



# Red Hat OpenStack Platform 16.1

## Overcloud Parameters

Parameters for customizing the core template collection for a Red Hat OpenStack Platform overcloud



# Red Hat OpenStack Platform 16.1 Overcloud Parameters

---

Parameters for customizing the core template collection for a Red Hat OpenStack Platform overcloud

OpenStack Team  
rhos-docs@redhat.com

## Legal Notice

Copyright © 2023 Red Hat, Inc.

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

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

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

## Abstract

This guide lists common parameters that are used in the deployment of OpenStack using the Orchestration service (heat). This guide is not a comprehensive resource, and not all parameters that are listed can be used in a supported configuration. The following parameters and definitions are extracted from source code, and provided here as reference only.

---

## Table of Contents

MAKING OPEN SOURCE MORE INCLUSIVE .....	3
PROVIDING FEEDBACK ON RED HAT DOCUMENTATION .....	4
CHAPTER 1. OVERCLOUD PARAMETERS .....	5
CHAPTER 2. CORE OVERCLOUD PARAMETERS .....	6
CHAPTER 3. ROLE-BASED PARAMETERS .....	10
CHAPTER 4. DEBUG PARAMETERS .....	12
CHAPTER 5. KERNEL PARAMETERS .....	14
CHAPTER 6. BARE METAL (IRONIC) PARAMETERS .....	16
CHAPTER 7. BLOCK STORAGE (CINDER) PARAMETERS .....	22
CHAPTER 8. CEPH STORAGE PARAMETERS .....	27
CHAPTER 9. COMPUTE (NOVA) PARAMETERS .....	33
CHAPTER 10. DASHBOARD (HORIZON) PARAMETERS .....	48
CHAPTER 11. IDENTITY (KEYSTONE) PARAMETERS .....	50
CHAPTER 12. IMAGE STORAGE (GLANCE) PARAMETERS .....	54
CHAPTER 13. KEY MANAGER (BARBICAN) PARAMETERS .....	58
CHAPTER 14. LOAD BALANCER (OCTAVIA) PARAMETERS .....	61
CHAPTER 15. MESSAGING PARAMETERS .....	63
CHAPTER 16. NETWORKING (NEUTRON) PARAMETERS .....	64
CHAPTER 17. OBJECT STORAGE (SWIFT) PARAMETERS .....	70
CHAPTER 18. ORCHESTRATION (HEAT) PARAMETERS .....	72
CHAPTER 19. SHARED FILE SERVICE (MANILA) PARAMETERS .....	75
CHAPTER 20. TIME PARAMETERS .....	76
CHAPTER 21. UPGRADE PARAMETERS .....	77



## MAKING OPEN SOURCE MORE INCLUSIVE

Red Hat is committed to replacing problematic language in our code, documentation, and web properties. We are beginning with these four terms: master, slave, blacklist, and whitelist. Because of the enormity of this endeavor, these changes will be implemented gradually over several upcoming releases. For more details, see [our CTO Chris Wright's message](#).

## PROVIDING FEEDBACK ON RED HAT DOCUMENTATION

We appreciate your input on our documentation. Tell us how we can make it better.

### Using the Direct Documentation Feedback (DDF) function

Use the **Add Feedback** DDF function for direct comments on specific sentences, paragraphs, or code blocks.

1. View the documentation in the *Multi-page HTML* format.
2. Ensure that you see the **Feedback** button in the upper right corner of the document.
3. Highlight the part of text that you want to comment on.
4. Click **Add Feedback**.
5. Complete the **Add Feedback** field with your comments.
6. Optional: Add your email address so that the documentation team can contact you for clarification on your issue.
7. Click **Submit**.



## CHAPTER 1. OVERCLOUD PARAMETERS

You can modify overcloud features with overcloud parameters. To set a parameter, include the chosen parameter and its value in an environment file under the **parameter\_defaults** section and include the environment file with your **openstack overcloud deploy** command.

## CHAPTER 2. CORE OVERCLOUD PARAMETERS

You can modify general overcloud configuration with the core overcloud parameters.

Parameter	Description
<b>AddVipsToEtcHosts</b>	Set to true to append per network VIPs to <b>/etc/hosts</b> on each node. The default value is <b>True</b> .
<b>CloudDomain</b>	The DNS domain used for the hosts. This must match the <code>overcloud_domain_name</code> configured on the undercloud. The default value is <b>localdomain</b> .
<b>CloudName</b>	The DNS name of this cloud. The default value is <b>overcloud.localdomain</b> .
<b>CloudNameCtlplane</b>	The DNS name of this cloud's control plane endpoint. The default value is <b>overcloud.ctlplane.localdomain</b> .
<b>CloudNameInternal</b>	The DNS name of this cloud's internal API endpoint. The default value is <b>overcloud.internalapi.localdomain</b> .
<b>CloudNameStorage</b>	The DNS name of this cloud's storage endpoint. For example, <code>ci-overcloud.storage.tripleo.org</code> . The default value is <b>overcloud.storage.localdomain</b> .
<b>CloudNameStorageManagement</b>	The DNS name of this cloud's storage management endpoint. The default value is <b>overcloud.storagegmt.localdomain</b> .
<b>ControlFixedIPs</b>	Defines a fixed VIP for the Control Plane. Value uses the following format: <b>[[ip_address:'1.2.3.4']]</b>
<b>ControlPlaneSubnet</b>	The name of the undercloud OpenStack Networking (neutron) control plane subnet. The default value is <b>ctlplane-subnet</b> .
<b>ControlPlaneSubnetCidr</b>	The subnet CIDR of the control plane network. The parameter is automatically resolved from the <b>ctlplane</b> subnet's <b>cidr</b> attribute.
<b>DeployIdentifier</b>	Setting this to a unique value will re-run any deployment tasks that perform configuration on a OpenStack Orchestration (heat) <b>stack-update</b> .
<b>DeploymentServerBlacklist</b>	List of server hostnames to blacklist from any triggered deployments.

Parameter	Description
<b>DnsSearchDomains</b>	A list of DNS search domains to be added (in order) to resolv.conf.
<b>EndpointMapOverride</b>	Can be used to override the calculated EndpointMap.
<b>ExternalVirtualFixedIPs</b>	Control the IP allocation for the ExternalVirtualInterface port. For example, <code>[{ip_address:'1.2.3.4'}]</code> .
<b>ExtraConfig</b>	Additional hiera configuration to inject into the cluster.
<b>ExtraHostFileEntries</b>	List of extra hosts entries to be appended to /etc/hosts.
<b>GlobalConfigExtraMapData</b>	Map of extra global_config_settings data to set on each node.
<b>HypervisorNeutronPhysicalBridge</b>	An Open vSwitch bridge to create on each hypervisor. This defaults to <b>br-ex</b> , which is the same as the control plane nodes. This ensures uniform configuration of the Open vSwitch agent. Typically should not need to be changed. The default value is <b>br-ex</b> .
<b>HypervisorNeutronPublicInterface</b>	What interface to add to the <b>HypervisorNeutronPhysicalBridge</b> . The default value is <b>nic1</b> .
<b>InternalApiVirtualFixedIPs</b>	Control the IP allocation for the InternalApiVirtualInterface port. Value uses the following format: <code>[{ip_address:'1.2.3.4'}]</code>
<b>NetworkDeploymentActions</b>	OpenStack Orchestration (heat) action when to apply network configuration changes. The default value is <b>['CREATE']</b> .
<b>NeutronControlPlaneID</b>	ID or name for Control Plane ctlplane network. The default value is <b>ctlplane</b> .
<b>NeutronPhysicalBridge</b>	An OVS bridge to create for accessing external networks. The default value is <b>br-ex</b> .
<b>NeutronPublicInterface</b>	The interface to attach to the external bridge. The default value is <b>nic1</b> .

Parameter	Description
<b>NodeCreateBatchSize</b>	Maximum batch size for creating nodes. It is recommended to not exceed a batch size of 32 nodes. The default value is <b>30</b> .
<b>NovaAdditionalCell</b>	Whether this is an cell additional to the default cell. The default value is <b>False</b> .
<b>NovaLocalMetadataPerCell</b>	Indicates that the nova-metadata API service has been deployed per-cell, so that we can have better performance and data isolation in a multi-cell deployment. Users should consider the use of this configuration depending on how OpenStack Networking (neutron) is setup. If networks span cells, you might need to run nova-metadata API service globally. If your networks are segmented along cell boundaries, then you can run nova-metadata API service per cell. When running nova-metadata API service per cell, you should also configure each OpenStack Networking (neutron) metadata-agent to point to the corresponding nova-metadata API service. The default value is <b>False</b> .
<b>OVNDBsVirtualFixedIPs</b>	Control the IP allocation for the virtual IP used by OVN DBs. For example, <code>[{ip_address:'1.2.3.4'}]</code> .
<b>PublicVirtualFixedIPs</b>	Control the IP allocation for the PublicVirtualInterface port. Value uses the following format: <code>[{ip_address:'1.2.3.4'}]</code>
<b>RabbitCookieSalt</b>	Salt for the RabbitMQ cookie. Change to force the randomly generated RabbitMQ cookie to change. The default value is <b>unset</b> .
<b>RedisVirtualFixedIPs</b>	Control the IP allocation for the virtual IP used by Redis. Value uses the following format: <code>[{ip_address:'1.2.3.4'}]</code>
<b>RootStackName</b>	The name of the stack/plan.
<b>ServerMetadata</b>	Extra properties or metadata passed to OpenStack Compute (nova) for the created nodes in the overcloud. Accessible through the OpenStack Compute (nova) metadata API.
<b>StorageMgmtVirtualFixedIPs</b>	Control the IP allocation for the StorageMgmtVirtualInterface port. Value uses the following format: <code>[{ip_address:'1.2.3.4'}]</code>

Parameter	Description
<b>StorageVirtualFixedIPs</b>	Control the IP allocation for the StorageVirtualInterface port. Value uses the following format: <b>[{ip_address:'1.2.3.4'}]</b>
<b>UndercloudHostsEntries</b>	List of undercloud hosts entries to be appended to /etc/hosts. The value is populated with the HEAT_HOSTS entries on the undercloud by tripleoclient when running deploy.
<b>UpdateIdentifier</b>	Set to a previously unused value during <b>stack-update</b> triggers package update on all nodes.

## CHAPTER 3. ROLE-BASED PARAMETERS

Substitute `_ROLE_` with the name of the role. For example, for `_ROLE_Count` use `ControllerCount`.

Parameter	Description
<code>_ROLE_AnyErrorsFatal</code>	Sets the <code>any_errors_fatal</code> value when running <code>config-download</code> Ansible playbooks. The default value is <code>True</code> .
<code>_ROLE_ControlPlaneSubnet</code>	Name of the subnet on <code>ctlplane</code> network for this role. The default value is <code>ctlplane-subnet</code> .
<code>_ROLE_Count</code>	The number of nodes to deploy in a role. The default value is <code>1</code> .
<code>_ROLE_ExtraConfig</code>	Role specific additional hiera configuration to inject into the cluster.
<code>_ROLE_ExtraGroupVars</code>	Optional extra Ansible group vars.
<code>_ROLE_HostnameFormat</code>	Format for node hostnames. Note that <code>%index%</code> is translated into the index of the node (e.g 0/1/2) and <code>%stackname%</code> is replaced with the stack name (e.g <code>overcloud</code> ). The default value is <code>%stackname%_-_role_%index%</code> .
<code>_ROLE_MaxFailPercentage</code>	Sets the <code>max_fail_percentage</code> value when running <code>config-download</code> Ansible playbooks. The default value is <code>0</code> .
<code>_ROLE_NetworkDeploymentActions</code>	OpenStack Orchestration (heat) action when to apply network configuration changes.
<code>_ROLE_Parameters</code>	Optional Role Specific parameters to be provided to service.
<code>_ROLE_RemovalPolicies</code>	List of resources to be removed from the role's <code>ResourceGroup</code> when doing an update that requires removal of specific resources.
<code>_ROLE_RemovalPoliciesMode</code>	How to handle change to <code>RemovalPolicies</code> for <code>ROLE</code> <code>ResourceGroup</code> when doing an update. Default mode <code>append</code> will append to the existing blocklist and <code>update</code> would replace the blocklist. The default value is <code>append</code> .
<code>_ROLE_SchedulerHints</code>	Optional scheduler hints to pass to OpenStack Compute (nova).

Parameter	Description
<b>_ROLE_ServiceNetMap</b>	Role specific ServiceNetMap overrides, the map provided will be merged with the global ServiceNetMap when passing the ServiceNetMap to the <i>ROLE_ServiceChain</i> resource and the <i>_ROLE</i> resource group. For example: _ROLE_ServiceNetMap: NovaLibvirtNetwork: internal_api_leaf2.
<b>_ROLE_Services</b>	A list of service resources (configured in the OpenStack Orchestration (heat) resource_registry) which represent nested stacks for each service that should get installed on the <i>ROLE</i> role.

## CHAPTER 4. DEBUG PARAMETERS

These parameters allow you to set debug mode on a per-service basis. The **Debug** parameter acts as a global parameter for all services and the per-service parameters can override the effects of global parameter on individual services.

Parameter	Description
<b>BarbicanDebug</b>	Set to True to enable debugging OpenStack Key Manager (barbican) service.
<b>CinderDebug</b>	Set to True to enable debugging on OpenStack Block Storage (cinder) services.
<b>ConfigDebug</b>	Whether to run configuration management (e.g. Puppet) in debug mode. The default value is <b>False</b> .
<b>Debug</b>	Set to True to enable debugging on all services. The default value is <b>False</b> .
<b>GlanceDebug</b>	Set to True to enable debugging OpenStack Image Storage (glance) service.
<b>HeatDebug</b>	Set to True to enable debugging OpenStack Orchestration (heat) services.
<b>HorizonDebug</b>	Set to True to enable debugging OpenStack Dashboard (horizon) service.
<b>IronicDebug</b>	Set to True to enable debugging OpenStack Bare Metal (ironic) services.
<b>KeystoneDebug</b>	Set to True to enable debugging OpenStack Identity (keystone) service.
<b>ManilaDebug</b>	Set to True to enable debugging OpenStack Shared File Systems (manila) services.
<b>MemcachedDebug</b>	Set to True to enable debugging Memcached service.
<b>NeutronDebug</b>	Set to True to enable debugging OpenStack Networking (neutron) services.
<b>NovaDebug</b>	Set to True to enable debugging OpenStack Compute (nova) services.
<b>OctaviaDebug</b>	Set to True to enable debugging OpenStack Load Balancing-as-a-Service (octavia) services.



---

Parameter	Description
<b>SaharaDebug</b>	Set to True to enable debugging OpenStack Clustering (sahara) services.

## CHAPTER 5. KERNEL PARAMETERS

Parameter	Description
<b>BridgeNfCallArpTables</b>	Configures sysctl net.bridge.bridge-nf-call-arp tables key. The default value is <b>1</b> .
<b>BridgeNfCallIp6Tables</b>	Configures sysctl net.bridge.bridge-nf-call-ip6tables key. The default value is <b>1</b> .
<b>BridgeNfCallIpTables</b>	Configures sysctl net.bridge.bridge-nf-call-iptables key. The default value is <b>1</b> .
<b>ExtraKernelModules</b>	Hash of extra kernel modules to load.
<b>ExtraKernelPackages</b>	List of extra kernel related packages to install.
<b>ExtraSysctlSettings</b>	Hash of extra sysctl settings to apply.
<b>FsAioMaxNumber</b>	The kernel allocates aio memory on demand, and this number limits the number of parallel aio requests; the only drawback of a larger limit is that a malicious guest could issue parallel requests to cause the kernel to set aside memory. Set this number at least as large as $128 * (\text{number of virtual disks on the host})$ . Libvirt uses a default of 1M requests to allow 8k disks, with at most 64M of kernel memory if all disks hit an aio request at the same time. The default value is <b>0</b> .
<b>InotifyInstancesMax</b>	Configures sysctl fs.inotify.max_user_instances key. The default value is <b>1024</b> .
<b>IsolCpusList</b>	List of logical CPU ids which need to be isolated from the host processes. This input is provided to the tuned profile cpu-partitioning to configure systemd and repin interrupts (IRQ repinning).
<b>KernelArgs</b>	Kernel Args to apply to the host.
<b>KernelArgsDeferReboot</b>	When enabled, nodes will not get rebooted automatically by tripleo-kernel if KernelArgs have been changed. Operators will have to manually reboot nodes after the deployment or update process is completed. Leaving this enabled during a scale-up operation will prevent reboot on new nodes. This might cause deployment errors as the KernelArgs will not be applied. The default value is <b>False</b> .

Parameter	Description
<b>KernelDisableIPv6</b>	Configures <code>sysctl net.ipv6.{default/all}.disable_ipv6</code> keys. The default value is <b>0</b> .
<b>KernelIpForward</b>	Configures <code>net.ipv4.ip_forward</code> key. The default value is <b>1</b> .
<b>KernelIpNonLocalBind</b>	Configures <code>net.ipv{4,6}.ip_nonlocal_bind</code> key. The default value is <b>1</b> .
<b>KernelPidMax</b>	Configures <code>sysctl kernel.pid_max</code> key. The default value is <b>1048576</b> .
<b>NeighbourGcThreshold1</b>	Configures <code>sysctl net.ipv4.neigh.default.gc_thresh1</code> value. This is the minimum number of entries to keep in the ARP cache. The garbage collector will not run if there are fewer than this number of entries in the cache. The default value is <b>1024</b> .
<b>NeighbourGcThreshold2</b>	Configures <code>sysctl net.ipv4.neigh.default.gc_thresh2</code> value. This is the soft maximum number of entries to keep in the ARP cache. The garbage collector will allow the number of entries to exceed this for 5 seconds before collection will be performed. The default value is <b>2048</b> .
<b>NeighbourGcThreshold3</b>	Configures <code>sysctl net.ipv4.neigh.default.gc_thresh3</code> value. This is the hard maximum number of entries to keep in the ARP cache. The garbage collector will always run if there are more than this number of entries in the cache. The default value is <b>4096</b> .
<b>NodeRebootWaitTimeout</b>	Timeout in seconds to specify the wait time for ansible node reboot. The default value is <b>900</b> .
<b>RoleNetIpMap</b>	NO DESCRIPTION
<b>TunedProfileName</b>	Tuned Profile to apply to the host. The default value is <b>throughput-performance</b> .

## CHAPTER 6. BARE METAL (IRONIC) PARAMETERS

Parameter	Description
<b>AdditionalArchitectures</b>	List of additional architectures to enable.
<b>ApacheCertificateKeySize</b>	Override the private key size used when creating the certificate for this service.
<b>CertificateKeySize</b>	Specifies the private key size used when creating the certificate. The default value is <b>2048</b> .
<b>IPALmageURLs</b>	IPA image URLs, the format should be ["http://path/to/kernel", "http://path/to/ramdisk"].
<b>IronicAutomatedClean</b>	Enables or disables automated cleaning. Disabling automated cleaning might result in security problems and deployment failures on rebuilds. Do not set to <b>False</b> unless you understand the consequences of disabling this feature. The default value is <b>True</b> .
<b>IronicCleaningDiskErase</b>	Type of disk cleaning before and between deployments. <b>full</b> for full cleaning. <b>metadata</b> to clean only disk metadata (partition table). The default value is <b>full</b> .
<b>IronicCleaningNetwork</b>	Name or UUID of the overcloud network used for cleaning bare metal nodes. Set to <b>provisioning</b> during the initial deployment (when no networks are created yet) and change to an actual UUID in a post-deployment stack update. The default value is <b>provisioning</b> .
<b>IronicConductorGroup</b>	The name of an OpenStack Bare Metal (ironic) Conductor Group.
<b>IronicConfigureSwiftTempUrlKey</b>	Whether to configure Swift temporary URLs for use with the "direct" and "ansible" deploy interfaces. The default value is <b>True</b> .
<b>IronicCorsAllowedOrigin</b>	Indicate whether this resource may be shared with the domain received in the request "origin" header.
<b>IronicDefaultBootOption</b>	How to boot the bare metal instances. Set to <b>local</b> to use local bootloader (requires grub2 for partition images). Set to <b>netboot</b> to make the instances boot from controllers using PXE/iPXE. The default value is <b>local</b> .

Parameter	Description
<b>IronicDefaultDeployInterface</b>	Deploy interface implementation to use by default. Leave empty to use the hardware type default.
<b>IronicDefaultInspectInterface</b>	Inspect interface implementation to use by default. Leave empty to use the hardware type default.
<b>IronicDefaultNetworkInterface</b>	Network interface implementation to use by default. Set to <b>flat</b> to use one flat provider network. Set to <b>neutron</b> to make OpenStack Bare Metal (ironic) interact with the OpenStack Networking (neutron) ML2 driver to enable other network types and certain advanced networking features. Requires <b>IronicProvisioningNetwork</b> to be correctly set. The default value is <b>flat</b> .
<b>IronicDefaultRescueInterface</b>	Default rescue implementation to use. The "agent" rescue requires a compatible ramdisk to be used. The default value is <b>agent</b> .
<b>IronicDefaultResourceClass</b>	Default resource class to use for new nodes.
<b>IronicDeployLogsStorageBackend</b>	Backend to use to store ramdisk logs, either "local" or "swift". The default value is <b>local</b> .
<b>IronicDhcpv6StatefulAddressCount</b>	Number of IPv6 addresses to allocate for ports created for provisioning, cleaning, rescue or inspection on DHCPv6-stateful networks. Different stages of the chain-loading process will request addresses with different CLID/IAID. Due to non-identical identifiers multiple addresses must be reserved for the host to ensure each step of the boot process can successfully lease addresses. The default value is <b>4</b> .
<b>IronicEnabledBiosInterfaces</b>	Enabled BIOS interface implementations. Each hardware type must have at least one valid implementation enabled. The default value is <b>['no-bios']</b> .
<b>IronicEnabledBootInterfaces</b>	Enabled boot interface implementations. Each hardware type must have at least one valid implementation enabled. The default value is <b>['ipxe', 'pxe']</b> .
<b>IronicEnabledConsoleInterfaces</b>	Enabled console interface implementations. Each hardware type must have at least one valid implementation enabled. The default value is <b>['ipmitool-socat', 'no-console']</b> .

Parameter	Description
<b>IronicEnabledDeployInterfaces</b>	Enabled deploy interface implementations. Each hardware type must have at least one valid implementation enabled. The default value is <b>['iscsi', 'direct']</b> .
<b>IronicEnabledHardwareTypes</b>	Enabled OpenStack Bare Metal (ironic) hardware types. The default value is <b>['ipmi', 'redfish']</b> .
<b>IronicEnabledInspectInterfaces</b>	Enabled inspect interface implementations. Each hardware type must have at least one valid implementation enabled. The default value is <b>['no-inspect']</b> .
<b>IronicEnabledManagementInterfaces</b>	Enabled management interface implementations. Each hardware type must have at least one valid implementation enabled. The default value is <b>['ipmitool', 'noop', 'redfish']</b> .
<b>IronicEnabledNetworkInterfaces</b>	Enabled network interface implementations. Each hardware type must have at least one valid implementation enabled. The default value is <b>['flat', 'neutron']</b> .
<b>IronicEnabledPowerInterfaces</b>	Enabled power interface implementations. Each hardware type must have at least one valid implementation enabled. The default value is <b>['ipmitool', 'redfish']</b> .
<b>IronicEnabledRaidInterfaces</b>	Enabled RAID interface implementations. Each hardware type must have at least one valid implementation enabled. The default value is <b>['no-raid', 'agent']</b> .
<b>IronicEnabledRescueInterfaces</b>	Enabled rescue interface implementations. Each hardware type must have at least one valid implementation enabled. The default value is <b>['no-rescue', 'agent']</b> .
<b>IronicEnabledStorageInterfaces</b>	Enabled storage interface implementations. Each hardware type must have at least one valid implementation enabled. The default value is <b>['cinder', 'noop']</b> .
<b>IronicEnabledVendorInterfaces</b>	Enabled vendor interface implementations. Each hardware type must have at least one valid implementation enabled. The default value is <b>['ipmitool', 'no-vendor']</b> .

Parameter	Description
<b>IroniEnableStagingDrivers</b>	Whether to enable use of staging drivers. The default value is <b>False</b> .
<b>IroniForcePowerStateDuringSync</b>	Whether to force power state during sync. The default value is <b>True</b> .
<b>IroniImageDownloadSource</b>	Image delivery method for the "direct" deploy interface. Use "swift" for the Object Storage temporary URLs, use "http" for the local HTTP server (the same as for iPXE). The default value is <b>swift</b> .
<b>IroniInspectorCollectors</b>	Comma-separated list of IPA inspection collectors. The default value is <b>default,logs</b> .
<b>IroniInspectorDiscoveryDefaultDriver</b>	The default driver to use for newly discovered nodes (requires IroniInspectorEnableNodeDiscovery set to True). This driver is automatically added to enabled_drivers. The default value is <b>ipmi</b> .
<b>IroniInspectorEnableNodeDiscovery</b>	Makes ironic-inspector enroll any unknown node that PXE-boots introspection ramdisk in OpenStack Bare Metal (ironic). The default driver to use for new nodes is specified by the IroniInspectorDiscoveryDefaultDriver parameter. Introspection rules can also be used to specify it. The default value is <b>False</b> .
<b>IroniInspectorExtraProcessingHooks</b>	Comma-separated list of processing hooks to append to the default list. The default value is <b>extra_hardware,lldp_basic,local_link_connection</b> .
<b>IroniInspectorInterface</b>	Network interface on which inspection dnsmasq will listen. The default value is <b>br-ex</b> .
<b>IroniInspectorIpRange</b>	Temporary IP range that will be given to nodes during the inspection process. This should not overlap with any range that OpenStack Networking (neutron) DHCP allocates, but it has to be routeable back to <b>ironic-inspector</b> . This option has no meaningful defaults, and thus is required.
<b>IroniInspectorIPXEEnabled</b>	Whether to use iPXE for inspection. The default value is <b>True</b> .
<b>IroniInspectorKernelArgs</b>	Kernel args for the OpenStack Bare Metal (ironic) inspector. The default value is <b>ipa-inspection-dhcp-all-interfaces=1 ipa-collect-lldp=1 ipa-debug=1</b> .

Parameter	Description
<b>IronicInspectorSubnets</b>	Temporary IP ranges that will be given to nodes during the inspection process. These ranges should not overlap with any range that OpenStack Networking (neutron) DHCP provides, but they need to be routeable back to the <b>ironic-inspector</b> API. This option has no meaningful defaults and is required.
<b>IronicInspectorUseSwift</b>	Whether to use Swift for storing introspection data. The default value is <b>True</b> .
<b>IronicIpVersion</b>	The IP version that will be used for PXE booting. The default value is <b>4</b> .
<b>IronicIPXEEnabled</b>	Whether to use iPXE instead of PXE for deployment. The default value is <b>True</b> .
<b>IronicIPXEPort</b>	Port to use for serving images when iPXE is used. The default value is <b>8088</b> .
<b>IronicIPXETimeout</b>	IPXE timeout in second. Set to 0 for infinite timeout. The default value is <b>60</b> .
<b>IronicIPXEUefiSnpOnly</b>	Whether to use SNP (Simple Network Protocol) iPXE EFI, or not. When set to true <b>ipxe-snponly</b> EFI is used. The default value is <b>True</b> .
<b>IronicPassword</b>	The password for the Bare Metal service and database account.
<b>IronicPowerStateChangeTimeout</b>	Number of seconds to wait for power operations to complete, i.e., so that a baremetal node is in the desired power state. If timed out, the power operation is considered a failure. The default value is <b>60</b> .
<b>IronicProvisioningNetwork</b>	Name or UUID of the overcloud network used for provisioning of bare metal nodes if <b>IronicDefaultNetworkInterface</b> is set to <b>neutron</b> . Set to <b>provisioning</b> during the initial deployment (when no networks are created yet) and change to an actual UUID in a post-deployment stack update. The default value is <b>provisioning</b> .



Parameter	Description
<b>IronicRescuingNetwork</b>	Name or UUID of the <b>overcloud</b> network used for rescuing of bare metal nodes, if <code>IronicDefaultRescueInterface</code> is not set to "no-rescue". The default value of "provisioning" can be left during the initial deployment (when no networks are created yet) and should be changed to an actual UUID in a post-deployment stack update. The default value is <b>provisioning</b> .
<b>IronicRpcTransport</b>	The remote procedure call transport between conductor and API processes, such as a messaging broker or JSON RPC.
<b>MemcacheUseAdvancedPool</b>	Use the advanced (eventlet safe) memcached client pool. The default value is <b>True</b> .

## CHAPTER 7. BLOCK STORAGE (CINDER) PARAMETERS

Parameter	Description
<b>ApacheCertificateKeySize</b>	Override the private key size used when creating the certificate for this service.
<b>CephClusterFSID</b>	The Ceph cluster FSID. Must be a UUID.
<b>CephClusterName</b>	The Ceph cluster name. The default value is <b>ceph</b> .
<b>CephExternalMultiConfig</b>	List of maps describing extra overrides which will be applied when configuring extra external Ceph clusters. If this list is non-empty, <b>ceph-ansible</b> will run an extra count(list) times using the same parameters as the first run except each parameter within each map will override the defaults. If the following were used, the second run would configure the overcloud to also use the ceph2 cluster with all the previous parameters except <code>/etc/ceph/ceph2.conf</code> would have a <code>mon_host</code> entry containing the value of <code>external_cluster_mon_ips</code> below, and not the default <code>CephExternalMonHost</code> . Subsequent <b>ceph-ansible</b> runs are restricted to just ceph clients. <code>CephExternalMultiConfig</code> may not be used to deploy additional internal Ceph clusters within one OpenStack Orchestration (heat) stack. The map for each list should contain not <code>tripleo-heat-template</code> parameters but <b>ceph-ansible</b> parameters. - cluster: <code>ceph2</code> fsid: <code>e2cba068-5f14-4b0f-b047-acf375c0004a</code> external_cluster_mon_ips: <code>172.18.0.5,172.18.0.6,172.18.0.7</code> keys: - name: <code>"client.openstack"</code> caps: mgr: <code>"allow *"</code> mon: <code>"profile rbd"</code> osd: <code>"osd: profile rbd pool=volumes, profile rbd pool=backups, profile rbd pool=vms, profile rbd pool=images"</code> key: <code>"AQCwmeRcAAAAABAA6SQU/bGqFjlfLro5KxrB1Q="</code> mode: <code>"0600"</code> dashboard_enabled: <code>false</code> .
<b>CertificateKeySize</b>	Specifies the private key size used when creating the certificate. The default value is <b>2048</b> .
<b>CinderApiWsgiTimeout</b>	The number of seconds until a OpenStack Block Storage (cinder) API WSGI connection times out. The default value is <b>60</b> .
<b>CinderBackupOptVolumes</b>	List of optional volumes to be mounted.
<b>CinderCronDbPurgeAge</b>	Cron to move deleted instances to another table - Age. The default value is <b>30</b> .

Parameter	Description
<b>CinderCronDbPurgeDestination</b>	Cron to move deleted instances to another table - Log destination. The default value is <b>/var/log/cinder/cinder-rowsflush.log</b> .
<b>CinderCronDbPurgeHour</b>	Cron to move deleted instances to another table - Hour. The default value is <b>0</b> .
<b>CinderCronDbPurgeMaxDelay</b>	Cron to move deleted instances to another table - Max Delay. The default value is <b>3600</b> .
<b>CinderCronDbPurgeMinute</b>	Cron to move deleted instances to another table - Minute. The default value is <b>1</b> .
<b>CinderCronDbPurgeMonth</b>	Cron to move deleted instances to another table - Month. The default value is <b>*</b> .
<b>CinderCronDbPurgeMonthday</b>	Cron to move deleted instances to another table - Month Day. The default value is <b>*</b> .
<b>CinderCronDbPurgeUser</b>	Cron to move deleted instances to another table - User. The default value is <b>cinder</b> .
<b>CinderCronDbPurgeWeekday</b>	Cron to move deleted instances to another table - Week Day. The default value is <b>*</b> .
<b>CinderDefaultVolumeType</b>	The name of the OpenStack Block Storage (cinder) default volume type. The default value is <b>tripleo</b> .
<b>CinderEnableDBPurge</b>	Whether to create cron job for purging soft deleted rows in OpenStack Block Storage (cinder) database. The default value is <b>True</b> .
<b>CinderEnableiscsiBackend</b>	Whether to enable or not the Iscsi backend for OpenStack Block Storage (cinder). The default value is <b>True</b> .
<b>CinderEnableNfsBackend</b>	Whether to enable or not the NFS backend for OpenStack Block Storage (cinder). The default value is <b>False</b> .
<b>CinderEnableRbdBackend</b>	Whether to enable or not the Rbd backend for OpenStack Block Storage (cinder). The default value is <b>False</b> .

Parameter	Description
<b>CinderEtcdLocalConnect</b>	When running OpenStack Block Storage (cinder) A/A, whether to connect to Etcd via the local IP for the Etcd network. If set to true, the ip on the local node will be used. If set to false, the VIP on the Etcd network will be used instead. Defaults to false. The default value is <b>False</b> .
<b>CinderImageConversionNfsOptions</b>	NFS mount options when using an NFS share for the OpenStack Block Storage (cinder) image conversion directory. The default value is <b>_netdev,bg,intr,context=system_u:object_r:container_file_t:s0</b> .
<b>CinderImageConversionNfsShare</b>	When set, the NFS share to be used for the OpenStack Block Storage (cinder) image conversion directory.
<b>CinderISCSIAvailabilityZone</b>	The availability zone of the Iscsi OpenStack Block Storage (cinder) backend. When set, it overrides the default CinderStorageAvailabilityZone.
<b>CinderISCSIHelper</b>	The iSCSI helper to use with cinder. The default value is <b>lioadm</b> .
<b>CinderISCSIProtocol</b>	Whether to use TCP ( <i>iscsi</i> ) or iSER RDMA ( <i>iser</i> ) for iSCSI. The default value is <b>iscsi</b> .
<b>CinderLVMLoopDeviceSize</b>	The size of the loopback file used by the cinder LVM driver. The default value is <b>10280</b> .
<b>CinderNasSecureFileOperations</b>	Controls whether security enhanced NFS file operations are enabled. Valid values are <i>auto</i> , <i>true</i> or <i>false</i> . Effective when CinderEnableNfsBackend is true. The default value is <b>False</b> .
<b>CinderNasSecureFilePermissions</b>	Controls whether security enhanced NFS file permissions are enabled. Valid values are <i>auto</i> , <i>true</i> or <i>false</i> . Effective when CinderEnableNfsBackend is true. The default value is <b>False</b> .
<b>CinderNfsAvailabilityZone</b>	The availability zone of the NFS OpenStack Block Storage (cinder) backend. When set, it overrides the default CinderStorageAvailabilityZone.

Parameter	Description
<b>CinderNfsMountOptions</b>	Mount options for NFS mounts used by OpenStack Block Storage (cinder) NFS backend. Effective when CinderEnableNfsBackend is true. The default value is <b>context=system_u:object_r:container_file_t:s0</b> .
<b>CinderNfsServers</b>	NFS servers used by OpenStack Block Storage (cinder) NFS backend. Effective when CinderEnableNfsBackend is true.
<b>CinderNfsSnapshotSupport</b>	Whether to enable support for snapshots in the NFS driver. Effective when CinderEnableNfsBackend is true. The default value is <b>True</b> .
<b>CinderPassword</b>	The password for the cinder service and database account.
<b>CinderRbdAvailabilityZone</b>	The availability zone of the RBD OpenStack Block Storage (cinder) backend. When set, it overrides the default CinderStorageAvailabilityZone.
<b>CinderRbdExtraPools</b>	List of extra Ceph pools for use with RBD backends for OpenStack Block Storage (cinder). An extra OpenStack Block Storage (cinder) RBD backend driver is created for each pool in the list. This is in addition to the standard RBD backend driver associated with the CinderRbdPoolName.
<b>CinderRbdFlattenVolumeFromSnapshot</b>	Whether RBD volumes created from a snapshot should be flattened in order to remove a dependency on the snapshot. The default value is <b>False</b> .
<b>CinderRbdPoolName</b>	The Ceph pool to use for cinder volumes. The default value is <b>volumes</b> .
<b>CinderRpcResponseTimeout</b>	Cinder's RPC response timeout, in seconds. The default value is <b>60</b> .
<b>CinderStorageAvailabilityZone</b>	The OpenStack Block Storage (cinder) service's storage availability zone. The default value is <b>nova</b> .
<b>CinderVolumeCluster</b>	The cluster name used for deploying the cinder-volume service in an active-active (A/A) configuration. This configuration requires the OpenStack Block Storage (cinder) backend drivers support A/A, and the cinder-volume service not be managed by pacemaker. If these criteria are not met then the cluster name must be left blank.

Parameter	Description
<b>CinderVolumeOptEnvVars</b>	Hash of optional environment variables.
<b>CinderVolumeOptVolumes</b>	List of optional volumes to be mounted.
<b>CinderWorkers</b>	Set the number of workers for the block storage service. Note that more workers creates a larger number of processes on systems, which results in excess memory consumption. It is recommended to choose a suitable non-default value on systems with high CPU core counts. <b>0</b> sets to the OpenStack internal default, which is equal to the number of CPU cores on the node. The default value is equal to the number of vCPU cores on the physical node.
<b>DockerCinderVolumeUlimit</b>	Ulimit for OpenStack Block Storage (cinder) Volume Container. The default value is <b>['nofile=131072']</b> .
<b>EnableEtcInternalTLS</b>	Controls whether etcd and the cinder-volume service use TLS for cinder's lock manager, even when the rest of the internal API network is using TLS. The default value is <b>False</b> .
<b>MemcacheUseAdvancedPool</b>	Use the advanced (eventlet safe) memcached client pool. The default value is <b>True</b> .
<b>MultipathdEnable</b>	Whether to enable the multipath daemon. The default value is <b>False</b> .
<b>NotificationDriver</b>	Driver or drivers to handle sending notifications. The default value is <b>noop</b> .

## CHAPTER 8. CEPH STORAGE PARAMETERS

Parameter	Description
<b>AllInOneUpgrade</b>	Parameter used for pausing all ceph osds and skip transfer data check during a special case of upgrade called AllInOne. This causes workload outage. The default value is <b>False</b> .
<b>CephAnsibleDisksConfig</b>	Disks configuration settings for <b>ceph-ansible</b> . The default value is <b>{'devices': ['/dev/vdb'], 'osd_scenario': 'lvm', 'osd_objectstore': 'bluestore'}</b> .
<b>CephAnsibleEnvironmentVariables</b>	Mapping of Ansible environment variables to override defaults.
<b>CephAnsibleExtraConfig</b>	Extra vars for the <b>ceph-ansible</b> playbook.
<b>CephAnsiblePlaybook</b>	List of paths to the <b>ceph-ansible</b> playbooks to execute. If not specified, the playbook will be determined automatically depending on type of operation being performed (deploy/update/upgrade). The default value is <b>['default']</b> .
<b>CephAnsiblePlaybookVerbosity</b>	The number of <b>-v</b> , <b>-vv</b> , etc. passed to ansible-playbook command. The default value is <b>1</b> .
<b>CephAnsibleRepo</b>	The repository that should be used to install the right <b>ceph-ansible</b> package. This value can be used by tripleo-validations to double check the right <b>ceph-ansible</b> version is installed. The default value is <b>rhceph-4-tools-for-rhel-8-x86_64-rpms</b> .
<b>CephAnsibleSkipTags</b>	List of <b>ceph-ansible</b> tags to skip. The default value is <b>package-install,with_pkg</b> .
<b>CephAnsibleWarning</b>	In particular scenarios we want this validation to show the warning but don't fail because the package is installed on the system but repos are disabled. The default value is <b>True</b> .
<b>CephCertificateKeySize</b>	Override the private key size used when creating the certificate for this service.
<b>CephClientKey</b>	The Ceph client key. Currently only used for external Ceph deployments to create the openstack user keyring. Can be created with: <b>ceph-authtool --gen-print-key</b>

Parameter	Description
<b>CephClusterFSID</b>	The Ceph cluster FSID. Must be a UUID.
<b>CephClusterName</b>	The Ceph cluster name. The default value is <b>ceph</b> .
<b>CephConfigOverrides</b>	Extra configuration settings to dump into ceph.conf.
<b>CephDashboardAdminPassword</b>	Admin password for the dashboard component.
<b>CephDashboardAdminRO</b>	Parameter used to set a read-only admin user. The default value is <b>True</b> .
<b>CephDashboardAdminUser</b>	Admin user for the dashboard component. The default value is <b>admin</b> .
<b>CephEnableDashboard</b>	Parameter used to trigger the dashboard deployment. The default value is <b>False</b> .
<b>CephExternalMonHost</b>	List of externally managed Ceph Mon Host IPs. Only used for external Ceph deployments.
<b>CephExternalMultiConfig</b>	<p>List of maps describing extra overrides which will be applied when configuring extra external Ceph clusters. If this list is non-empty, <b>ceph-ansible</b> will run an extra count(list) times using the same parameters as the first run except each parameter within each map will override the defaults. If the following were used, the second run would configure the overcloud to also use the ceph2 cluster with all the previous parameters except</p> <pre>/etc/ceph/ceph2.conf would have a mon_host entry containing the value of external_cluster_mon_ips below, and not the default CephExternalMonHost. Subsequent <b>ceph-ansible</b> runs are restricted to just ceph clients. CephExternalMultiConfig may not be used to deploy additional internal Ceph clusters within one OpenStack Orchestration (heat) stack. The map for each list should contain not tripleo-heat-template parameters but <b>ceph-ansible</b> parameters. - cluster: ceph2 fsid: e2cba068-5f14-4b0f-b047-acf375c0004a external_cluster_mon_ips: 172.18.0.5,172.18.0.6,172.18.0.7 keys: - name: "client.openstack" caps: mgr: "allow *" mon: "profile rbd" osd: "osd: profile rbd pool=volumes, profile rbd pool=backups, profile rbd pool=vms, profile rbd pool=images" key: "AQCwmeRcAAAAABAA6SQU/bGqFjflLro5KxrB1Q=" mode: "0600" dashboard_enabled: false.</pre>



Parameter	Description
<b>CephExtraKeys</b>	List of maps describing extra keys which will be created on the deployed Ceph cluster. Uses <b>ceph-ansible/library/ceph_key.py</b> ansible module. Each item in the list must be in the following example format - name: "client.glance" caps: mgr: "allow *" mon: "profile rbd" osd: "profile rbd pool=images" key: "AQBRgQ9eAAAAABAav84zEilJYZPNuJ0lwn9Ndg==" mode: "0600".
<b>CephGrafanaAdminPassword</b>	Admin password for grafana component.
<b>CephIPv6</b>	Enables Ceph daemons to bind to IPv6 addresses. The default value is <b>False</b> .
<b>CephManilaClientKey</b>	The Ceph client key. Can be created with: <b>ceph- authtool --gen-print-key</b>
<b>CephMsgRSecureMode</b>	Enable Ceph msgR2 secure mode to enable on-wire encryption between Ceph daemons and also between Ceph clients and daemons. The default value is <b>False</b> .
<b>CephOsdPercentageMin</b>	The minimum percentage of Ceph OSDs which must be running and in the Ceph cluster, according to ceph osd stat, for the deployment not to fail. Used to catch deployment errors early. Set this value to 0 to disable this check. The default value is <b>66</b> .
<b>CephPoolDefaultPgNum</b>	Default placement group size to use for the RBD pools. The default value is <b>128</b> .
<b>CephPoolDefaultSize</b>	Default minimum replication for RBD copies. The default value is <b>3</b> .
<b>CephPools</b>	Override settings for one of the predefined pools or to create additional ones. Example: { <b>"volumes": { "size": 5, "pg_num": 128, "pgp_num": 128 } }</b>
<b>CephRbdMirrorConfigure</b>	Perform mirror configuration between local and remote pool. The default value is <b>True</b> .
<b>CephRbdMirrorCopyAdminKey</b>	Copy the admin key to all nodes. The default value is <b>False</b> .
<b>CephRbdMirrorPool</b>	Name of the local pool to mirror to remote cluster.

Parameter	Description
<b>CephRbdMirrorRemoteCluster</b>	The name given to the remote Ceph cluster from the local cluster. Keys reside in the <code>/etc/ceph</code> directory. The default value is <b>not-ceph</b> .
<b>CephRbdMirrorRemoteUser</b>	The rbd-mirror daemon needs a user to authenticate with the remote cluster. By default, this key should be available under <code>/etc/ceph/&lt;remote_cluster&gt;.client.&lt;remote_user&gt;.keyring</code> .
<b>CephRgwCertificateKeySize</b>	Override the private key size used when creating the certificate for this service.
<b>CephRgwClientName</b>	The client name for the RADOSGW service." The default value is <b>radosgw</b> .
<b>CephRgwKey</b>	The cephx key for the RADOSGW client. Can be created with <code>ceph-authtool --gen-print-key</code> .
<b>CephValidationDelay</b>	Interval (in seconds) in between validation checks. The default value is <b>30</b> .
<b>CephValidationRetries</b>	Number of retry attempts for Ceph validation. The default value is <b>40</b> .
<b>CertificateKeySize</b>	Specifies the private key size used when creating the certificate. The default value is <b>2048</b> .
<b>CinderBackupBackend</b>	The short name of the OpenStack Block Storage (cinder) Backup backend to use. The default value is <b>swift</b> .
<b>CinderBackupRbdPoolName</b>	Pool to use if Block Storage (cinder) Backup is enabled. The default value is <b>backups</b> .
<b>CinderEnableRbdBackend</b>	Whether to enable or not the Rbd backend for OpenStack Block Storage (cinder). The default value is <b>False</b> .
<b>CinderRbdExtraPools</b>	List of extra Ceph pools for use with RBD backends for OpenStack Block Storage (cinder). An extra OpenStack Block Storage (cinder) RBD backend driver is created for each pool in the list. This is in addition to the standard RBD backend driver associated with the <code>CinderRbdPoolName</code> .
<b>CinderRbdPoolName</b>	Pool to use for Block Storage (cinder) service. The default value is <b>volumes</b> .

Parameter	Description
<b>DeploymentServerBlacklist</b>	List of server hostnames to blacklist from any triggered deployments.
<b>GlanceBackend</b>	The short name of the OpenStack Image Storage (glance) backend to use. Set to <b>rbd</b> to use Ceph Storage. The default value is <b>swift</b> .
<b>GlanceRbdPoolName</b>	Pool to use for Image Storage (glance) service. The default value is <b>images</b> .
<b>GnocchiBackend</b>	The short name of the OpenStack Telemetry Metrics (gnocchi) backend to use. Should be one of swift, rbd, file or s3. The default value is <b>swift</b> .
<b>GnocchiRbdPoolName</b>	Pool to use for Telemetry storage. The default value is <b>metrics</b> .
<b>LocalCephAnsibleFetchDirectoryBackup</b>	Filesystem path on undercloud to persist a copy of the data from the <b>ceph-ansible</b> fetch directory. Used as an alternative to backing up the <code>fetch_directory</code> in Swift. Path must be writable and readable by the user running ansible from config-download, e.g. the <code>mistral</code> user in the <code>mistral-executor</code> container is able to read/write to <code>/var/lib/mistral/ceph_fetch</code> .
<b>ManilaCephFSCephFSAuthId</b>	The CephFS user ID for Shared Filesystem Service (manila). The default value is <b>manila</b> .
<b>ManilaCephFSDataPoolName</b>	Pool to use for file share storage. The default value is <b>manila_data</b> .
<b>ManilaCephFSDataPoolPGNum</b>	Placement group count for the CephFS data pool for file share storage. The default value is <b>128</b> .
<b>ManilaCephFSMetadataPoolName</b>	Pool to use for file share metadata storage. The default value is <b>manila_metadata</b> .
<b>ManilaCephFSMetadataPoolPGNum</b>	Placement group count for the CephFS metadata pool for file share storage. The default value is <b>128</b> .
<b>ManilaCephFSShareBackendName</b>	Backend name of the CephFS share for file share storage. The default value is <b>cephfs</b> .
<b>NodeExporterContainerImage</b>	Ceph NodeExporter container image.

Parameter	Description
<b>NovaEnableRbdBackend</b>	Whether to enable the Ceph backend for Compute (nova). The default value is <b>False</b> .
<b>NovaRbdPoolName</b>	Pool to use for Compute storage. The default value is <b>vms</b> .

## CHAPTER 9. COMPUTE (NOVA) PARAMETERS

Parameter	Description
<b>ApacheCertificateKeySize</b>	Override the private key size used when creating the certificate for this service.
<b>CertificateKeySize</b>	Specifies the private key size used when creating the certificate. The default value is <b>2048</b> .
<b>ContainerCpusetCpus</b>	Limit the specific CPUs or cores a container can use. The default value is <b>all</b> .
<b>ContainerNovaLibvirtPidsLimit</b>	Tune nova_libvirt container PID limit (set to 0 for unlimited) (defaults to 65536). The default value is <b>65536</b> .
<b>ContainerNovaLibvirtUlimit</b>	Ulimit for OpenStack Compute (nova) Libvirt Container. The default value is <b>['nofile=131072', 'nproc=126960']</b> .
<b>DockerInsecureRegistryAddress</b>	Optional. The IP Address and Port of an insecure docker namespace that will be configured in /etc/sysconfig/docker. The value can be multiple addresses separated by commas.
<b>DockerNovaComputeUlimit</b>	Ulimit for OpenStack Compute (nova) Compute Container. The default value is <b>['nofile=131072', 'memlock=67108864']</b> .
<b>DockerNovaMigrationSshdPort</b>	Port that dockerized nova migration target sshd service binds to. The default value is <b>2022</b> .
<b>EnableCache</b>	Enable caching with memcached. The default value is <b>True</b> .
<b>EnableConfigPurge</b>	Remove configuration that is not generated by the director. Used to avoid configuration remnants after upgrades. The default value is <b>False</b> .
<b>EnableInstanceHA</b>	Whether to enable an Instance Ha configurarion or not. This setup requires the Compute role to have the PacemakerRemote service added to it. The default value is <b>False</b> .
<b>EnableSQLAlchemyCollectd</b>	Set to true to enable the SQLAlchemy-collectd server plugin. The default value is <b>False</b> .
<b>ExtractedPlacementEnabled</b>	Set to True when deploying the extracted Placement service. The default value is <b>False</b> .

Parameter	Description
<b>GlanceMultistoreConfig</b>	Dictionary of settings when configuring additional glance backends. The hash key is the backend ID, and the value is a dictionary of parameter values unique to that backend. Multiple rbd backends are allowed, but cinder, file and swift backends are limited to one each. Example: # Default glance store is rbd. GlanceBackend: rbd GlanceStoreDescription: <i>Default rbd store</i> # GlanceMultistoreConfig specifies a second rbd backend, plus a cinder # backend. GlanceMultistoreConfig: rbd2_store: GlanceBackend: rbd GlanceStoreDescription: <i>Second rbd store</i> CephClusterName: ceph2 # Override CephClientUserName if this cluster uses a different # client name. CephClientUserName: client2 cinder_store: GlanceBackend: cinder GlanceStoreDescription: <i>OpenStack Block Storage (cinder) store</i> .
<b>InstanceNameTemplate</b>	Template string to be used to generate instance names. The default value is <b>instance-%08x</b> .
<b>InternalTLSNbdCAFile</b>	Specifies the CA cert to use for NBD TLS. The default value is <b>/etc/ipa/ca.crt</b> .
<b>InternalTLSQemuCAFile</b>	Specifies the CA cert to use for qemu. The default value is <b>/etc/ipa/ca.crt</b> .
<b>InternalTLSVncCAFile</b>	Specifies the CA cert to use for VNC TLS. The default value is <b>/etc/ipa/ca.crt</b> .
<b>InternalTLSVncProxyCAFile</b>	Specifies the CA cert to use for VNC TLS. The default value is <b>/etc/ipa/ca.crt</b> .
<b>KernelArgs</b>	Kernel Args to apply to the host.
<b>LibvirtCACert</b>	This specifies the CA certificate to use for TLS in libvirt. This file will be symlinked to the default CA path in libvirt, which is /etc/pki/CA/cacert.pem. Note that due to limitations GNU TLS, which is the TLS backend for libvirt, the file must be less than 65K (so we can't use the system's CA bundle). This parameter should be used if the default (which comes from the InternalTLSCAFile parameter) is not desired. The current default reflects TripleO's default CA, which is FreeIPA. It will only be used if internal TLS is enabled.
<b>LibvirtCertificateKeySize</b>	Override the private key size used when creating the certificate for this service.

Parameter	Description
<b>LibvirtEnabledPerfEvents</b>	This is a performance event list which could be used as monitor. For example: <b>cmt,mbml,mbmt</b> . Make sure you are using Red Hat Enterprise Linux 7.4 as the base and <b>libvirt</b> version is 1.3.3 or above. Also ensure you have enabled the notifications and are using hardware with a CPU that supports the <b>cmt</b> flag.
<b>LibvirtLogFilters</b>	Defines a filter to select a different logging level for a given category log outputs, as specified in <a href="https://libvirt.org/logging.html">https://libvirt.org/logging.html</a> . The default value is <b>1:libvirt 1:qemu 1:conf 1:security 3:event 3:json 3:file 3:object 1:util</b> .
<b>LibvirtNbdCACert</b>	This specifies the CA certificate to use for NBD TLS. This file will be symlinked to the default CA path, which is <code>/etc/pki/libvirt-nbd/ca-cert.pem</code> . This parameter should be used if the default (which comes from the <code>InternalTLSNbdCAFile</code> parameter) is not desired. The current default reflects TripleO's default CA, which is FreeIPA. It will only be used if internal TLS is enabled.
<b>LibvirtTLSPassword</b>	The password for the libvirt service when TLS is enabled.
<b>LibvirtTLPriority</b>	Override the compile time default TLS priority string. The default value is <b>NORMAL:-VERS-SSL3.0:-VERS-TLS-ALL:+VERS-TLS1.2</b> .
<b>LibvirtVncCACert</b>	This specifies the CA certificate to use for VNC TLS. This file will be symlinked to the default CA path, which is <code>/etc/pki/libvirt-vnc/ca-cert.pem</code> . This parameter should be used if the default (which comes from the <code>InternalTLSVncProxyCAFile</code> parameter) is not desired. The current default reflects TripleO's default CA, which is FreeIPA. It will only be used if internal TLS is enabled.
<b>LibvirtVNCClientCertificateKeySize</b>	Override the private key size used when creating the certificate for this service.
<b>LibvirtVNCServerCertificateKeySize</b>	Override the private key size used when creating the certificate for this service.

Parameter	Description
<b>MemcachedTLS</b>	Set to True to enable TLS on Memcached service. Because not all services support Memcached TLS, during the migration period, Memcached will listen on 2 ports - on the port set with MemcachedPort parameter (above) and on 11211, without TLS. The default value is <b>False</b> .
<b>MemcacheUseAdvancedPool</b>	Use the advanced (eventlet safe) memcached client pool. The default value is <b>True</b> .
<b>MigrationSshKey</b>	SSH key for migration. Expects a dictionary with keys <i>public_key</i> and <i>private_key</i> . Values should be identical to SSH public/private key files. The default value is <b>{'public_key': '', 'private_key': ''}</b> .
<b>MigrationSshPort</b>	Target port for migration over ssh. The default value is <b>2022</b> .
<b>MultipathdEnable</b>	Whether to enable the multipath daemon. The default value is <b>False</b> .
<b>MysqllPv6</b>	Enable IPv6 in MySQL. The default value is <b>False</b> .
<b>NeutronMetadataProxySharedSecret</b>	Shared secret to prevent spoofing.
<b>NeutronPhysnetNUMANodesMapping</b>	Map of physnet name as key and NUMA nodes as value. For example: <b>NeutronPhysnetNUMANodesMapping: {'foo': [0, 1], 'bar': [1]}</b> where <b>foo</b> and <b>bar</b> are physnet names and corresponding values are list of associated <b>numa_nodes</b> .
<b>NeutronTunnelNUMANodes</b>	Used to configure NUMA affinity for all tunneled networks.
<b>NotificationDriver</b>	Driver or drivers to handle sending notifications. The default value is <b>noop</b> .
<b>NovaAdditionalCell</b>	Whether this is an cell additional to the default cell. The default value is <b>False</b> .
<b>NovaAllowResizeToSameHost</b>	Allow destination machine to match source for resize. The default value is <b>False</b> .
<b>NovaApiMaxLimit</b>	Max number of objects returned per API query. The default value is <b>1000</b> .



Parameter	Description
<b>NovaAutoDisabling</b>	Max number of consecutive build failures before the nova-compute will disable itself. The default value is <b>10</b> .
<b>NovaComputeCpuDedicatedSet</b>	A comma-separated list or range of physical host CPU numbers to which processes for pinned instance CPUs can be scheduled. For example, <b>NovaComputeCpuDedicatedSet: [4-12,^8,15]</b> reserves cores from 4-12 and 15, excluding 8. If setting this option, do not set the deprecated <b>NovaVcpuPinSet</b> parameter.
<b>NovaComputeCpuSharedSet</b>	If the deprecated <b>NovaVcpuPinSet</b> option is not set, then <b>NovaComputeCpuSharedSet</b> is set to a comma-separated list or range of physical host CPU numbers used to provide vCPU inventory, determine the host CPUs that unpinned instances can be scheduled to, and determine the host CPUs that instance emulator threads should be offloaded to for instances configured with the share emulator thread policy, <b>hw:emulator_threads_policy=share</b> . If the deprecated <b>NovaVcpuPinSet</b> is set, then <b>NovaComputeCpuSharedSet</b> is set to a list or range of host CPU cores used to determine the host CPUs that instance emulator threads should be offloaded to for instances configured with the share emulator thread policy ( <b>hw:emulator_threads_policy=share</b> ). In this case, <b>NovaVcpuPinSet</b> is used to provide vCPU inventory and to determine the host CPUs that both pinned and unpinned instances can be scheduled to. For example, <b>NovaComputeCpuSharedSet: [4-12,^8,15]</b> reserves cores from 4-12 and 15, excluding 8.
<b>NovaComputeDisableIrqBalance</b>	Whether to disable irqbalance on compute nodes or not. Especially in Realtime Compute role one wants to keep it disabled. The default value is <b>False</b> .
<b>NovaComputeEnableKsm</b>	Whether to enable KSM on compute nodes or not. Especially in NFV use case one wants to keep it disabled. The default value is <b>False</b> .
<b>NovaComputeLibvirtType</b>	Libvirt domain type. Defaults to <i>kvm</i> . The default value is <b>kvm</b> .
<b>NovaComputeOptEnvVars</b>	List of optional environment variables.
<b>NovaComputeOptVolumes</b>	List of optional volumes.

Parameter	Description
<b>NovaCPUAllocationRatio</b>	Virtual CPU to physical CPU allocation ratio. The default value is <b>0.0</b> .
<b>NovaCronArchiveDeleteAllCells</b>	Archive deleted instances from all cells. The default value is <b>True</b> .
<b>NovaCronArchiveDeleteRowsAge</b>	Cron to archive deleted instances - Age. This will define the retention policy when archiving the deleted instances entries in days. 0 means archive data older than today in shadow tables. The default value is <b>90</b> . The default value is <b>90</b> .
<b>NovaCronArchiveDeleteRowsDestination</b>	Cron to move deleted instances to another table - Log destination. The default value is <b>/var/log/nova/nova-rowsflush.log</b> .
<b>NovaCronArchiveDeleteRowsHour</b>	Cron to move deleted instances to another table - Hour. The default value is <b>0</b> .
<b>NovaCronArchiveDeleteRowsMaxDelay</b>	Cron to move deleted instances to another table - Max Delay. The default value is <b>3600</b> .
<b>NovaCronArchiveDeleteRowsMaxRows</b>	Cron to move deleted instances to another table - Max Rows. The default value is <b>1000</b> .
<b>NovaCronArchiveDeleteRowsMinute</b>	Cron to move deleted instances to another table - Minute. The default value is <b>1</b> .
<b>NovaCronArchiveDeleteRowsMonth</b>	Cron to move deleted instances to another table - Month. The default value is <b>*</b> .
<b>NovaCronArchiveDeleteRowsMonthday</b>	Cron to move deleted instances to another table - Month Day. The default value is <b>*</b> .
<b>NovaCronArchiveDeleteRowsPurge</b>	Purge shadow tables immediately after scheduled archiving. The default value is <b>False</b> .
<b>NovaCronArchiveDeleteRowsUntilComplete</b>	Cron to move deleted instances to another table - Until complete. The default value is <b>True</b> .
<b>NovaCronArchiveDeleteRowsUser</b>	Cron to move deleted instances to another table - User. The default value is <b>nova</b> .
<b>NovaCronArchiveDeleteRowsWeekday</b>	Cron to move deleted instances to another table - Week Day. The default value is <b>*</b> .

Parameter	Description
<b>NovaCronPurgeShadowTablesAge</b>	Cron to purge shadow tables - Age This will define the retention policy when purging the shadow tables in days. 0 means, purge data older than today in shadow tables. The default value is <b>14</b> .
<b>NovaCronPurgeShadowTablesAllCells</b>	Cron to purge shadow tables - All cells. The default value is <b>True</b> .
<b>NovaCronPurgeShadowTablesDestination</b>	Cron to purge shadow tables - Log destination. The default value is <b>/var/log/nova/nova-rowspurge.log</b> .
<b>NovaCronPurgeShadowTablesHour</b>	Cron to purge shadow tables - Hour. The default value is <b>5</b> .
<b>NovaCronPurgeShadowTablesMaxDelay</b>	Cron to purge shadow tables - Max Delay. The default value is <b>3600</b> .
<b>NovaCronPurgeShadowTablesMinute</b>	Cron to purge shadow tables - Minute. The default value is <b>0</b> .
<b>NovaCronPurgeShadowTablesMonth</b>	Cron to purge shadow tables - Month. The default value is <b>*</b> .
<b>NovaCronPurgeShadowTablesMonthday</b>	Cron to purge shadow tables - Month Day. The default value is <b>*</b> .
<b>NovaCronPurgeShadowTablesUser</b>	Cron to purge shadow tables - User. The default value is <b>nova</b> .
<b>NovaCronPurgeShadowTablesVerbose</b>	Cron to purge shadow tables - Verbose. The default value is <b>False</b> .
<b>NovaCronPurgeShadowTablesWeekday</b>	Cron to purge shadow tables - Week Day. The default value is <b>*</b> .
<b>NovaCrossAZAttach</b>	Whether instances can attach cinder volumes from a different availability zone. The default value is <b>True</b> .
<b>NovaDbSyncTimeout</b>	Timeout for OpenStack Compute (nova) database synchronization in seconds. The default value is <b>300</b> .
<b>NovaDefaultFloatingPool</b>	Default pool for floating IP addresses. The default value is <b>public</b> .

Parameter	Description
<b>NovaDisableImageDownloadToRbd</b>	Refuse to boot an instance if it would require downloading from glance and uploading to ceph instead of a COW clone. The default value is <b>False</b> .
<b>NovaDiskAllocationRatio</b>	Virtual disk to physical disk allocation ratio. The default value is <b>0.0</b> .
<b>NovaEnableDBArchive</b>	Whether to create cron job for archiving soft deleted rows in OpenStack Compute (nova) database. The default value is <b>True</b> .
<b>NovaEnableDBPurge</b>	Whether to create cron job for purging soft deleted rows in OpenStack Compute (nova) database. The default value is <b>True</b> .
<b>NovaEnableNUMALiveMigration</b>	Whether to enable or not the live migration for NUMA topology instances. The default value is <b>False</b> .
<b>NovaGlanceEnableRbdDownload</b>	Enable download of OpenStack Image Storage (glance) images directly via RBD. The default value is <b>False</b> .
<b>NovaGlanceRbdDownloadMultistoreID</b>	The hash key, which is the backend ID, of the GlanceMultistoreConfig to be used for the role where NovaGlanceEnableRbdDownload is enabled and defaults should be overridden. If CephClientUserName or GlanceRbdPoolName are not set in the GlanceMultistoreConfig, the global values of those parameters will be used.
<b>NovaHWMachineType</b>	Specifies the default machine type for each host architecture. Red Hat recommends setting the default to the lowest RHEL minor release in your environment, for backwards compatibility during live migration. The default value is <b>x86_64=pc-i440fx-rhel7.6.0,aarch64=virt-rhel7.6.0,ppc64=pseries-rhel7.6.0,ppc64le=pseries-rhel7.6.0</b> .
<b>NovalImageCacheTTL</b>	Time in seconds that nova compute should continue caching an image once it is no longer used by any instances on the host. The default value is <b>86400</b> .
<b>NovalImageTypeExcludeList</b>	List of image formats that should not be advertised as supported by the compute service.

Parameter	Description
<b>NovaLibvirtCPUMode</b>	The libvirt CPU mode to configure. Defaults to <i>host-model</i> if <i>virt_type</i> is set to <i>kvm</i> , otherwise defaults to <i>none</i> . The default value is <b>host-model</b> .
<b>NovaLibvirtCPUModelExtraFlags</b>	This allows specifying granular CPU feature flags when specifying CPU models. Only has effect if <i>cpu_mode</i> is not set to <i>none</i> .
<b>NovaLibvirtCPUModels</b>	The named libvirt CPU model (see names listed in <i>/usr/share/libvirt/cpu_map.xml</i> ). Only has effect if <i>cpu_mode</i> ="custom" and <i>virt_type</i> ="kvm qemu".
<b>NovaLibvirtFileBackedMemory</b>	Available capacity in MiB for file-backed memory. When configured, the <b>NovaRAMAllocationRatio</b> and <b>NovaReservedHostMemory</b> parameters must be set to 0. The default value is <b>0</b> .
<b>NovaLibvirtMaxQueues</b>	Add parameter to configure the libvirt <i>max_queues</i> . The maximum number of virtio queue pairs that can be enabled when creating a multiqueue guest. The number of virtio queues allocated will be the lesser of the CPUs requested by the guest and the max value defined. Default 0 corresponds to not set. The default value is <b>0</b> .
<b>NovaLibvirtMemStatsPeriodSeconds</b>	A number of seconds to memory usage statistics period, zero or negative value mean to disable memory usage statistics. The default value is <b>10</b> .
<b>NovaLibvirtNumPciePorts</b>	Set <b>num_pcie_ports</b> to specify the number of PCIe ports an instance will get. Libvirt allows a custom number of PCIe ports ( <i>pcie-root-port</i> controllers) a target instance will get. Some will be used by default, rest will be available for hotplug use. The default value is <b>16</b> .
<b>NovaLibvirtOptVolumes</b>	List of optional volumes to be mounted.
<b>NovaLibvirtRxQueueSize</b>	Virtio-net RX queue size. Valid values are 256, 512, 1024. The default value is <b>512</b> .
<b>NovaLibvirtTxQueueSize</b>	Virtio-net TX queue size. Valid values are 256, 512, 1024. The default value is <b>512</b> .
<b>NovaLibvirtVolumeUseMultipath</b>	Whether to enable or not the multipath connection of the volumes. The default value is <b>False</b> .

Parameter	Description
<b>NovaLiveMigrationPermitAutoConverge</b>	Defaults to "True" to slow down the instance CPU until the memory copy process is faster than the instance's memory writes when the migration performance is slow and might not complete. Auto converge will only be used if this flag is set to True and post copy is not permitted or post copy is unavailable due to the version of libvirt and QEMU. The default value is <b>True</b> .
<b>NovaLiveMigrationPermitPostCopy</b>	If "True" activates the instance on the destination node before migration is complete, and to set an upper bound on the memory that needs to be transferred. Post copy gets enabled per default if the compute roles is not a realtime role or disabled by this parameter. The default value is <b>True</b> .
<b>NovaLiveMigrationWaitForVIFPlug</b>	Whether to wait for <b>network-vif-plugged</b> events before starting guest transfer. The default value is <b>True</b> .
<b>NovaLocalMetadataPerCell</b>	Indicates that the nova-metadata API service has been deployed per-cell, so that we can have better performance and data isolation in a multi-cell deployment. Users should consider the use of this configuration depending on how OpenStack Networking (neutron) is setup. If networks span cells, you might need to run nova-metadata API service globally. If your networks are segmented along cell boundaries, then you can run nova-metadata API service per cell. When running nova-metadata API service per cell, you should also configure each OpenStack Networking (neutron) metadata-agent to point to the corresponding nova-metadata API service. The default value is <b>False</b> .
<b>NovaNfsEnabled</b>	Whether to enable or not the NFS backend for OpenStack Compute (nova). The default value is <b>False</b> .
<b>NovaNfsOptions</b>	NFS mount options for nova storage (when NovaNfsEnabled is true). The default value is <b>context=system_u:object_r:nfs_t:s0</b> .
<b>NovaNfsShare</b>	NFS share to mount for nova storage (when NovaNfsEnabled is true).

Parameter	Description
<b>NovaNfsVersion</b>	NFS version used for nova storage (when <code>NovaNfsEnabled</code> is true). Since NFSv3 does not support full locking a NFSv4 version need to be used. To not break current installations the default is the previous hard coded version 4. The default value is <b>4</b> .
<b>NovaOVSBridge</b>	Name of integration bridge used by Open vSwitch. The default value is <b>br-int</b> .
<b>NovaOVSDDBConnection</b>	OVS DB connection string to used by OpenStack Compute (nova).
<b>NovaPassword</b>	The password for the OpenStack Compute (nova) service and database account.
<b>NovaPCIPassthrough</b>	YAML list of PCI passthrough whitelist parameters.
<b>NovaPMEMMappings</b>	PMEM namespace mappings as backend for vPMEM feature. This parameter sets Nova's <b>pmem_namespaces</b> configuration options. PMEM namespaces needs to be create manually or with conjunction with <b>NovaPMEMNamespaces</b> parameter. Requires format: \$LABEL:\$NSNAME[ \$NSNAME] [\$LABEL:\$NSNAME[ \$NSNAME]].
<b>NovaPMEMNamespaces</b>	Creates PMEM namespaces on the host server using <b>ndctl</b> tool through Ansible. Requires format: \$SIZE:\$NSNAME[\$SIZE:\$NSNAME...]. \$SIZE supports the suffixes "k" or "K" for KiB, "m" or "M" for MiB, "g" or "G" for GiB and "t" or "T" for TiB. NOTE: This requires properly configured NVDIMM regions and enough space for requested namespaces.
<b>NovaRAMAllocationRatio</b>	Virtual RAM to physical RAM allocation ratio. The default value is <b>1.0</b> .
<b>NovaReservedHostMemory</b>	Reserved RAM for host processes. The default value is <b>4096</b> .
<b>NovaReservedHugePages</b>	A list of valid key=value which reflect NUMA node ID, page size (Default unit is KiB) and number of pages to be reserved. Example - NovaReservedHugePages: ["node:0,size:2048,count:64","node:1,size:1GB,count:1"] will reserve on NUMA node 0 64 pages of 2MiB and on NUMA node 1, 1 page of 1GiB.

Parameter	Description
<b>NovaResumeGuestsShutdownTimeout</b>	Number of seconds we're willing to wait for a guest to shut down. If this is 0, then there is no time out (use with caution, as guests might not respond to a shutdown request). The default value is 300 seconds (5 minutes). The default value is <b>300</b> .
<b>NovaResumeGuestsStateOnHostBoot</b>	Whether to start running instance on compute host reboot. The default value is <b>False</b> .
<b>NovaSchedulerAvailableFilters</b>	List of available filters for OpenStack Compute (nova) to use to filter nodes.
<b>NovaSchedulerDefaultFilters</b>	An array of filters OpenStack Compute (nova) uses to filter a node. OpenStack Compute applies these filters in the order they are listed. Place your most restrictive filters first to make the filtering process more efficient.
<b>NovaSchedulerDiscoverHostsInCellsInterval</b>	This value controls how often (in seconds) the scheduler should attempt to discover new hosts that have been added to cells. The default value of -1 disables the periodic task completely. It is recommended to set this parameter for deployments using OpenStack Bare Metal (ironic). The default value is <b>-1</b> .
<b>NovaSchedulerEnableIsolatedAggregateFiltering</b>	This setting allows the scheduler to restrict hosts in aggregates based on matching required traits in the aggregate metadata and the instance flavor/image. If an aggregate is configured with a property with key trait:\$TRAIT_NAME and value required, the instance flavor extra_specs and/or image metadata must also contain trait:\$TRAIT_NAME=required to be eligible to be scheduled to hosts in that aggregate. The default value is <b>False</b> .
<b>NovaSchedulerLimitTenantsToPlacementAggregate</b>	This value allows to have tenant isolation with placement. It ensures hosts in tenant-isolated host aggregate and availability zones will only be available to specific set of tenants. The default value is <b>False</b> .
<b>NovaSchedulerMaxAttempts</b>	Maximum number of attempts the scheduler will make when deploying the instance. You should keep it greater or equal to the number of bare metal nodes you expect to deploy at once to work around potential race conditions when scheduling. The default value is <b>3</b> .



Parameter	Description
<b>NovaSchedulerPlacementAggregateRequiredForTenants</b>	This setting, when <b>NovaSchedulerLimitTenantsToPlacementAggregate</b> is true, controls whether or not a tenant with no aggregate affinity will be allowed to schedule to any available node. If aggregates are used to limit some tenants but not all, then this should be False. If all tenants should be confined via aggregate, then this should be True. The default value is <b>False</b> .
<b>NovaSchedulerQueryImageType</b>	This setting causes the scheduler to ask placement only for compute hosts that support the disk_format of the image used in the request. The default value is <b>True</b> .
<b>NovaSchedulerQueryPlacementForAvailabilityZone</b>	This setting allows the scheduler to look up a host aggregate with metadata key of availability zone set to the value provided by incoming request, and request result from placement be limited to that aggregate. The default value is <b>False</b> .
<b>NovaSchedulerWorkers</b>	Number of workers for OpenStack Compute (nova) Scheduler services. The default value is <b>0</b> .
<b>NovaStatedirOwnershipSkip</b>	List of paths relative to nova_statedir to ignore when recursively setting the ownership and selinux context. The default value is <b>['triliovault-mounts']</b> .
<b>NovaSyncPowerStateInterval</b>	Interval to sync power states between the database and the hypervisor. Set to -1 to disable. Setting this to 0 will run at the default rate. The default value is <b>0</b> .
<b>NovaVcpuPinSet</b>	(Deprecated) A list or range of physical CPU cores to reserve for virtual machine processes. For example, <b>NovaVcpuPinSet: [4-12,^8]</b> reserves cores from 4-12 excluding 8. This parameter has been deprecated. Use <b>NovaComputeCpuDedicatedSet</b> and <b>NovaComputeCpuSharedSet</b> instead.
<b>NovaVNCCertificateKeySize</b>	Override the private key size used when creating the certificate for this service.
<b>NovaVNCProxySSLCiphers</b>	OpenSSL cipher preference string that specifies what ciphers to allow for TLS connections from clients. See the man page for the OpenSSL <i>ciphers</i> command for details of the cipher preference string format and allowed values.

Parameter	Description
<b>NovaVNCProxySSLMinimumVersion</b>	Minimum allowed SSL/TLS protocol version. Valid values are <i>default</i> , <i>tlsv1_1</i> , <i>tlsv1_2</i> , and <i>tlsv1_3</i> . A value of <i>default</i> will use the underlying system OpenSSL defaults. The default value is <b>default</b> .
<b>NovaWorkers</b>	Number of workers for the Compute's Conductor service. Note that more workers creates a larger number of processes on systems, which results in excess memory consumption. <b>0</b> sets to the OpenStack internal default, which is equal to the number of CPU cores on the node. The default value is <b>0</b> .
<b>OvsDpdkSocketMemory</b>	Sets the amount of hugepage memory to assign per NUMA node. It is recommended to use the socket closest to the PCIe slot used for the desired DPDK NIC. The format should be in "<socket 0 mem>, <socket 1 mem>, <socket n mem>", where the value is specified in MB. For example: "1024,0".
<b>PlacementAPIInterface</b>	Endpoint interface to be used for the placement API. The default value is <b>internal</b> .
<b>PlacementPassword</b>	The password for the Placement service and database account.
<b>QemuCACert</b>	This specifies the CA certificate to use for qemu. This file will be symlinked to the default CA path, which is <code>/etc/pki/qemu/ca-cert.pem</code> . This parameter should be used if the default (which comes from the <code>InternalTLSQemuCAFile</code> parameter) is not desired. The current default reflects TripleO's default CA, which is FreeIPA. It will only be used if internal TLS is enabled.
<b>QemuClientCertificateKeySize</b>	Override the private key size used when creating the certificate for this service.
<b>QemuDefaultTLSVerify</b>	Whether to enable or disable TLS client certificate verification. Enabling this option will reject any client who does not have a certificate signed by the CA in <code>/etc/pki/qemu/ca-cert.pem</code> . The default value is <b>False</b> .
<b>QemuMemoryBackingDir</b>	Directory used for memoryBacking source if configured as file. NOTE: big files will be stored here.
<b>QemuServerCertificateKeySize</b>	Override the private key size used when creating the certificate for this service.

Parameter	Description
<b>StackUpdateType</b>	Type of update, to differentiate between UPGRADE and UPDATE cases when StackAction is UPDATE (both are the same stack action).
<b>UpgradeLevelNovaCompute</b>	OpenStack Compute upgrade level.
<b>UseTLSTransportForNbd</b>	If set to true and if EnableInternalTLS is enabled, it will enable TLS transport for libvirt NBD and configure the relevant keys for libvirt. The default value is <b>True</b> .
<b>UseTLSTransportForVnc</b>	If set to true and if EnableInternalTLS is enabled, it will enable TLS transport for libvirt VNC and configure the relevant keys for libvirt. The default value is <b>True</b> .
<b>VerifyGlanceSignatures</b>	Whether to verify image signatures. The default value is <b>False</b> .
<b>VhostuserSocketGroup</b>	The vhost-user socket directory group name. Defaults to <i>qemu</i> . When vhostuser mode is <i>dpdkvhostuserclient</i> (which is the default mode), the vhost socket is created by qemu. The default value is <b>qemu</b> .

## CHAPTER 10. DASHBOARD (HORIZON) PARAMETERS

Parameter	Description
<b>HorizonAllowedHosts</b>	A list of IP/Hostname for the server OpenStack Dashboard (horizon) is running on. Used for header checks. The default value is <code>*</code> .
<b>HorizonCustomizationModule</b>	OpenStack Dashboard (horizon) has a global overrides mechanism available to perform customizations.
<b>HorizonHelpURL</b>	On top of dashboard there is a Help button. This button could be used to re-direct user to vendor documentation or dedicated help portal. The default value is <a href="https://access.redhat.com/documentation/en-us/red_hat_openstack_platform">https://access.redhat.com/documentation/en-us/red_hat_openstack_platform</a> .
<b>HorizonPasswordValidator</b>	Regex for password validation.
<b>HorizonPasswordValidatorHelp</b>	Help text for password validation.
<b>HorizonSecret</b>	Secret key for the webserver.
<b>HorizonSecureCookies</b>	Set <code>CSRF_COOKIE_SECURE</code> / <code>SESSION_COOKIE_SECURE</code> in OpenStack Dashboard (horizon). The default value is <b>False</b> .
<b>HorizonVhostExtraParams</b>	Extra parameters for OpenStack Dashboard (horizon) vhost configuration. The default value is <code>{'add_listen': True, 'priority': 10, 'access_log_format': '%a %l %u %t \"%r\" %&gt;s %b \"%%{{Referer}}i\" \"%%{{User-Agent}}i\"', 'options': ['FollowSymLinks', 'MultiViews']}</code> .
<b>MemcachedIPv6</b>	Enable IPv6 features in Memcached. The default value is <b>False</b> .
<b>TimeZone</b>	The timezone to be set on the overcloud. The default value is <b>UTC</b> .
<b>WebSSOChoices</b>	Specifies the list of SSO authentication choices to present. Each item is a list of an SSO choice identifier and a display message. The default value is <code>[['OIDC', 'OpenID Connect']]</code> .
<b>WebSSOEnable</b>	Enable support for Web Single Sign-On. The default value is <b>False</b> .

Parameter	Description
<b>WebSSOIDPMapping</b>	Specifies a mapping from SSO authentication choice to identity provider and protocol. The identity provider and protocol names must match the resources defined in keystone. The default value is <b>{'OIDC': ['myidp', 'openid']}</b> .
<b>WebSSOInitialChoice</b>	The initial authentication choice to select by default. The default value is <b>OIDC</b> .

## CHAPTER 11. IDENTITY (KEYSTONE) PARAMETERS

Parameter	Description
<b>AdminEmail</b>	The email for the OpenStack Identity (keystone) admin account. The default value is <b>admin@example.com</b> .
<b>AdminToken</b>	The OpenStack Identity (keystone) secret and database password.
<b>ApacheCertificateKeySize</b>	Override the private key size used when creating the certificate for this service.
<b>CertificateKeySize</b>	Specifies the private key size used when creating the certificate. The default value is <b>2048</b> .
<b>EnableCache</b>	Enable caching with memcached. The default value is <b>True</b> .
<b>EnablePublicTLS</b>	Whether to enable TLS on the public interface or not. The default value is <b>True</b> .
<b>KeystoneAuthMethods</b>	A list of methods used for authentication.
<b>KeystoneChangePasswordUponFirstUse</b>	Enabling this option requires users to change their password when the user is created, or upon administrative reset.
<b>KeystoneCorsAllowedOrigin</b>	Indicate whether this resource may be shared with the domain received in the request "origin" header.
<b>KeystoneCredential0</b>	The first OpenStack Identity (keystone) credential key. Must be a valid key.
<b>KeystoneCredential1</b>	The second OpenStack Identity (keystone) credential key. Must be a valid key.
<b>KeystoneDisableUserAccountDaysInactive</b>	The maximum number of days a user can go without authenticating before being considered "inactive" and automatically disabled (locked).
<b>KeystoneEnableMember</b>	Create the <i>member</i> role, useful for undercloud deployment. The default value is <b>False</b> .
<b>KeystoneFederationEnable</b>	Enable support for federated authentication. The default value is <b>False</b> .

Parameter	Description
<b>KeystoneFernetKeys</b>	Mapping containing OpenStack Identity (keystone) fernet keys and their paths.
<b>KeystoneFernetMaxActiveKeys</b>	The maximum active keys in the OpenStack Identity (keystone) fernet key repository. The default value is <b>5</b> .
<b>KeystoneLDAPBackendConfigs</b>	Hash containing the configurations for the LDAP backends configured in keystone.
<b>KeystoneLDAPDomainEnable</b>	Trigger to call <code>ldap_backend puppet keystone define</code> . The default value is <b>False</b> .
<b>KeystoneLockoutDuration</b>	The number of seconds a user account will be locked when the maximum number of failed authentication attempts (as specified by <code>KeystoneLockoutFailureAttempts</code> ) is exceeded.
<b>KeystoneLockoutFailureAttempts</b>	The maximum number of times that a user can fail to authenticate before the user account is locked for the number of seconds specified by <code>KeystoneLockoutDuration</code> .
<b>KeystoneMinimumPasswordAge</b>	The number of days that a password must be used before the user can change it. This prevents users from changing their passwords immediately in order to wipe out their password history and reuse an old password.
<b>KeystoneNotificationFormat</b>	The OpenStack Identity (keystone) notification format. The default value is <b>basic</b> .
<b>KeystoneNotificationTopics</b>	OpenStack Identity (keystone) notification topics to enable.
<b>KeystoneOpenIdcClientId</b>	The client ID to use when handshaking with your OpenID Connect provider.
<b>KeystoneOpenIdcClientSecret</b>	The client secret to use when handshaking with your OpenID Connect provider.
<b>KeystoneOpenIdcCryptoPassphrase</b>	Passphrase to use when encrypting data for OpenID Connect handshake. The default value is <b>openstack</b> .
<b>KeystoneOpenIdcEnable</b>	Enable support for OpenIDC federation. The default value is <b>False</b> .

Parameter	Description
<b>KeystoneOpenIdcEnableOAuth</b>	Enable OAuth 2.0 integration. The default value is <b>False</b> .
<b>KeystoneOpenIdcIdpName</b>	The name associated with the IdP in OpenStack Identity (keystone).
<b>KeystoneOpenIdcIntrospectionEndpoint</b>	OAuth 2.0 introspection endpoint for mod_auth_openidc.
<b>KeystoneOpenIdcProviderMetadataUrl</b>	The url that points to your OpenID Connect provider metadata.
<b>KeystoneOpenIdcRemotIdAttribute</b>	Attribute to be used to obtain the entity ID of the Identity Provider from the environment. The default value is <b>HTTP_OIDC_ISS</b> .
<b>KeystoneOpenIdcResponseType</b>	Response type to be expected from the OpenID Connect provider. The default value is <b>id_token</b> .
<b>KeystonePasswordExpiresDays</b>	The number of days for which a password will be considered valid before requiring it to be changed.
<b>KeystonePasswordRegex</b>	The regular expression used to validate password strength requirements.
<b>KeystonePasswordRegexDescription</b>	Describe your password regular expression here in language for humans.
<b>KeystoneSSLCertificate</b>	OpenStack Identity (keystone) certificate for verifying token validity.
<b>KeystoneSSLCertificateKey</b>	OpenStack Identity (keystone) key for signing tokens.
<b>KeystoneTokenProvider</b>	The OpenStack Identity (keystone) token format. The default value is <b>fernet</b> .
<b>KeystoneTrustedDashboards</b>	A list of dashboard URLs trusted for single sign-on.
<b>KeystoneUniqueLastPasswordCount</b>	This controls the number of previous user password iterations to keep in history, in order to enforce that newly created passwords are unique.



Parameter	Description
<b>KeystoneWorkers</b>	Set the number of workers for the OpenStack Identity (keystone) service. Note that more workers creates a larger number of processes on systems, which results in excess memory consumption. It is recommended to choose a suitable non-default value on systems with high CPU core counts. <b>0</b> sets to the OpenStack internal default, which is equal to the number of CPU cores on the node. The default value is equal to the number of vCPU cores on the physical node.
<b>ManageKeystoneFernetKeys</b>	Whether director should manage the OpenStack Identity (keystone) fernet keys or not. If set to True, the fernet keys will get the values from the saved keys repository in OpenStack Workflow (mistral) from the <b>KeystoneFernetKeys</b> variable. If set to false, only the stack creation initializes the keys, but subsequent updates will not touch them. The default value is <b>True</b> .
<b>MemcachedTLS</b>	Set to True to enable TLS on Memcached service. Because not all services support Memcached TLS, during the migration period, Memcached will listen on 2 ports - on the port set with MemcachedPort parameter (above) and on 11211, without TLS. The default value is <b>False</b> .
<b>NotificationDriver</b>	Driver or drivers to handle sending notifications. The default value is <b>noop</b> .
<b>PublicSSLCertificateAutogenerated</b>	Whether the public SSL certificate was autogenerated or not. The default value is <b>False</b> .
<b>PublicTLSCAFile</b>	Specifies the default CA cert to use if TLS is used for services in the public network.
<b>SSLCertificate</b>	The content of the SSL certificate (without Key) in PEM format.
<b>TokenExpiration</b>	Set a token expiration time in seconds. The default value is <b>3600</b> .

## CHAPTER 12. IMAGE STORAGE (GLANCE) PARAMETERS

Parameter	Description
<b>CephClusterName</b>	The Ceph cluster name. The default value is <b>ceph</b> .
<b>GlanceApiOptVolumes</b>	List of optional volumes to be mounted.
<b>GlanceBackend</b>	The short name of the OpenStack Image Storage (glance) backend to use. Should be one of swift, rbd, cinder, or file. The default value is <b>swift</b> .
<b>GlanceBackendID</b>	The default backend's identifier. The default value is <b>default_backend</b> .
<b>GlanceCacheEnabled</b>	Enable OpenStack Image Storage (glance) Image Cache. The default value is <b>False</b> .
<b>GlanceCinderMountPointBase</b>	The mount point base when glance is using cinder as store and cinder backend is NFS. This mount point is where the NFS volume is mounted on the glance node. The default value is <b>/var/lib/glance/mnt</b> .
<b>GlanceDiskFormats</b>	List of allowed disk formats in Glance; all formats are allowed when left unset.
<b>GlanceEnabledImportMethods</b>	List of enabled Image Import Methods. Valid values in the list are <i>glance-direct</i> and <i>web-download</i> . The default value is <b>web-download</b> .
<b>GlanceIgnoreUserRoles</b>	List of user roles to be ignored for injecting image metadata properties. The default value is <b>admin</b> .
<b>GlanceImageCacheDir</b>	Base directory that the Image Cache uses. The default value is <b>/var/lib/glance/image-cache</b> .
<b>GlanceImageCacheMaxSize</b>	The upper limit on cache size, in bytes, after which the cache-pruner cleans up the image cache. The default value is <b>10737418240</b> .
<b>GlanceImageCacheStallTime</b>	The amount of time, in seconds, to let an image remain in the cache without being accessed. The default value is <b>86400</b> .
<b>GlanceImageConversionOutputFormat</b>	Desired output format for image conversion plugin. The default value is <b>raw</b> .

Parameter	Description
<b>GlanceImageImportPlugins</b>	List of enabled Image Import Plugins. Valid values in the list are <i>image_conversion</i> , <i>inject_metadata</i> , <i>no_op</i> . The default value is <b>['no_op']</b> .
<b>GlanceImageMemberQuota</b>	Maximum number of image members per image. Negative values evaluate to unlimited. The default value is <b>128</b> .
<b>GlanceImagePrefetcherInterval</b>	The interval in seconds to run periodic job <i>cache_images</i> . The default value is <b>300</b> .
<b>GlanceInjectMetadataProperties</b>	Metadata properties to be injected in image.
<b>GlanceLogFile</b>	The filepath of the file to use for logging messages from OpenStack Image Storage (glance).
<b>GlanceMultistoreConfig</b>	Dictionary of settings when configuring additional glance backends. The hash key is the backend ID, and the value is a dictionary of parameter values unique to that backend. Multiple rbd backends are allowed, but cinder, file and swift backends are limited to one each. Example: # Default glance store is rbd. GlanceBackend: rbd GlanceStoreDescription: <i>Default rbd store</i> # GlanceMultistoreConfig specifies a second rbd backend, plus a cinder # backend. GlanceMultistoreConfig: rbd2_store: GlanceBackend: rbd GlanceStoreDescription: <i>Second rbd store</i> CephClusterName: ceph2 # Override CephClientUserName if this cluster uses a different # client name. CephClientUserName: client2 cinder_store: GlanceBackend: cinder GlanceStoreDescription: <i>OpenStack Block Storage (cinder) store</i> .
<b>GlanceNetappNfsEnabled</b>	When using <b>GlanceBackend: file</b> , Netapp mounts NFS share for image storage. The default value is <b>False</b> .
<b>GlanceNfsEnabled</b>	When using <b>GlanceBackend: file</b> , mount NFS share for image storage. The default value is <b>False</b> .
<b>GlanceNfsOptions</b>	NFS mount options for image storage when <b>GlanceNfsEnabled</b> is true. The default value is <b>_netdev,bg,intr,context=system_u:object_r:s_virt_sandbox_file_t:s0</b> .
<b>GlanceNfsShare</b>	NFS share to mount for image storage when <b>GlanceNfsEnabled</b> is true.

Parameter	Description
<b>GlanceNodeStagingUri</b>	URI that specifies the staging location to use when importing images. The default value is <b>file:///var/lib/glance/staging</b> .
<b>GlanceNotifierStrategy</b>	Strategy to use for OpenStack Image Storage (glance) notification queue. The default value is <b>noop</b> .
<b>GlancePassword</b>	The password for the image storage service and database account.
<b>GlanceShowMultipleLocations</b>	Whether to show multiple image locations e.g for copy-on-write support on RBD or Netapp backends. Potential security risk, see glance.conf for more information. The default value is <b>False</b> .
<b>GlanceSparseUploadEnabled</b>	When using GlanceBackend <i>file</i> and <i>rbd</i> to enable or not sparse upload. The default value is <b>False</b> .
<b>GlanceStagingNfsOptions</b>	NFS mount options for NFS image import staging. The default value is <b>_netdev,bg,intr,context=system_u:object_r:s_virt_sandbox_file_t:s0</b> .
<b>GlanceStagingNfsShare</b>	NFS share to mount for image import staging.
<b>GlanceStoreDescription</b>	User facing description for the OpenStack Image Storage (glance) backend. The default value is <b>Default glance store backend.</b>
<b>GlanceWorkers</b>	Set the number of workers for the image storage service. Note that more workers creates a larger number of processes on systems, which results in excess memory consumption. It is recommended to choose a suitable non-default value on systems with high CPU core counts. <b>0</b> sets to the OpenStack internal default, which is equal to the number of CPU cores on the node.
<b>MemcacheUseAdvancedPool</b>	Use the advanced (eventlet safe) memcached client pool. The default value is <b>True</b> .
<b>MultipathdEnable</b>	Whether to enable the multipath daemon. The default value is <b>False</b> .
<b>NetappShareLocation</b>	Netapp share to mount for image storage (when GlanceNetappNfsEnabled is true).

Parameter	Description
<b>NotificationDriver</b>	Driver or drivers to handle sending notifications. The default value is <b>noop</b> .

## CHAPTER 13. KEY MANAGER (BARBICAN) PARAMETERS

Parameter	Description
<b>ApacheCertificateKeySize</b>	Override the private key size used when creating the certificate for this service.
<b>ATOSVars</b>	Hash of atos-hsm role variables used to install ATOS client software.
<b>BarbicanDogtagStoreGlobalDefault</b>	Whether this plugin is the global default plugin. The default value is <b>False</b> .
<b>BarbicanDogtagStoreHost</b>	Hostname of the Dogtag server.
<b>BarbicanDogtagStoreNSSPassword</b>	Password for the NSS DB.
<b>BarbicanDogtagStorePEMPath</b>	Path for the PEM file used to authenticate requests. The default value is <b>/etc/barbican/kra_admin_cert.pem</b> .
<b>BarbicanDogtagStorePort</b>	Port for the Dogtag server. The default value is <b>8443</b> .
<b>BarbicanKmipStoreGlobalDefault</b>	Whether this plugin is the global default plugin. The default value is <b>False</b> .
<b>BarbicanKmipStoreHost</b>	Host for KMIP device.
<b>BarbicanKmipStorePassword</b>	Password to connect to KMIP device.
<b>BarbicanKmipStorePort</b>	Port for KMIP device.
<b>BarbicanKmipStoreUsername</b>	Username to connect to KMIP device.
<b>BarbicanPassword</b>	The password for the OpenStack Key Manager (barbican) service account.
<b>BarbicanPkcs11AlwaysSetCkaSensitive</b>	Always set CKA_SENSITIVE=CK_TRUE. The default value is <b>True</b> .
<b>BarbicanPkcs11CryptoAESGCMGenerateIV</b>	Generate IVs for CKM_AES_GCM encryption mechanism. The default value is <b>True</b> .
<b>BarbicanPkcs11CryptoATOSEnabled</b>	Enable ATOS for PKCS11. The default value is <b>False</b> .
<b>BarbicanPkcs11CryptoEnabled</b>	Enable PKCS11. The default value is <b>False</b> .

Parameter	Description
<b>BarbicanPkcs11CryptoEncryptionMechanism</b>	Cryptoki Mechanism used for encryption. The default value is <b>CKM_AES_CBC</b> .
<b>BarbicanPkcs11CryptoGlobalDefault</b>	Whether this plugin is the global default plugin. The default value is <b>False</b> .
<b>BarbicanPkcs11CryptoHMACKeygenMechanism</b>	Cryptoki Mechanism used to generate Master HMAC Key. The default value is <b>CKM_AES_KEY_GEN</b> .
<b>BarbicanPkcs11CryptoHMACKeyType</b>	Cryptoki Key Type for Master HMAC key. The default value is <b>CKK_AES</b> .
<b>BarbicanPkcs11CryptoHMACLabel</b>	Label for the HMAC key.
<b>BarbicanPkcs11CryptoLibraryPath</b>	Path to vendor PKCS11 library.
<b>BarbicanPkcs11CryptoLogin</b>	Password (PIN) to login to PKCS#11 session.
<b>BarbicanPkcs11CryptoLunasaEnabled</b>	Enable Luna SA HSM for PKCS11. The default value is <b>False</b> .
<b>BarbicanPkcs11CryptoMKEKLabel</b>	Label for Master KEK.
<b>BarbicanPkcs11CryptoMKEKLength</b>	Length of Master KEK in bytes. The default value is <b>256</b> .
<b>BarbicanPkcs11CryptoOsLockingOk</b>	Set CKF_OS_LOCKING_OK flag when initializing the client library. The default value is <b>False</b> .
<b>BarbicanPkcs11CryptoRewrapKeys</b>	Cryptoki Mechanism used to generate Master HMAC Key. The default value is <b>False</b> .
<b>BarbicanPkcs11CryptoSlotId</b>	Slot Id for the PKCS#11 token to be used. The default value is <b>0</b> .
<b>BarbicanPkcs11CryptoThalesEnabled</b>	Enable Thales for PKCS11. The default value is <b>False</b> .
<b>BarbicanPkcs11CryptoTokenLabel</b>	(DEPRECATED) Use <code>BarbicanPkcs11CryptoTokenLabels</code> instead.
<b>BarbicanPkcs11CryptoTokenLabels</b>	List of comma separated labels for the tokens to be used. This is typically a single label, but some devices may require more than one label for Load Balancing and High Availability configurations.
<b>BarbicanPkcs11CryptoTokenSerialNumber</b>	Serial number for PKCS#11 token to be used.

Parameter	Description
<b>BarbicanSimpleCryptoGlobalDefault</b>	Whether this plugin is the global default plugin. The default value is <b>False</b> .
<b>BarbicanSimpleCryptoKek</b>	KEK used to encrypt secrets.
<b>BarbicanWorkers</b>	Set the number of workers for barbican::wsgi::apache. The default value is % <b>{::processorcount}</b> .
<b>CertificateKeySize</b>	Specifies the private key size used when creating the certificate. The default value is <b>2048</b> .
<b>LunasaClientIPNetwork</b>	(Optional) When set OpenStack Key Manager (barbican) nodes will be registered with the HSMs using the IP from this network instead of the FQDN.
<b>LunasaVars</b>	Hash of lunasa-hsm role variables used to install Lunasa client software.
<b>MemcacheUseAdvancedPool</b>	Use the advanced (eventlet safe) memcached client pool. The default value is <b>True</b> .
<b>NotificationDriver</b>	Driver or drivers to handle sending notifications. The default value is <b>noop</b> .
<b>ThalesHSMNetworkName</b>	The network that the HSM is listening on. The default value is <b>internal_api</b> .
<b>ThalesVars</b>	Hash of thales-hsm role variables used to install Thales client software.



## CHAPTER 14. LOAD BALANCER (OCTAVIA) PARAMETERS

Parameter	Description
<b>OctaviaAdminLogFacility</b>	The syslog "LOG_LOCAL" facility to use for the administrative log messages. The default value is <b>1</b> .
<b>OctaviaAdminLogTargets</b>	List of syslog endpoints, host:port comma separated list, to receive administrative log messages.
<b>OctaviaAmphoraExpiryAge</b>	The interval in seconds after which an unused Amphora will be considered expired and cleaned up. If left to 0, the configuration will not be set and the system will use the service defaults. The default value is <b>0</b> .
<b>OctaviaAmphoraSshKeyFile</b>	Public key file path. User will be able to SSH into amphorae with the provided key. User may, in most cases, also elevate to root from user <i>centos</i> (CentOS), <i>ubuntu</i> (Ubuntu) or <i>cloud-user</i> (RHEL) (depends on how amphora image was created). Logging in to amphorae provides a convenient way to e.g. debug load balancing services.
<b>OctaviaAmphoraSshKeyName</b>	SSH key name. The default value is <b>octavia-ssh-key</b> .
<b>OctaviaAntiAffinity</b>	Flag to indicate if anti-affinity feature is turned on. The default value is <b>True</b> .
<b>OctaviaCaCert</b>	OpenStack Load Balancing-as-a-Service (octavia) CA certificate data. If provided, this will create or update a file on the host with the path provided in <code>OctaviaCaCertFile</code> with the certificate data.
<b>OctaviaCaKey</b>	The private key for the certificate provided in <code>OctaviaCaCert</code> . If provided, this will create or update a file on the host with the path provided in <code>OctaviaCaKeyFile</code> with the key data.
<b>OctaviaCaKeyPassphrase</b>	CA private key passphrase.
<b>OctaviaClientCert</b>	OpenStack Load Balancing-as-a-Service (octavia) client certificate data. If provided, this will create or update a file on the host with the path provided in <code>OctaviaClientCertFile</code> with the certificate data.
<b>OctaviaConnectionLogging</b>	When false, tenant connection flows will not be logged. The default value is <b>True</b> .

Parameter	Description
<b>OctaviaDisableLocalLogStorage</b>	When true, logs will not be stored on the amphora filesystem. This includes all kernel, system, and security logs. The default value is <b>False</b> .
<b>OctaviaEnableDriverAgent</b>	Set to false if the driver agent needs to be disabled for some reason. The default value is <b>True</b> .
<b>OctaviaFlavorId</b>	OpenStack Compute (nova) flavor ID to be used when creating the nova flavor for amphora. The default value is <b>65</b> .
<b>OctaviaForwardAllLogs</b>	When true, all log messages from the amphora will be forwarded to the administrative log endpoints, including non-load balancing related logs. The default value is <b>False</b> .
<b>OctaviaGenerateCerts</b>	Enable internal generation of certificates for secure communication with amphorae for isolated private clouds or systems where security is not a concern. Otherwise, use OctaviaCaCert, OctaviaCaKey, OctaviaCaKeyPassphrase, OctaviaClientCert and OctaviaServerCertsKeyPassphrase to configure OpenStack Load Balancing-as-a-Service (octavia). The default value is <b>False</b> .
<b>OctaviaLoadBalancerTopology</b>	Load balancer topology configuration.
<b>OctaviaLogOffload</b>	When true, log messages from the amphora will be forwarded to the administrative log endpoints and will be stored with the controller logs. The default value is <b>False</b> .
<b>OctaviaTenantLogFacility</b>	The syslog "LOG_LOCAL" facility to use for the tenant traffic flow log messages. The default value is <b>0</b> .
<b>OctaviaTenantLogTargets</b>	List of syslog endpoints, host:port comma separated list, to receive tenant traffic flow log messages.
<b>OctaviaTimeoutClientData</b>	Frontend client inactivity timeout. The default value is <b>50000</b> .
<b>OctaviaTimeoutMemberData</b>	Backend member inactivity timeout. The default value is <b>50000</b> .

## CHAPTER 15. MESSAGING PARAMETERS

Parameter	Description
<b>RpcCertificateKeySize</b>	Override the private key size used when creating the certificate for this service.
<b>RpcPassword</b>	The password for messaging backend.
<b>RpcPort</b>	The network port for messaging backend. The default value is <b>5672</b> .
<b>RpcUserName</b>	The username for messaging backend. The default value is <b>guest</b> .
<b>RpcUseSSL</b>	Messaging client subscriber parameter to specify an SSL connection to the messaging host. The default value is <b>False</b> .

## CHAPTER 16. NETWORKING (NEUTRON) PARAMETERS

Parameter	Description
<b>CertificateKeySize</b>	Specifies the private key size used when creating the certificate. The default value is <b>2048</b> .
<b>ContainerOvnCertificateKeySize</b>	Override the private key size used when creating the certificate for this service.
<b>DhcpAgentNotification</b>	Enables DHCP agent notifications. The default value is <b>True</b> .
<b>DockerAdditionalSockets</b>	Additional domain sockets for the docker daemon to bind to (useful for mounting into containers that launch other containers). The default value is <b>['/var/lib/openstack/docker.sock']</b> .
<b>DockerInsecureRegistryAddress</b>	Optional. The IP Address and Port of an insecure docker namespace that will be configured in <code>/etc/sysconfig/docker</code> . The value can be multiple addresses separated by commas.
<b>EnableVLANTransparency</b>	If True, then allow plugins that support it to create VLAN transparent networks. The default value is <b>False</b> .
<b>MemcacheUseAdvancedPool</b>	Use the advanced (eventlet safe) memcached client pool. The default value is <b>True</b> .
<b>NeutronAgentDownTime</b>	Seconds to regard the agent as down; should be at least twice <code>NeutronGlobalReportInterval</code> , to be sure the agent is down for good. The default value is <b>600</b> .
<b>NeutronAllowL3AgentFailover</b>	Allow automatic l3-agent failover. The default value is <b>True</b> .
<b>NeutronApiOptEnvVars</b>	Hash of optional environment variables.
<b>NeutronApiOptVolumes</b>	List of optional volumes to be mounted.
<b>NeutronBridgeMappings</b>	The logical to physical bridge mappings to use. The default ( <b>datacentre:br-ex</b> ) maps <b>br-ex</b> (the external bridge on hosts) to a physical name <b>datacentre</b> , which provider networks can use (for example, the default floating network). If changing this, either use different post-install network scripts or be sure to keep <b>datacentre</b> as a mapping network name. The default value is <b>datacentre:br-ex</b> .

Parameter	Description
<b>NeutronCertificateKeySize</b>	Override the private key size used when creating the certificate for this service.
<b>NeutronCorePlugin</b>	The core plugin for networking. The value should be the entrypoint to be loaded from <b>neutron.core_plugins</b> namespace. The default value is <b>ml2</b> .
<b>NeutronDBSyncExtraParams</b>	String of extra command line parameters to append to the <b>neutron-db-manage upgrade head</b> command.
<b>NeutronDefaultAvailabilityZones</b>	Comma-separated list of default network availability zones to be used by OpenStack Networking (neutron) if its resource is created without availability zone hints. If not set, no AZs will be configured for OpenStack Networking (neutron) network services.
<b>NeutronDhcpAgentsPerNetwork</b>	The number of DHCP agents to schedule per network. The default value is <b>0</b> .
<b>NeutronDhcpLoadType</b>	Additional to the availability zones aware network scheduler. The default value is <b>networks</b> .
<b>NeutronDnsDomain</b>	Domain to use for building the hostnames. The default value is <b>openstacklocal</b> .
<b>NeutronEnableDVR</b>	Enable Distributed Virtual Router.
<b>NeutronEnableIcmpSnooping</b>	Enable IGMP Snooping. The default value is <b>False</b> .
<b>NeutronFirewallDriver</b>	Firewall driver for realizing OpenStack Networking (neutron) security group function. The default value is <b>iptables_hybrid</b> .
<b>NeutronFlatNetworks</b>	Sets the flat network name to configure in plugins. The default value is <b>datacentre</b> .
<b>NeutronGeneveMaxHeaderSize</b>	Geneve encapsulation header size. The default value is <b>38</b> .

Parameter	Description
<b>NeutronGlobalPhysnetMtu</b>	MTU of the underlying physical network. OpenStack Networking (neutron) uses this value to calculate MTU for all virtual network components. For flat and VLAN networks, OpenStack Networking uses this value without modification. For overlay networks such as VXLAN, OpenStack Networking automatically subtracts the overlay protocol overhead from this value. The default value is <b>0</b> .
<b>NeutronGlobalReportInterval</b>	Seconds between nodes reporting state to server; should be less than NeutronAgentDownTime, best if it is half or less than NeutronAgentDownTime. The default value is <b>300</b> .
<b>NeutronMechanismDrivers</b>	The mechanism drivers for the OpenStack Networking (neutron) tenant network. The default value is <b>ovn</b> .
<b>NeutronMetadataProxySharedSecret</b>	Shared secret to prevent spoofing.
<b>NeutronML2PhysicalNetworkMtus</b>	A list of mappings of physical networks to MTU values. The format of the mapping is <b>&lt;physnet&gt;: &lt;mtu val&gt;</b> . This mapping allows you to specify a physical network MTU value that differs from the default <b>segment_mtu</b> value in ML2 plugin and overwrites values from <b>global_physnet_mtu</b> for the selected network.
<b>NeutronNetworkSchedulerDriver</b>	The network schedule driver to use for availability zones. The default value is <b>neutron.scheduler.dhcp_agent_scheduler.AZ AwareWeightScheduler</b> .
<b>NeutronNetworkType</b>	The tenant network type for OpenStack Networking (neutron). The default value is <b>geneve</b> .
<b>NeutronNetworkVLANRanges</b>	The OpenStack Networking (neutron) ML2 and Open vSwitch VLAN mapping range to support. Defaults to permitting any VLAN on the <b>datacentre</b> physical network (See <b>NeutronBridgeMappings</b> ). The default value is <b>datacentre:1:1000</b> .
<b>NeutronOverlayIPVersion</b>	IP version used for all overlay network endpoints. The default value is <b>4</b> .
<b>NeutronOVNLoggingBurstLimit</b>	Maximum number of packets per rate_limit. The default value is <b>25</b> .
<b>NeutronOVNLoggingLocalOutputLogBase</b>	Output logfile path on agent side, default syslog file.

Parameter	Description
<b>NeutronOVNLoggingRateLimit</b>	Maximum number of packets logging per second. The default value is <b>100</b> .
<b>NeutronOvsIntegrationBridge</b>	Name of Open vSwitch bridge to use.
<b>NeutronOvsVnicTypeBlacklist</b>	Comma-separated list of VNIC types for which support in OpenStack Networking (neutron) is administratively prohibited by the OVS mechanism driver.
<b>NeutronPassword</b>	The password for the OpenStack Networking (neutron) service and database account.
<b>NeutronPluginExtensions</b>	Comma-separated list of enabled extension plugins. The default value is <b>qos,port_security,dns</b> .
<b>NeutronPluginMI2PuppetTags</b>	Puppet resource tag names that are used to generate configuration files with puppet. The default value is <b>neutron_plugin_ml2</b> .
<b>NeutronPortQuota</b>	Number of ports allowed per tenant, and minus means unlimited. The default value is <b>500</b> .
<b>NeutronRouterSchedulerDriver</b>	The router schedule driver to use for availability zones. The default value is <b>neutron.scheduler.l3_agent_scheduler.AZLeastRoutersScheduler</b> .
<b>NeutronRpcWorkers</b>	Sets the number of RPC workers for the OpenStack Networking (neutron) service. If not specified, it'll take the value of NeutronWorkers and if this is not specified either, the default value results in the configuration being left unset and a system-dependent default will be chosen (usually 1).
<b>NeutronServicePlugins</b>	Comma-separated list of service plugin endpoints. The default value is <b>qos,ovn-router,trunk,segments,port_forwarding,log</b> .
<b>NeutronSriovVnicTypeBlacklist</b>	Comma-separated list of VNIC types for which support in OpenStack Networking (neutron) is administratively prohibited by the SR-IOV mechanism driver.
<b>NeutronTunnelIdRanges</b>	Comma-separated list of <tun_min>:<tun_max> tuples enumerating ranges of GRE tunnel IDs that are available for tenant network allocation. The default value is <b>['1:4094']</b> .

Parameter	Description
<b>NeutronTypeDrivers</b>	Comma-separated list of network type driver entrypoints to be loaded. The default value is <b>geneve,vxlan,vlan,flat</b> .
<b>NeutronVhostuserSocketDir</b>	The vhost-user socket directory for OVS.
<b>NeutronVniRanges</b>	Comma-separated list of <vni_min>:<vni_max> tuples enumerating ranges of VXLAN VNI IDs that are available for tenant network allocation. The default value is <b>['1:65536']</b> .
<b>NeutronWorkers</b>	Sets the number of API and RPC workers for the OpenStack Networking service. Note that more workers creates a larger number of processes on systems, which results in excess memory consumption. It is recommended to choose a suitable non-default value on systems with high CPU core counts. <b>0</b> sets to the OpenStack internal default, which is equal to the number of CPU cores on the node.
<b>NotificationDriver</b>	Driver or drivers to handle sending notifications. The default value is <b>noop</b> .
<b>OVNAvailabilityZone</b>	The az options to configure in ovs db. eg. [az-0, az-1, az-2].
<b>OVNCMSOptions</b>	The CMS options to configure in ovs db.
<b>OVNControllerUpdateTimeout</b>	During update, how long we wait for the container to be updated, in seconds. The default value is <b>600</b> .
<b>OVNDbConnectionTimeout</b>	Timeout in seconds for the OVSDB connection transaction. The default value is <b>180</b> .
<b>OVNDnsServers</b>	List of servers to use as as dns forwarders.
<b>OVNEnableHaproxyDockerWrapper</b>	Generate a wrapper script so that haproxy is launched in a separate container. The default value is <b>True</b> .
<b>OVNEncapType</b>	Type of encapsulation used in OVN. It can be "geneve" or "vxlan". The default value is <b>geneve</b> .
<b>OVNIntegrationBridge</b>	Name of the OVS bridge to use as integration bridge by OVN Controller. The default value is <b>br-int</b> .



Parameter	Description
<b>OvnMetadataCertificateKeySize</b>	Override the private key size used when creating the certificate for this service.
<b>OVNMetadataEnabled</b>	Whether Metadata Service has to be enabled. The default value is <b>True</b> .
<b>OVNNeutronSyncMode</b>	The synchronization mode of OVN with OpenStack Networking (neutron) DB. The default value is <b>log</b> .
<b>OVNNorthboundServerPort</b>	Port of the OVN Northbound DB server. The default value is <b>6641</b> .
<b>OVNOpenflowProbeInterval</b>	The inactivity probe interval of the OpenFlow connection to the OpenvSwitch integration bridge, in seconds. The default value is <b>60</b> .
<b>OVNQosDriver</b>	OVN notification driver for OpenStack Networking (neutron) QOS service plugin. The default value is <b>ovn-qos</b> .
<b>OVNRemoteProbeInterval</b>	Probe interval in ms. The default value is <b>60000</b> .
<b>OVNSouthboundServerPort</b>	Port of the OVN Southbound DB server. The default value is <b>6642</b> .
<b>OVNVifType</b>	Type of VIF to be used for ports. The default value is <b>ovs</b> .
<b>OvsHwOffload</b>	Enable OVS Hardware Offload. This feature supported from OVS 2.8.0. The default value is <b>False</b> .
<b>TenantNetPhysnetMtu</b>	MTU of the underlying physical network. OpenStack Networking (neutron) uses this value to calculate MTU for all virtual network components. For flat and VLAN networks, OpenStack Networking (neutron) uses this value without modification. For overlay networks such as VXLAN, OpenStack Networking (neutron) automatically subtracts the overlay protocol overhead from this value. (The mtu setting of the Tenant network in network_data.yaml control's this parameter.). The default value is <b>1500</b> .

## CHAPTER 17. OBJECT STORAGE (SWIFT) PARAMETERS

Parameter	Description
<b>MemcachedTLS</b>	Set to True to enable TLS on Memcached service. Because not all services support Memcached TLS, during the migration period, Memcached will listen on 2 ports - on the port set with MemcachedPort parameter (above) and on 11211, without TLS. The default value is <b>False</b> .
<b>SwiftAccountWorkers</b>	Number of workers for Swift account service. The default value is <b>0</b> .
<b>SwiftCeilometerIgnoreProjects</b>	Comma-separated list of project names to ignore. The default value is <b>['service']</b> .
<b>SwiftCeilometerPipelineEnabled</b>	Set to <b>False</b> to disable the object storage proxy ceilometer pipeline. The default value is <b>False</b> .
<b>SwiftContainerSharderEnabled</b>	Set to True to enable Swift container sharder service. The default value is <b>False</b> .
<b>SwiftContainerWorkers</b>	Number of workers for Swift account service. The default value is <b>0</b> .
<b>SwiftCorsAllowedOrigin</b>	Indicate whether this resource may be shared with the domain received in the request "origin" header.
<b>SwiftEncryptionEnabled</b>	Set to True to enable data-at-rest encryption in Swift. The default value is <b>False</b> .
<b>SwiftHashSuffix</b>	A random string to be used as a salt when hashing to determine mappings in the ring.
<b>SwiftMinPartHours</b>	The minimum time (in hours) before a partition in a ring can be moved following a rebalance. The default value is <b>1</b> .
<b>SwiftMountCheck</b>	Check if the devices are mounted to prevent accidentally writing to the root device. The default value is <b>False</b> .
<b>SwiftObjectWorkers</b>	Number of workers for Swift account service. The default value is <b>0</b> .
<b>SwiftPartPower</b>	Partition power to use when building object storage rings. The default value is <b>10</b> .

Parameter	Description
<b>SwiftPassword</b>	The password for the object storage service account.
<b>SwiftProxyNodeTimeout</b>	Timeout for requests going from <b>swift-proxy</b> to account, container, and object services. The default value is <b>60</b> .
<b>SwiftRawDisks</b>	Additional raw devices to use for the object storage backend. For example: <b>{sdb: {}}</b>
<b>SwiftReplicas</b>	Number of replicas to use in the object storage rings. The default value is <b>3</b> .
<b>SwiftRingBuild</b>	Whether to manage object storage rings or not. The default value is <b>True</b> .
<b>SwiftRingGetTempurl</b>	A temporary Swift URL to download rings from.
<b>SwiftRingPutTempurl</b>	A temporary Swift URL to upload rings to.
<b>SwiftUseLocalDir</b>	Use a local directory for object storage services when building rings. The default value is <b>True</b> .
<b>SwiftWorkers</b>	Number of workers for object storage service. Note that more workers creates a larger number of processes on systems, which results in excess memory consumption. It is recommended to choose a suitable non-default value on systems with high CPU core counts. <b>0</b> sets to the OpenStack internal default, which is equal to the number of CPU cores on the node. The default value is <b>0</b> .

## CHAPTER 18. ORCHESTRATION (HEAT) PARAMETERS

Parameter	Description
<b>ApacheCertificateKeySize</b>	Override the private key size used when creating the certificate for this service.
<b>CertificateKeySize</b>	Specifies the private key size used when creating the certificate. The default value is <b>2048</b> .
<b>EnableCache</b>	Enable caching with memcached. The default value is <b>True</b> .
<b>HeatApiOptEnvVars</b>	Hash of optional environment variables.
<b>HeatApiOptVolumes</b>	List of optional volumes to be mounted.
<b>HeatAuthEncryptionKey</b>	Auth encryption key for heat-engine.
<b>HeatConfigureDelegatedRoles</b>	Create delegated roles. The default value is <b>False</b> .
<b>HeatConvergenceEngine</b>	Enables the heat engine with the convergence architecture. The default value is <b>True</b> .
<b>HeatCorsAllowedOrigin</b>	Indicate whether this resource may be shared with the domain received in the request "origin" header.
<b>HeatCronPurgeDeletedAge</b>	Cron to purge database entries marked as deleted and older than \$age - Age. The default value is <b>30</b> .
<b>HeatCronPurgeDeletedAgeType</b>	Cron to purge database entries marked as deleted and older than \$age - Age type. The default value is <b>days</b> .
<b>HeatCronPurgeDeletedDestination</b>	Cron to purge database entries marked as deleted and older than \$age - Log destination. The default value is <b>/dev/null</b> .
<b>HeatCronPurgeDeletedEnsure</b>	Cron to purge database entries marked as deleted and older than \$age - Ensure. The default value is <b>present</b> .
<b>HeatCronPurgeDeletedHour</b>	Cron to purge database entries marked as deleted and older than \$age - Hour. The default value is <b>0</b> .
<b>HeatCronPurgeDeletedMaxDelay</b>	Cron to purge database entries marked as deleted and older than \$age - Max Delay. The default value is <b>3600</b> .

Parameter	Description
<b>HeatCronPurgeDeletedMinute</b>	Cron to purge database entries marked as deleted and older than \$age - Minute. The default value is <b>1</b> .
<b>HeatCronPurgeDeletedMonth</b>	Cron to purge database entries marked as deleted and older than \$age - Month. The default value is <b>*</b> .
<b>HeatCronPurgeDeletedMonthday</b>	Cron to purge database entries marked as deleted and older than \$age - Month Day. The default value is <b>*</b> .
<b>HeatCronPurgeDeletedUser</b>	Cron to purge database entries marked as deleted and older than \$age - User. The default value is <b>heat</b> .
<b>HeatCronPurgeDeletedWeekday</b>	Cron to purge database entries marked as deleted and older than \$age - Week Day. The default value is <b>*</b> .
<b>HeatEnableDBPurge</b>	Whether to create cron job for purging soft deleted rows in the OpenStack Orchestration (heat) database. The default value is <b>True</b> .
<b>HeatEngineOptEnvVars</b>	Hash of optional environment variables.
<b>HeatEngineOptVolumes</b>	List of optional volumes to be mounted.
<b>HeatEnginePluginDirs</b>	An array of directories to search for plug-ins.
<b>HeatMaxJsonBodySize</b>	Maximum raw byte size of the OpenStack Orchestration (heat) API JSON request body. The default value is <b>4194304</b> .
<b>HeatMaxNestedStackDepth</b>	Maximum number of nested stack depth. The default value is <b>6</b> .
<b>HeatMaxResourcesPerStack</b>	Maximum resources allowed per top-level stack. -1 stands for unlimited. The default value is <b>1000</b> .
<b>HeatPassword</b>	The password for the Orchestration service and database account.
<b>HeatReauthenticationAuthMethod</b>	Allow reauthentication on token expiry, such that long-running tasks may complete. Note this defeats the expiry of any provided user tokens.
<b>HeatStackDomainAdminPassword</b>	The admin password for the OpenStack Orchestration (heat) domain in OpenStack Identity (keystone).

Parameter	Description
<b>HeatWorkers</b>	Number of workers for OpenStack Orchestration (heat) service. Note that more workers creates a larger number of processes on systems, which results in excess memory consumption. It is recommended to choose a suitable non-default value on systems with high CPU core counts. <b>0</b> sets to the OpenStack internal default, which is equal to the number of CPU cores on the node. The default value is <b>0</b> .
<b>HeatYaqlLimitIterators</b>	The maximum number of elements in collection yaql expressions can take for its evaluation. The default value is <b>1000</b> .
<b>HeatYaqlMemoryQuota</b>	The maximum size of memory in bytes that yaql expressions can take for its evaluation. The default value is <b>100000</b> .
<b>MemcachedTLS</b>	Set to True to enable TLS on Memcached service. Because not all services support Memcached TLS, during the migration period, Memcached will listen on 2 ports - on the port set with MemcachedPort parameter (above) and on 11211, without TLS. The default value is <b>False</b> .
<b>MemcacheUseAdvancedPool</b>	Use the advanced (eventlet safe) memcached client pool. The default value is <b>True</b> .
<b>NotificationDriver</b>	Driver or drivers to handle sending notifications. The default value is <b>noop</b> .

## CHAPTER 19. SHARED FILE SERVICE (MANILA) PARAMETERS

Parameter	Description
<b>ApacheCertificateKeySize</b>	Override the private key size used when creating the certificate for this service.
<b>CephClusterName</b>	The Ceph cluster name. The default value is <b>ceph</b> .
<b>CertificateKeySize</b>	Specifies the private key size used when creating the certificate. The default value is <b>2048</b> .
<b>ManilaCephClientUserName</b>	Ceph client username for manila integration. The default value is <b>manila</b> .
<b>ManilaCephFSCephFSProtocolHelperType</b>	Protocol type ( <i>CEPHFS</i> or <i>NFS</i> ) when cephfs back end is enabled. Set via manila cephfs environment files. The default value is <b>CEPHFS</b> .
<b>ManilaEnabledShareProtocols</b>	List of protocols to be allowed for share creation in manila. When not set, the list is inferred via the storage back end/s enabled.
<b>ManilaIPv6</b>	Set to True to enable IPv6 access in manila. The default value is <b>False</b> .
<b>ManilaPassword</b>	The password for the shared file service account.
<b>ManilaWorkers</b>	Set the number of workers for manila::wsgi::apache. The default value is equal to the number of vCPU cores on the physical node.
<b>MemcacheUseAdvancedPool</b>	Use the advanced (eventlet safe) memcached client pool. The default value is <b>True</b> .
<b>NotificationDriver</b>	Driver or drivers to handle sending notifications. The default value is <b>noop</b> .

## CHAPTER 20. TIME PARAMETERS

Parameter	Description
<b>ChronyAclRules</b>	Access Control List of NTP clients. By default no clients are permitted. The default value is <b>['deny all']</b> .
<b>ChronyGlobalPoolOptions</b>	Default pool options for the configured NTP pools in chrony.conf. If this is specified, NtpIburstEnable, MaxPoll, and MinPoll are ignored.
<b>ChronyGlobalServerOptions</b>	Default server options for the configured NTP servers in chrony.conf. If this is specified, NtpIburstEnable, MaxPoll, and MinPoll are ignored.
<b>EnablePackageInstall</b>	Set to true to enable package installation at deploy time. The default value is <b>false</b> .
<b>MaxPoll</b>	Specify maximum poll interval of upstream servers for NTP messages, in seconds to the power of two. Allowed values are 4 to 17. The default value is <b>10</b> .
<b>MinPoll</b>	Specify minimum poll interval of upstream servers for NTP messages, in seconds to the power of two. The minimum poll interval defaults to 6 (64 s). Allowed values are 4 to 17. The default value is <b>6</b> .
<b>NtpIburstEnable</b>	Specifies whether to enable the iburst option for every NTP peer. If iburst is enabled, when the NTP server is unreachable NTP will send a burst of eight packages instead of one. This is designed to speed up the initial synchronization. The default value is <b>True</b> .
<b>NtpPool</b>	NTP pool list. Defaults to [], so only NtpServer is used by default.
<b>NtpServer</b>	NTP servers list. The default value is <b>['0.pool.ntp.org', '1.pool.ntp.org', '2.pool.ntp.org', '3.pool.ntp.org']</b> .
<b>TimeZone</b>	The timezone to be set on the overcloud. The default value is <b>UTC</b> .



## CHAPTER 21. UPGRADE PARAMETERS

Parameter	Description
<b>UpgradeInitCommand</b>	Command or script snippet to run on all overcloud nodes to initialize the upgrade process. For example, a repository switch.
<b>UpgradeInitCommonCommand</b>	Common commands required by the upgrades process. This should not normally be modified by the operator and is set and unset in the major-upgrade-composable-steps.yaml and major-upgrade-converge.yaml environment files.
<b>UpgradeLeappCommandOptions</b>	Additional command line options to append to the Leapp command.
<b>UpgradeLeappDebug</b>	Print debugging output when running Leapp. The default value is <b>True</b> .
<b>UpgradeLeappDevelSkip</b>	Skip Leapp checks by setting env variables when running Leapp in development/testing. For example, LEAPP_DEVEL_SKIP_RHSM=1.
<b>UpgradeLeappEnabled</b>	Use Leapp for operating system upgrade. The default value is <b>False</b> .
<b>UpgradeLeappPostRebootDelay</b>	Maximum (seconds) to wait for machine to reboot and respond to a test command. The default value is <b>120</b> .
<b>UpgradeLeappRebootTimeout</b>	Timeout (seconds) for the OS upgrade phase via Leapp. The default value is <b>3600</b> .
<b>UpgradeLeappToInstall</b>	List of packages to install after Leapp upgrade.
<b>UpgradeLeappToRemove</b>	List of packages to remove during Leapp upgrade.