.
wSDL_location = None	(String) Optional vim service WSDL location e.g <a href="http://<server>/vimService.wsdl">http://<server>/vimService.wsdl . Optional over-ride to default location for bug work-arounds.

Table 7.21. Description of XenAPI configuration options

Configuration option = Default value	Description
[xenapi]	
connection_password = None	(String) Password for connection to XenServer/Xen Cloud Platform.
connection_url = None	(String) URL for connection to XenServer/Xen Cloud Platform.
connection_username = root	(String) Username for connection to XenServer/Xen Cloud Platform.

7.1.2. New, Updated, and Deprecated Options in Newton for Telemetry

Table 7.22. New options

Option = default value	(Type) Help string
[api] panko_is_enabled = None	(BoolOpt) Set True to redirect events URLs to Panko. Default autodetection by querying keystone.
[api] panko_url = None	(StrOpt) The endpoint of Panko to redirect events URLs to Panko API. Default autodetection by querying keystone.
[coordination] max_retry_interval = 30	(IntOpt) Maximum number of seconds between retry to join partitioning group
[coordination] retry_backoff = 1	(IntOpt) Retry backoff factor when retrying to connect with coordination backend
[database] sql_expire_samples_only = False	(BoolOpt) Indicates if exirer expires only samples. If set true, expired samples will be deleted, but residual resource and meter definition data will remain.
[dispatcher_http] verify_ssl = None	(StrOpt) The path to a server certificate or directory if the system CAs are not used or if a self-signed certificate is used. Set to False to ignore SSL cert verification.
[hardware] readonly_user_auth_proto = None	(StrOpt) SNMPd v3 authentication algorithm of all the nodes running in the cloud
[hardware] readonly_user_priv_password = None	(StrOpt) SNMPd v3 encryption password of all the nodes running in the cloud.
[hardware] readonly_user_priv_proto = None	(StrOpt) SNMPd v3 encryption algorithm of all the nodes running in the cloud

Table 7.23. New default values

Option	Previous default value	New default value
[DEFAULT] event_dispatchers	['database']	[]
[DEFAULT] host	localhost	<your_hostname>
[notification] batch_size	1	100
[notification] batch_timeout	None	5

Table 7.24. Deprecated options

Deprecated option	New Option
[DEFAULT] use_syslog	None
[hyperv] force_volumeutils_v1	None

7.2. TELEMETRY ALARMING CONFIGURATION OPTIONS

The Telemetry Alarming services trigger alarms when the collected metering or event data break the defined rules. Their options are configured in the `/etc/aodh/aodh.conf` file.

7.2.1. Description of Configuration Options

The following tables provide a comprehensive list of the Telemetry Alarming configuration options.

Table 7.25. Description of API configuration options

Configuration option = Default value	Description
[api]	
alarm_max_actions = -1	(Integer) Maximum count of actions for each state of an alarm, non-positive number means no limit.
enable_combination_alarms = False	(Boolean) DEPRECATED: Enable deprecated combination alarms. Combination alarms are deprecated. This option and combination alarms will be removed in Aodh 5.0.
paste_config = /usr/lib/python/site-packages/aodh/aodh/api/api-paste.ini	(String) Configuration file for WSGI definition of API.
project_alarm_quota = None	(Integer) Maximum number of alarms defined for a project.
user_alarm_quota = None	(Integer) Maximum number of alarms defined for a user.
[oslo_middleware]	
enable_proxy_headers_parsing = False	(Boolean) Whether the application is behind a proxy or not. This determines if the middleware should parse the headers or not.
max_request_body_size = 114688	(Integer) The maximum body size for each request, in bytes.

Configuration option = Default value	Description
secure_proxy_ssl_header = X-Forwarded-Proto	(String) DEPRECATED: The HTTP Header that will be used to determine what the original request protocol scheme was, even if it was hidden by a SSL termination proxy.

Table 7.26. Description of common configuration options

Configuration option = Default value	Description
[DEFAULT]	
additional_ingestion_lag = 0	(Integer) The number of seconds to extend the evaluation windows to compensate the reporting/ingestion lag.
evaluation_interval = 60	(Integer) Period of evaluation cycle, should be >= than configured pipeline interval for collection of underlying meters.
event_alarm_cache_ttl = 60	(Integer) TTL of event alarm caches, in seconds. Set to 0 to disable caching.
executor_thread_pool_size = 64	(Integer) Size of executor thread pool.
http_timeout = 600	(Integer) Timeout seconds for HTTP requests. Set it to None to disable timeout.
notifier_topic = alarming	(String) The topic that aodh uses for alarm notifier messages.
record_history = True	(Boolean) Record alarm change events.
rest_notifier_ca_bundle_certificate_path = None	(String) SSL CA_BUNDLE certificate for REST notifier
rest_notifier_certificate_file =	(String) SSL Client certificate file for REST notifier.
rest_notifier_certificate_key =	(String) SSL Client private key file for REST notifier.
rest_notifier_max_retries = 0	(Integer) Number of retries for REST notifier
rest_notifier_ssl_verify = True	(Boolean) Whether to verify the SSL Server certificate when calling alarm action.
[database]	

Configuration option = Default value	Description
alarm_history_time_to_live = -1	(Integer) Number of seconds that alarm histories are kept in the database for (<= 0 means forever).
[evaluator]	
workers = 1	(Integer) Number of workers for evaluator service. default value is 1.
[healthcheck]	
backends =	(List) Additional backends that can perform health checks and report that information back as part of a request.
detailed = False	(Boolean) Show more detailed information as part of the response
disable_by_file_path = None	(String) Check the presence of a file to determine if an application is running on a port. Used by DisableByFileHealthcheck plugin.
disable_by_file_paths =	(List) Check the presence of a file based on a port to determine if an application is running on a port. Expects a "port:path" list of strings. Used by DisableByFilesPortsHealthcheck plugin.
path = /healthcheck	(String) DEPRECATED: The path to respond to healthcheck requests on.
[listener]	
batch_size = 1	(Integer) Number of notification messages to wait before dispatching them.
batch_timeout = None	(Integer) Number of seconds to wait before dispatching samples when batch_size is not reached (None means indefinitely).
event_alarm_topic = alarm.all	(String) The topic that aodh uses for event alarm evaluation.
workers = 1	(Integer) Number of workers for listener service. default value is 1.
[notifier]	
batch_size = 1	(Integer) Number of notification messages to wait before dispatching them.

Configuration option = Default value	Description
batch_timeout = None	(Integer) Number of seconds to wait before dispatching samples when batch_size is not reached (None means indefinitely).
workers = 1	(Integer) Number of workers for notifier service. default value is 1.
[service_credentials]	
interface = public	(String) Type of endpoint in Identity service catalog to use for communication with OpenStack services.
region_name = None	(String) Region name to use for OpenStack service endpoints.
[service_types]	
zaqar = messaging	(String) Message queue service type.

Table 7.27. Description of coordination configuration options

Configuration option = Default value	Description
[coordination]	
backend_url = None	(String) The backend URL to use for distributed coordination. If left empty, per-deployment central agent and per-host compute agent won't do workload partitioning and will only function correctly if a single instance of that service is running.
check_watchers = 10.0	(Floating point) Number of seconds between checks to see if group membership has changed
heartbeat = 1.0	(Floating point) Number of seconds between heartbeats for distributed coordination.
max_retry_interval = 30	(Integer) Maximum number of seconds between retry to join partitioning group
retry_backoff = 1	(Integer) Retry backoff factor when retrying to connect with coordination backend

Table 7.28. Description of Redis configuration options

Configuration option = Default value	Description
[matchmaker_redis]	
check_timeout = 20000	(Integer) Time in ms to wait before the transaction is killed.
host = 127.0.0.1	(String) DEPRECATED: Host to locate redis. Replaced by [DEFAULT]/transport_url
password =	(String) DEPRECATED: Password for Redis server (optional). Replaced by [DEFAULT]/transport_url
port = 6379	(Port number) DEPRECATED: Use this port to connect to redis host. Replaced by [DEFAULT]/transport_url
sentinel_group_name = oslo-messaging-zeromq	(String) Redis replica set name.
sentinel_hosts =	(List) DEPRECATED: List of Redis Sentinel hosts (fault tolerance mode), e.g., [host:port, host1:port ...] Replaced by [DEFAULT]/transport_url
socket_timeout = 10000	(Integer) Timeout in ms on blocking socket operations.
wait_timeout = 2000	(Integer) Time in ms to wait between connection attempts.

7.2.2. New, Updated, and Deprecated Options in Ocata for Telemetry Alarming

Table 7.29. New options

Option = default value	(Type) Help string
[api] auth_mode = keystone	(StrOpt) Authentication mode to use. Unset to disable authentication

Table 7.30. New default values

Option	Previous default value	New default value
[api] paste_config	api_paste.ini	/usr/lib/python/site-packages/aodh/aodh/api/api-paste.ini

Table 7.31. Deprecated options

Deprecated option	New Option
[DEFAULT] <code>rpc_thread_pool_size</code>	[DEFAULT] <code>executor_thread_pool_size</code>
[DEFAULT] <code>use_syslog</code>	None