



Red Hat OpenStack Platform 17.0

Storage Guide

Understanding, using, and managing persistent storage in OpenStack

Red Hat OpenStack Platform 17.0 Storage Guide

Understanding, using, and managing persistent storage in OpenStack

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

This guide details the different procedures for using and managing persistent storage in a Red Hat OpenStack Platform environment. It also includes procedures for configuring and managing the respective OpenStack service of each persistent storage type.

Table of Contents

MAKING OPEN SOURCE MORE INCLUSIVE	5
PROVIDING FEEDBACK ON RED HAT DOCUMENTATION	6
CHAPTER 1. INTRODUCTION TO PERSISTENT STORAGE IN RED HAT OPENSTACK PLATFORM (RHOSP)	7
1.1. SCALABILITY AND BACK-END STORAGE	8
1.2. STORAGE ACCESSIBILITY AND ADMINISTRATION	8
1.3. STORAGE SECURITY	9
1.4. STORAGE REDUNDANCY AND DISASTER RECOVERY	9
CHAPTER 2. CONFIGURING THE BLOCK STORAGE SERVICE (CINDER)	11
2.1. BLOCK STORAGE SERVICE BACK ENDS	11
2.2. ACTIVE-ACTIVE BLOCK STORAGE FOR HIGH AVAILABILITY	12
2.2.1. Enabling active-active Block Storage	12
2.2.2. Maintenance commands for active-active Block Storage configurations	13
2.2.3. Volume manage and unmanage	13
2.2.4. Volume migration on a clustered service	14
2.2.5. Initiating Block Storage service maintenance	14
2.3. GROUP VOLUME CONFIGURATION WITH VOLUME TYPES	15
2.3.1. Listing back-end driver properties	16
2.3.2. Creating and configuring a volume type	17
2.3.3. Editing a volume type	18
2.3.4. Creating and configuring private volume types	18
2.3.5. Defining a project-specific default volume type	20
2.4. CREATING AND CONFIGURING AN INTERNAL PROJECT FOR THE BLOCK STORAGE SERVICE (CINDER)	21
2.5. CONFIGURING THE IMAGE-VOLUME CACHE	22
2.6. BLOCK STORAGE SERVICE (CINDER) QUALITY OF SERVICE SPECIFICATIONS	23
2.6.1. Consumers of QoS specifications	23
2.6.2. Block Storage QoS property keys	24
2.6.2.1. Set the IO request size for IOPS limits	26
2.6.2.2. IOPS limits that scale according to volume size	27
2.6.3. Creating and configuring a QoS specification with the Dashboard	28
2.6.4. Creating and configuring a QoS specification with the CLI	29
2.6.5. Associating a QoS specification with a volume type by using the Dashboard	30
2.6.6. Associating a QoS specification with a volume type by using the CLI	31
2.6.7. Disassociating a QoS specification from a volume type with the Dashboard	32
2.6.8. Disassociating a QoS specification from volume types with the CLI	32
2.7. BLOCK STORAGE SERVICE (CINDER) VOLUME ENCRYPTION	33
2.7.1. Configuring Block Storage service volume encryption with the Dashboard	33
2.7.2. Configuring Block Storage service volume encryption with the CLI	34
2.7.3. Automatic deletion of volume image encryption key	35
2.8. DEPLOYING AVAILABILITY ZONES FOR BLOCK STORAGE VOLUME BACK ENDS	36
2.9. BLOCK STORAGE SERVICE (CINDER) CONSISTENCY GROUPS	37
2.9.1. Configuring Block Storage service consistency groups	37
2.9.2. Creating Block Storage consistency groups with the Dashboard	39
2.9.3. Managing Block Storage service consistency groups with the Dashboard	39
2.9.4. Creating and managing consistency group snapshots for the Block Storage service	40
2.9.5. Cloning Block Storage service consistency groups	41
2.10. CONFIGURING THE DEFAULT BLOCK STORAGE SCHEDULER FILTERS	41
2.11. ENABLING LVM2 FILTERING ON OVERCLOUD NODES	42
2.12. MULTIPATH CONFIGURATION	43

2.12.1. Using director to configure multipath	44
2.12.1.1. Multipath heat template parameters	45
2.12.2. Verifying multipath configuration	46
CHAPTER 3. PERFORMING BASIC OPERATIONS WITH THE BLOCK STORAGE SERVICE (CINDER)	48
3.1. CREATING BLOCK STORAGE VOLUMES	48
3.2. EDITING A VOLUME NAME OR DESCRIPTION	50
3.3. RESIZING (EXTENDING) A BLOCK STORAGE SERVICE VOLUME	50
3.4. DELETING A BLOCK STORAGE SERVICE VOLUME	51
3.5. VOLUME ALLOCATION ON MULTIPLE BACK ENDS	52
3.6. ATTACHING A VOLUME TO AN INSTANCE	52
3.7. DETACHING A VOLUME FROM AN INSTANCE	53
3.8. CONFIGURING THE ACCESS RIGHTS TO A VOLUME	53
3.9. CHANGING A VOLUME OWNER WITH THE DASHBOARD	54
3.10. CHANGING A VOLUME OWNER WITH THE CLI	55
CHAPTER 4. PERFORMING ADVANCED OPERATIONS WITH THE BLOCK STORAGE SERVICE (CINDER)	57
4.1. CREATING VOLUME SNAPSHOTS	57
4.2. CREATING NEW VOLUMES FROM SNAPSHOTS	58
4.3. DELETING VOLUME SNAPSHOTS	59
4.4. RESTORING A VOLUME FROM A SNAPSHOT	59
4.5. UPLOADING A VOLUME TO THE IMAGE SERVICE (GLANCE)	61
4.6. VOLUMES THAT CAN BE ATTACHED TO MULTIPLE INSTANCES	61
4.6.1. Creating a multi-attach volume type	62
4.6.2. Multi-attach volume retyping	63
4.6.3. Creating a multi-attach volume	63
4.7. MOVING VOLUMES BETWEEN BACK ENDS	63
4.7.1. Moving available volumes	64
4.7.2. Moving in-use volumes	65
4.8. BLOCK STORAGE VOLUME RETYPING	65
4.8.1. Retyping a volume from the Dashboard	66
4.8.2. Retyping a volume from the CLI	67
4.9. MIGRATING A VOLUME BETWEEN BACK ENDS WITH THE DASHBOARD	67
4.10. MIGRATING A VOLUME BETWEEN BACK ENDS WITH THE CLI	68
4.11. ENCRYPTING UNENCRYPTED VOLUMES	70
4.12. PROTECTED AND UNPROTECTED SNAPSHOTS IN A RED HAT CEPH STORAGE BACK END	71
CHAPTER 5. CONFIGURING THE OBJECT STORAGE SERVICE (SWIFT)	72
5.1. OBJECT STORAGE RINGS	72
5.1.1. Checking cluster health	72
5.1.2. Increasing ring partition power	73
5.1.3. Partition power recommendation for the Object Storage service	74
5.1.4. Custom rings	75
5.2. CUSTOMIZING THE OBJECT STORAGE SERVICE	75
5.2.1. Configuring fast-post	75
5.2.2. Enabling at-rest encryption	76
5.2.3. Deploying a standalone Object Storage service cluster	76
5.2.4. Using external SAN disks	78
5.2.5. Disk recommendation for the Object Storage service	79
5.3. ADDING OR REMOVING OBJECT STORAGE NODES	79
5.3.1. Adding nodes to the overcloud	79
5.3.2. Scaling up bare-metal nodes	81
5.3.3. Defining dedicated Object Storage nodes	82
5.3.4. Updating and rebalancing the Object Storage rings	83

5.3.5. Syncing node changes and migrating data	84
5.3.6. Removing Object Storage nodes	85
5.3.7. Scaling down bare-metal nodes	85
5.4. CONTAINER MANAGEMENT IN THE OBJECT STORAGE SERVICE	87
5.4.1. Creating private and public containers	88
5.4.2. Creating pseudo folders for containers	88
5.4.3. Deleting containers from the Object Storage service	88
5.4.4. Uploading objects to containers	89
5.4.5. Copying objects between containers	89
5.4.6. Deleting objects from the Object Storage service	90
CHAPTER 6. CONFIGURING THE SHARED FILE SYSTEMS SERVICE (MANILA)	91
6.1. CONFIGURING SHARED FILE SYSTEMS SERVICE BACK ENDS	91
6.1.1. Configuring multiple back ends	91
6.1.2. Deploying multiple back ends	92
6.1.3. Confirming deployment of multiple back ends	93
6.1.4. Overriding allowed NAS protocols	94
6.1.5. Viewing back-end capabilities	94
6.2. CREATING SHARE TYPES	96
6.3. COMPARING COMMON CAPABILITIES OF SHARE TYPES	97
6.4. PLANNING NETWORKING FOR SHARED FILE SYSTEMS	97
6.5. ENSURING NETWORK CONNECTIVITY TO THE SHARE	98
6.6. CHANGING THE DEFAULT QUOTAS IN THE SHARED FILE SYSTEMS SERVICE	99
6.6.1. Updating quotas for projects, users, and share types	99
6.6.2. Resetting quotas for projects, users, and share types	101
6.6.3. Updating the default quota values for Shared File Systems service projects	102
CHAPTER 7. PERFORMING OPERATIONS WITH THE SHARED FILE SYSTEMS SERVICE (MANILA)	104
7.1. DISCOVERING SHARE TYPES	104
7.2. CREATING NFS OR NATIVE CEPHFS SHARES	104
7.3. LISTING SHARES AND EXPORTING INFORMATION	105
7.4. CREATING A SNAPSHOT OF DATA ON A SHARED FILE SYSTEM	105
7.4.1. Creating a share from a snapshot	106
7.4.2. Deleting a snapshot	107
7.5. CONNECTING TO A SHARED NETWORK TO ACCESS SHARES	108
7.6. CONFIGURING AN IPV6 INTERFACE BETWEEN THE NETWORK AND AN INSTANCE	110
7.7. GRANTING SHARE ACCESS FOR END-USER CLIENTS	110
7.7.1. Granting access to an NFS share	111
7.7.2. Granting access to a native CephFS share	112
7.7.3. Revoking access to a share	113
7.8. MOUNTING SHARES ON COMPUTE INSTANCES	113
7.8.1. Listing share export locations	114
7.8.2. Mounting NFS or native CephFS	114
7.9. DELETING SHARES	115
7.10. LISTING RESOURCE LIMITS OF THE SHARED FILE SYSTEMS SERVICE	116
7.11. TROUBLESHOOTING OPERATION FAILURES	116
7.11.1. Fixing create share or create share group failures	116
7.11.2. Debugging share mounting failures	123

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. INTRODUCTION TO PERSISTENT STORAGE IN RED HAT OPENSTACK PLATFORM (RHOSP)

Within Red Hat OpenStack Platform, storage is provided by three main services:

- Block Storage (**openstack-cinder**)
- Object Storage (**openstack-swift**)
- Shared File System Storage (**openstack-manila**)


These services provide different types of persistent storage, each with its own set of advantages in different use cases. This guide discusses the suitability of each for general enterprise storage requirements.

You can manage cloud storage by using either the RHOSP dashboard or the command-line clients. You can perform most procedures by using either method. However, you can complete some of the more advanced procedures only on the command line. This guide provides procedures for the dashboard where possible.



NOTE

For the complete suite of documentation for Red Hat OpenStack Platform, see [Red Hat OpenStack Platform Documentation](#).



IMPORTANT

This guide documents the use of **crudini** to apply some custom service settings. As such, you need to install the **crudini** package first:

```
# dnf install crudini -y
```

RHOSP recognizes two types of storage: *ephemeral* and *persistent*. Ephemeral storage is storage that is associated only to a specific Compute instance. Once that instance is terminated, so is its ephemeral storage. This type of storage is useful for basic runtime requirements, such as storing the instance's operating system.

Persistent storage, is designed to survive (persist) independent of any running instance. This storage is used for any data that needs to be reused, either by different instances or beyond the life of a specific instance. RHOSP uses the following types of persistent storage:

Volumes

The OpenStack Block Storage service (**openstack-cinder**) allows users to access block storage devices through *volumes*. Users can attach volumes to instances in order to augment their ephemeral storage with general-purpose persistent storage. Volumes can be detached and re-attached to instances at will, and can only be accessed through the instance they are attached to.

You can also configure instances to not use ephemeral storage. Instead of using ephemeral storage, you can configure the Block Storage service to write images to a volume. You can then use the volume as a bootable root volume for an instance.

Volumes also provide inherent redundancy and disaster recovery through backups and snapshots. In addition, you can also encrypt volumes for added security.

Containers

The OpenStack Object Storage service (`openstack-swift`) provides a fully-distributed storage solution used to store any kind of static data or binary object, such as media files, large datasets, and disk images. The Object Storage service organizes these objects by using containers.

Although the content of a volume can be accessed only through instances, the objects inside a container can be accessed through the Object Storage REST API. As such, the Object Storage service can be used as a repository by nearly every service within the cloud.

Shares

The Shared File Systems service (`openstack-manila`) provides the means to easily provision remote, shareable file systems, or *shares*. Shares allow projects within the cloud to openly share storage, and can be consumed by multiple instances simultaneously.

Each storage type is designed to address specific storage requirements. Containers are designed for wide access, and as such feature the highest throughput, access, and fault tolerance among all storage types. Container usage is geared more towards services.

On the other hand, volumes are used primarily for instance consumption. They do not enjoy the same level of access and performance as containers, but they do have a larger feature set and have more native security features than containers. Shares are similar to volumes in this regard, except that they can be consumed by multiple instances.

The following sections discuss each storage type's architecture and feature set in detail, within the context of specific storage criteria.

1.1. SCALABILITY AND BACK-END STORAGE

In general, a clustered storage solution provides greater back-end scalability. For example, when you use Red Hat Ceph as a Block Storage (`cinder`) back end, you can scale storage capacity and redundancy by adding more Ceph Object Storage Daemon (OSD) nodes. Block Storage, Object Storage (`swift`) and Shared File Systems Storage (`manila`) services support Red Hat Ceph Storage as a back end.

The Block Storage service can use multiple storage solutions as discrete back ends. At the back-end level, you can scale capacity by adding more back ends and restarting the service. The Block Storage service also features a large list of supported back-end solutions, some of which feature additional scalability features.

By default, the Object Storage service uses the file system on configured storage nodes, and it can use as much space as is available. The Object Storage service supports the XFS and ext4 file systems, and both can be scaled up to consume as much underlying block storage as is available. You can also scale capacity by adding more storage devices to the storage node.

The Shared File Systems service provisions file shares from designated storage pools that are managed by one or more third-party back-end storage systems. You can scale this shared storage by increasing the size or number of storage pools available to the service or by adding more third-party back-end storage systems to the deployment.

1.2. STORAGE ACCESSIBILITY AND ADMINISTRATION

Volumes are consumed only through instances, and can only be attached to and mounted within one instance at a time. Users can create snapshots of volumes, which they can be used for cloning or restoring a volume to a previous state. For more information, see [Section 1.4, "Storage redundancy and disaster recovery"](#). As a project administrator, you can use the Block Storage service to create *volume types*, which aggregate volume settings, such as size and back end. You can associate volume types with

Quality of Service (QoS) specifications to provide different levels of performance for your cloud users. Your users can specify the volume type they require when creating new volumes. For example, volumes that use higher performance QoS specifications could provide your users with more IOPS or your users could assign lighter workloads to volumes that use lower performance QoS specifications to conserve resources.

Like volumes, shares are consumed through instances. However, shares can be directly mounted within an instance, and do not need to be attached through the dashboard or CLI. Shares can also be mounted by multiple instances simultaneously. The Shared File Systems service also supports share snapshots and cloning; you can also create *share types* to aggregate settings (similar to volume types).

Objects in a container are accessible via API, and can be made accessible to instances and services within the cloud. This makes them ideal as object repositories for services; for example, the Image service (**openstack-glance**) can store its images in containers managed by the Object Storage service.

1.3. STORAGE SECURITY

The Block Storage service (cinder) provides basic data security through volume encryption. With this, you can configure a volume type to be encrypted through a static key; the key is then used to encrypt all volumes that are created from the configured volume type. For more information, see [Section 2.7, “Block Storage service \(cinder\) volume encryption”](#).

Object and container security is configured at the service and node level. The Object Storage service (swift) provides no native encryption for containers and objects. Rather, the Object Storage service prioritizes accessibility within the cloud, and as such relies solely on the cloud network security to protect object data.

The Shared File Systems service (manila) can secure shares through access restriction, whether by instance IP, user or group, or TLS certificate. In addition, some Shared File Systems service deployments can feature separate share servers to manage the relationship between share networks and shares; some share servers support, or even require, additional network security. For example, a CIFS share server requires the deployment of an LDAP, Active Directory, or Kerberos authentication service.

For more information about how to secure the Image service (glance), such as image signing and verification and metadata definition (metadef) API restrictions, see [The Image service \(glance\) in *Creating and Managing Images*](#).

1.4. STORAGE REDUNDANCY AND DISASTER RECOVERY

The Block Storage service (cinder) features volume backup and restoration, which provides basic disaster recovery for user storage. Use backups to protect volume contents. The service also supports snapshots. In addition to cloning, you can use snapshots to restore a volume to a previous state.

In a multi-back end environment, you can also migrate volumes between back ends. This is useful if you need to take a back end offline for maintenance. Backups are typically stored in a storage back end separate from their source volumes to help protect the data. This is not possible with snapshots because snapshots are dependent on their source volumes.

The Block Storage service also supports the creation of consistency groups to group volumes together for simultaneous snapshot creation. This provides a greater level of data consistency across multiple volumes. For more information, see [Section 2.9, “Block Storage service \(cinder\) consistency groups”](#).

The Object Storage service (swift) provides no built-in backup features. You must perform all backups at the file system or node level. The Object Storage service features more robust redundancy and fault tolerance, even the most basic deployment of the Object Storage service replicates objects multiple times. You can use failover features like **dm-multipath** to enhance redundancy.

The Shared File Systems service provides no built-in backup features for shares, but it does allow you to create snapshots for cloning and restoration.

CHAPTER 2. CONFIGURING THE BLOCK STORAGE SERVICE (CINDER)

The Block Storage service (cinder) manages the administration, security, scheduling, and overall management of all volumes. Volumes are used as the primary form of persistent storage for Compute instances.

For more information about volume backups, see the [Block Storage Backup Guide](#).



IMPORTANT

You must install host bus adapters (HBAs) on all Controller nodes and Compute nodes in any deployment that uses the Block Storage service and a Fibre Channel (FC) back end.

Block Storage is configured using the Block Storage REST API.



NOTE

Ensure that you are using Block Storage REST API version 3 because Block Storage no longer supports version 2. The default overcloud deployment does this for you by setting the environment variable **OS_VOLUME_API_VERSION=3.0**.

The Block Storage REST API preserves backward compatibility by using microversions to add enhancements. The **cinder** CLI uses the REST API version of 3.0, unless you specify a specific microversion. For instance, to specify the 3.17 microversion for a **cinder** command, add the **--os-volume-api-version 3.17** argument.



NOTE

The **openstack** CLI can only use the Block Storage REST API version of 3.0 because it does not support these microversions.

2.1. BLOCK STORAGE SERVICE BACK ENDS

Red Hat OpenStack Platform (RHOSP) is deployed using director. Doing so helps ensure the correct configuration of each service, including the Block Storage service (cinder) and, by extension, its back end. Director also has several integrated back-end configurations.

RHOSP supports [Red Hat Ceph Storage](#) and NFS as Block Storage service back ends. By default, the Block Storage service uses an LVM back end as a repository for volumes. While this back end is suitable for test environments, LVM is not supported in production environments.

For instructions on how to deploy Red Hat Ceph Storage with RHOSP, see [Deploying Red Hat Ceph Storage and OpenStack Platform together with director](#).

You can also configure the Block Storage service to use supported third-party storage appliances. Director includes the necessary components for deploying different back-end solutions.

For a complete list of supported Block Storage service back-end appliances and drivers, see [Cinder in Component, Plug-In, and Driver Support in Red Hat OpenStack Platform](#). All third-party back-end appliances and drivers have additional deployment guides. Review the appropriate deployment guide to determine if a back-end appliance or driver requires a plugin.

If you configured Block Storage to use multiple back ends, you must create a volume type for each back end. If you do not specify a back end when creating the volume, the Block Storage scheduler uses filters to select suitable back ends. For more information, see [Configuring the default Block Storage scheduler filters](#).

Additional resources

- [Creating and configuring a volume type](#)

2.2. ACTIVE-ACTIVE BLOCK STORAGE FOR HIGH AVAILABILITY

In active-passive mode, if the Block Storage service fails in a hyperconverged deployment, node fencing is undesirable. This is because node fencing can trigger storage to be rebalanced unnecessarily. Edge sites do not deploy Pacemaker, although Pacemaker is still present at the control site. Instead, edge sites deploy the Block Storage service in an active-active configuration to support highly available hyperconverged deployments.

Active-active deployments improve scaling, performance, and reduce response time by balancing workloads across all available nodes. Deploying the Block Storage service in an active-active configuration creates a highly available environment that maintains the management layer during partial network outages and single- or multi-node hardware failures. Active-active deployments allow a cluster to continue providing Block Storage services during a node outage.

Active-active deployments do not, however, enable workflows to resume automatically. If a service stops, individual operations running on the failed node will also fail during the outage. In this situation, confirm that the service is down and initiate a cleanup of resources that had in-flight operations.

2.2.1. Enabling active-active Block Storage

The **cinder-volume-active-active.yaml** file enables you to deploy the Block Storage service in an active-active configuration. This file ensures director uses the non-Pacemaker cinder-volume heat template and adds the **etcd** service to the deployment as a distributed lock manager (DLM).

The **cinder-volume-active-active.yaml** file also defines the active-active cluster name by assigning a value to the **CinderVolumeCluster** parameter. **CinderVolumeCluster** is a global Block Storage parameter. Therefore, you cannot include clustered (active-active) and non-clustered back ends in the same deployment.



IMPORTANT

Currently, active-active configuration for Block Storage works only with Ceph RADOS Block Device (RBD) back ends. If you plan to use multiple back ends, all back ends must support the active-active configuration. If a back end that does not support the active-active configuration is included in the deployment, that back end will not be available for storage. In an active-active deployment, you risk data loss if you save data on a back end that does not support the active-active configuration.

Prerequisites

- The undercloud is installed. For more information, see [Installing director](#) in *Director Installation and Usage*.

Procedure

1. Log in to the undercloud host as the **stack** user.
2. Source the **stackrc** undercloud credentials file:

```
$ source ~/stackrc
```

3. To enable active-active Block Storage service volumes, add this environment file to the stack with your other environment files and deploy the overcloud:

```
/usr/share/openstack-tripleo-heat-templates/environments/cinder-volume-active-active.yaml
```

2.2.2. Maintenance commands for active-active Block Storage configurations

After deploying an active-active Block Storage configuration, you can use the following commands to manage the clusters and their services.



NOTE

These commands need a Block Storage (cinder) REST API microversion of 3.17 or later.

User goal	Command
To see detailed information about all the services, such as: binary, host, zone, status, state, cluster, disabled reason, and the cluster name.	cinder --os-volume-api-version 3.17 service-list
To see detailed information about all the clusters, such as: name, binary, state, and status. NOTE: When deployed by director for the Ceph back end, the default cluster name is tripleo@tripleo_ceph .	cinder --os-volume-api-version 3.17 cluster-list
To see detailed information about a specific clustered service.	cinder --os-volume-api-version 3.17 cluster-show <cluster_name>
To enable a clustered service.	cinder --os-volume-api-version 3.17 cluster-enable <cluster_name>
To disable a clustered service.	cinder --os-volume-api-version 3.17 cluster-disable <cluster_name>

2.2.3. Volume manage and unmanage

The unmanage and manage mechanisms facilitate moving volumes from one service using version X to another service using version X+1. Both services remain running during this process.

In Block Storage (cinder) REST API microversion 3.17 or later, you can list the volumes and snapshots that can be managed in Block Storage clusters. To see these lists, use the **--cluster** argument with **cinder manageable-list** or **cinder snapshot-manageable-list**.

In Block Storage REST API microversion 3.16 and later, you can use the optional **--cluster** argument of the **cinder manage** command to add unmanaged volumes to a Block Storage cluster.

2.2.4. Volume migration on a clustered service

With Block Storage (cinder) REST API microversion 3.16 and later, the **cinder migrate** and **cinder-manage** commands use the **--cluster** argument to define the destination for active-active deployments.

When you migrate a volume on a Block Storage clustered service, use the optional **--cluster** argument and omit the **host** positional argument, because these arguments are mutually exclusive.

2.2.5. Initiating Block Storage service maintenance

All Block Storage volume services perform their own maintenance when they start.

In an environment with multiple volume services grouped in a cluster, you can clean up services that are not currently running.

The command **work-cleanup** triggers server cleanups. The command returns:

- A list of the services that the command can clean.
- A list of the services that the command cannot clean because they are not currently running in the cluster.

Prerequisites

- You must be a project administrator to initiate Block Storage service maintenance.
- Block Storage (cinder) REST API microversion 3.24 or later.

Procedure

1. Source the overcloud credentials file:

```
$ source ~/<credentials_file>
```

- Replace **<credentials_file>** with the name of your credentials file, for example, **overcloudrc**.

2. Run the following command to verify whether all of the services for a cluster are running:

```
$ cinder cluster-list --detailed
```

Alternatively, run the **cluster show** command.

3. If any services are not running, run the following command to identify those specific services:

```
$ cinder service-list
```

4. Run the following command to trigger the server cleanup:

```
$ cinder --os-volume-api-version 3.24 work-cleanup [--cluster <cluster-name>] [--host
<hostname>] [--binary <binary>] [--is-up <True|true|False|false>] [--disabled
<True|true|False|false>] [--resource-id <resource-id>] [--resource-type <Volume|Snapshot>]
```



NOTE

Filters, such as **--cluster**, **--host**, and **--binary**, define what the command cleans. You can filter on cluster name, host name, type of service, and resource type, including a specific resource. If you do not apply filtering, the command attempts to clean everything that can be cleaned.

The following example filters by cluster name:

```
$ cinder --os-volume-api-version 3.24 work-cleanup --cluster
tripleo@tripleo_ceph
```

2.3. GROUP VOLUME CONFIGURATION WITH VOLUME TYPES

With Red Hat OpenStack Platform you can create volume types so that you can apply associated settings to each volume type. You can assign the required volume type before and after you create a volume. For more information, see [Creating Block Storage volumes](#) and [Block Storage volume retyping](#). The following list shows some of the associated settings that you can apply to a volume type:

- The encryption of a volume. For more information, see [Block Storage service \(cinder\) volume encryption](#).
- The back end that a volume uses. For more information, see [Volume allocation on multiple back ends](#) and [Moving volumes between back ends](#).
- The associated list of Quality of Service (QoS) performance limits or QoS specification for a volume. For more information, see [Block Storage service \(cinder\) Quality of Service specifications](#).

Settings are associated with volume types using key-value pairs called Extra Specs. When you specify a volume type during volume creation, the Block Storage scheduler applies these key-value pairs as settings. You can associate multiple key-value pairs to the same volume type.

You can create volume types to provide different levels of performance for your cloud users:

- Add specific performance, resilience, and other Extra Specs as key-value pairs to each volume type.
- Associate different lists of QoS performance limits or QoS specifications to your volume types.

When your users create their volumes, they can select the appropriate volume type that fulfills their performance requirements.

If you create a volume and do not specify a volume type, then Block Storage uses the default volume type. You can use the Block Storage (cinder) configuration file to define the general default volume type that applies to all your projects (tenants). But if your deployment uses project-specific volume types, ensure that you define default volume types for each project. In this case, Block Storage uses the project-specific volume type instead of the general default volume type. For more information, see [Defining a project-specific default volume type](#).

Additional resources

- [Creating and configuring a volume type](#)
- [Block Storage service back ends](#)
- [Associating a QoS specification with a volume type by using the Dashboard](#)
- [Configuring Block Storage service volume encryption with the Dashboard](#)

2.3.1. Listing back-end driver properties

The properties associated with volume types use key-value pairs called Extra Specs. Each volume type back-end driver supports their own set of Extra Specs. For more information on which Extra Specs a driver supports, see the back-end driver documentation.

Alternatively, you can query the Block Storage host directly to list the well-defined standard Extra Specs of its back-end driver.

Prerequisites

- You must be a project administrator to query the Block Storage host directly.

Procedure

1. Source the overcloud credentials file:

```
$ source ~/<credentials_file>
```

- Replace **<credentials_file>** with the name of your credentials file, for example, **overcloudrc**.

2. Determine the host of **cinder-volume**:

```
$ cinder service-list
```

This command will return a list containing the host of each Block Storage service (**cinder-backup**, **cinder-scheduler**, and **cinder-volume**). For example:

```
+-----+-----+-----+-----+
| Binary | Host | Zone | Status ...
+-----+-----+-----+-----+
| cinder-backup | localhost.localdomain | nova | enabled ...
| cinder-scheduler | localhost.localdomain | nova | enabled ...
| cinder-volume | *localhost.localdomain@lvm* | nova | enabled ...
+-----+-----+-----+-----+
```

3. Display the driver capabilities to determine the supported Extra Specs of a Block Storage service:

```
$ cinder get-capabilities <volsvchost>
```

- Replace **<volsvchost>** with the host of **cinder-volume**. For example:

```
$ cinder get-capabilities localhost.localdomain@lvm
```

```

+-----+-----+
| Volume stats |           Value           |
+-----+-----+
| description  |           None           |
| display_name |           None           |
| driver_version |          3.0.0           |
| namespace    | OS::Storage::Capabilities::localhost.loc...
| pool_name    |           None           |
| storage_protocol |          iSCSI           |
| vendor_name  |          Open Source     |
| visibility   |           None           |
| volume_backend_name |          lvm           |
+-----+-----+
+-----+-----+
| Backend properties |           Value           |
+-----+-----+
| compression    | {u'type': u'boolean', u'description'...
| qos            | {u'type': u'boolean', u'des ...
| replication    | {u'type': u'boolean', u'description'...
| thin_provisioning | {u'type': u'boolean', u'description': u'S...
+-----+-----+

```

The **Backend properties** column shows a list of Extra Spec Keys that you can set, while the **Value** column provides information on valid corresponding values.

2.3.2. Creating and configuring a volume type

You can create volume types so that you can apply associated settings to each volume type. For instance, you can create volume types to provide different levels of performance for your cloud users:

- Add specific performance, resilience, and other Extra Specs as key-value pairs to each volume type.
- Associate different lists of QoS performance limits or QoS specifications to your volume types. For more information, see [Block Storage service \(cinder\) Quality of Service specifications](#).

When your users create their volumes, they can select the appropriate volume type that fulfills their performance requirements.

Prerequisites

- You must be a project administrator to create and configure volume types.
- Access to the Red Hat OpenStack Platform (RHOSP) Dashboard (horizon). For more information, see [Introduction to the OpenStack Dashboard](#).

Procedure

1. Log into the dashboard as an admin user.
2. Select **Admin > Volumes > Volume Types**
3. Click **Create Volume Type**
4. Enter the volume type name in the **Name** field.

5. Click **Create Volume Type**. The new type appears in the **Volume Types** table.
6. Select the volume type's **View Extra Specs** action.
7. Click **Create** and specify the **Key** and **Value**. The key-value pair must be valid; otherwise, specifying the volume type during volume creation will result in an error.
8. Click **Create**. The associated setting (key-value pair) now appears in the **Extra Specs** table.

By default, all volume types are accessible to all OpenStack projects. If you need to create volume types with restricted access, you will need to do so through the CLI. For instructions, see [Creating and configuring private volume types](#).

Next steps

- [Associating a QoS specification with a volume type by using the Dashboard](#)

2.3.3. Editing a volume type

Edit a volume type in the dashboard to modify the **Extra Specs** configuration of the volume type. You can also delete a volume type.

Prerequisites

- You must be a project administrator to edit or delete volume types.
- Access to the Red Hat OpenStack Platform (RHOSP) Dashboard (horizon). For more information, see [Introduction to the OpenStack Dashboard](#).

Procedure

1. Log into the dashboard as an admin user.
2. Select **Admin > Volumes > Volume Types**
3. In the **Volume Types** table, select the volume type's **View Extra Specs** action.
4. On the **Extra Specs** table of this page, you can:
 - Add a new setting to the volume type. To do this, click **Create** and specify the key/value pair of the new setting you want to associate to the volume type.
 - Edit an existing setting associated with the volume type by selecting the setting's **Edit** action.
 - Delete existing settings associated with the volume type by selecting the extra specs' check box and clicking **Delete Extra Specs** in this and the next dialog screen.

To delete a volume type, select its corresponding check boxes from the **Volume Types** table and click **Delete Volume Types**

2.3.4. Creating and configuring private volume types

By default, all volume types are available to all projects (tenants). You can create a restricted volume type by marking it **private**. To do so, set the **is-public** flag of the volume type to **false**, because the default value for this flag is true.

Private volume types are useful for restricting access to volumes with certain attributes. Typically, these are settings that should only be usable by specific projects. For instance, new back ends or ultra-high performance configurations that are being tested.

Prerequisites

- You must be a project administrator to create, view, or configure access for private volume types.

Procedure

1. Source the overcloud credentials file:

```
$ source ~/<credentials_file>
```

- Replace **<credentials_file>** with the name of your credentials file, for example, **overcloudrc**.

2. Create a new cinder volume type and set the **is-public** flag to **false**:

```
$ cinder type-create --is-public false <type_name>
```

- Replace **<type_name>** with the name that you want to call this new private volume type.

By default, private volume types are only accessible to their creators. However, admin users can find and view private volume types, by using the following command:

```
$ cinder type-list
```

This command lists the name and ID of both public and private volume types. You need the ID of the volume type to provide access to it.

Access to a private volume type is granted at the project level. You therefore need to know the ID of the required project. If you do not know this tenant ID but you do know a name of a user of this project, then run:



NOTE

If you are unsure of this user name, the **openstack user list** command lists the name and ID of all the configured users.

```
$ openstack user show <user_name>
```

- Replace **<user_name>** with the name of a user of the required project to display a list of the user details, including the **tenantId** of the project to which this user is associated.

To grant a project access to a private volume type, run:

```
$ cinder type-access-add --volume-type <type_id> --project-id <tenant_id>
```

- Replace **<type_id>** with the ID of the required private volume type.
- Replace **<tenant_id>** with the required tenant ID.

To view which projects have access to a private volume type, run:

```
$ cinder type-access-list --volume-type <type_id>
```

To remove a project from the access list of a private volume type, run:

```
$ cinder type-access-remove --volume-type <type_id> --project-id <tenant_id>
```

2.3.5. Defining a project-specific default volume type

Optional: For complex deployments, project administrators can define a default volume type for each project (tenant).

If you create a volume and do not specify a volume type, then Block Storage uses the default volume type.

You can use the **default_volume_type** option of the Block Storage (cinder) configuration file **cinder.conf** to define the general default volume type that applies to all your projects.

But if your Red Hat OpenStack Platform (RHOSP) deployment uses project-specific volume types, ensure that you define default volume types for each project. In this case, Block Storage uses the project-specific volume type instead of the general default volume type. The following RHOSP deployment examples need project-specific default volume types:

- A distributed RHOSP deployment spanning many availability zones (AZs). Each AZ is in its own project and has its own volume types.
- A RHOSP deployment for three different departments of a company. Each department is in its own project and has its own specialized volume type.

Prerequisites


- At least one volume type in each project that will be the project-specific default volume type. For more information, see [Creating and configuring a volume type](#).
- Block Storage REST API microversion 3.62 or later.
- Only project administrators can define, clear, or list default volume types for their projects.

Procedure

1. Source the overcloud credentials file:

```
$ source ~/<credentials_file>
```

- Replace **<credentials_file>** with the name of your credentials file, for example, **overcloudrc**.
2. Define, clear, or list the default volume type for a project:

**NOTE**

You must replace **<project_id>** in these commands, with the ID of the required project. To find the ID and name of each tenant, run the **openstack project list** command.

- To define the default volume type for a project:

```
$ cinder --os-volume-api-version 3.62 default-type-set <volume_type> <project_id>
```

- Replace **<volume_type>** with the name or ID of the required volume type. You can run the **cinder type-list** command to list the name and ID of all the volume types.

- To clear the default volume type for a project:

```
$ cinder --os-volume-api-version 3.62 default-type-unset <project_id>
```

- To list the default volume type for a project:

```
$ cinder --os-volume-api-version 3.62 default-type-list --project <project_id>
```

2.4. CREATING AND CONFIGURING AN INTERNAL PROJECT FOR THE BLOCK STORAGE SERVICE (CINDER)

Some Block Storage features (for example, the Image-Volume cache) require the configuration of an *internal tenant*. The Block Storage service uses this tenant/project to manage block storage items that do not necessarily need to be exposed to normal users. Examples of such items are images cached for frequent volume cloning or temporary copies of volumes being migrated.

Procedure

1. To configure an internal project, first create a generic project and user, both named **cinder-internal**. To do so, log in to the Controller node and run:

```
$ openstack project create --enable --description "Block Storage Internal Project" cinder-internal
+-----+-----+
| Property | Value |
+-----+-----+
| description | Block Storage Internal Tenant |
| enabled | True |
| id | cb91e1fe446a45628bb2b139d7dccaef |
| name | cinder-internal |
+-----+-----+
$ openstack user create --project cinder-internal cinder-internal
+-----+-----+
| Property | Value |
+-----+-----+
| email | None |
| enabled | True |
| id | 84e9672c64f041d6bfa7a930f558d946 |
| name | cinder-internal |
```

```
|project_id| cb91e1fe446a45628bb2b139d7dccaef |
|username|      cinder-internal      |
+-----+-----+

```

2.5. CONFIGURING THE IMAGE-VOLUME CACHE

The Block Storage service features an optional *Image-Volume cache* which can be used when creating volumes from images. This cache is designed to improve the speed of volume creation from frequently-used images. For information on how to create volumes from images, see [Creating Block Storage volumes](#).

When enabled, the Image-Volume cache stores a copy of an image the first time a volume is created from it. This stored image is cached locally to the Block Storage back end to help improve performance the next time the image is used to create a volume. The limit of the Image-Volume cache can be set to a size (in GB), number of images, or both.

The Image-Volume cache is supported by several back ends. If you are using a third-party back end, refer to its documentation for information on Image-Volume cache support.

Prerequisites

- An *internal tenant* has been configured for the Block Storage service. For more information, see [Creating and configuring an internal project for the Block Storage service \(cinder\)](#).
- The undercloud is installed. For more information, see [Installing director](#) in *Director Installation and Usage*.

Procedure

1. Log in to the undercloud host as the **stack** user.
2. Source the **stackrc** undercloud credentials file:

```
$ source ~/stackrc
```

3. To enable and configure the Image-Volume cache on a back end, you must add the following values to an **ExtraConfig** section of an environment file included in your overcloud deployment command:

```
parameter_defaults:
  ExtraConfig:
    cinder::config::cinder_config:
      DEFAULT/cinder_internal_tenant_project_id:
        value: TENANTID 1
      DEFAULT/cinder_internal_tenant_user_id:
        value: USERID 2
      BACKEND/image_volume_cache_enabled: 3
        value: True
      BACKEND/image_volume_cache_max_size_gb:
        value: MAXSIZE 4
      BACKEND/image_volume_cache_max_count:
        value: MAXNUMBER 5
```


- 1 Replace *TENANTID* with the ID of the **cinder-internal** project.
- 2 Replace *USERID* with the ID of the **cinder-internal** user.
- 3 Replace *BACKEND* with the name of the target back end (specifically, its **volume_backend_name** value).
- 4 By default, the Image-Volume cache size is only limited by the back end. Set *MAXSIZE* to the required size in GB.
- 5 Set *MAXNUMBER* to the maximum number of images.

The Block Storage service database uses a time stamp to track when each cached image was last used to create an image. If either or both *MAXSIZE* and *MAXNUMBER* are set, the Block Storage service will delete cached images as needed to make way for new ones. Cached images with the oldest time stamp are deleted first whenever the Image-Volume cache limits are met.

4. Save the updates to your environment file.
5. Add your environment file to the stack with your other environment files and deploy the overcloud.

2.6. BLOCK STORAGE SERVICE (CINDER) QUALITY OF SERVICE SPECIFICATIONS

You can apply performance limits to volumes that your cloud users create, by creating and associating Quality of Service (QoS) specifications to each volume type. For example, volumes that use higher performance QoS specifications could provide your users with more IOPS or users could assign lighter workloads to volumes that use lower performance QoS specifications to conserve resources.



NOTE

You must be a project administrator to create, configure, associate, and disassociate QoS specifications.

When you create a QoS specification you must choose the required consumer. The consumer determines where you want to apply the QoS limits and determines which QoS property keys are available to define the QoS limits. For more information about the available consumers, see [Consumers of QoS specifications](#).

You can create volume performance limits by setting the required QoS property keys to your deployment specific values. For more information on the QoS property keys provided by the Block Storage service (cinder), see [Block Storage QoS property keys](#).

To create a QoS specification and associate it with a volume type, complete the following tasks:

1. Create and configure the QoS specification.
2. Associate the QoS specification with a volume type.

You can create, configure, and associate a QoS specification to a volume type by using the Dashboard, or by using the CLI.

2.6.1. Consumers of QoS specifications

When you create a QoS specification you must choose the required consumer. The consumer determines where you want to apply the QoS limits and determines which QoS property keys are available to define the QoS limits. The Block Storage service (cinder) supports the following consumers of QoS specifications:

- **front-end:** The Compute service (nova) applies the QoS limits when the volume is attached to an instance. The Compute service supports all the QoS property keys provided by the Block Storage service.
- **back-end:** The back-end driver of the associated volume type applies the QoS limits. Each back-end driver supports their own set of QoS property keys. For more information on which QoS property keys the driver supports, see the back-end driver documentation. You would use the **back-end** consumer in cases where the **front-end** consumer is not supported. For instance, when attaching volumes to bare metal nodes through the Bare Metal Provisioning service (ironic).
- **both:** Both consumers apply the QoS limits, where possible. This consumer type therefore supports the following QoS property keys:
 - When a volume is attached to an instance, then you can use every QoS property key that both the Compute service and the back-end driver supports.
 - When the volume is not attached to an instance, then you can only use the QoS property keys that the back-end driver supports.

2.6.2. Block Storage QoS property keys

The Block Storage service provides you with QoS property keys so that you can limit the performance of the volumes that your cloud users create. These limits use the following two industry standard measurements of storage volume performance:

- Input/output operations per second (IOPS)
- Data transfer rate, measured in bytes per second

The consumer of the QoS specification determines which QoS property keys are supported. For more information, see [Consumers of QoS specifications](#).

Block Storage cannot perform error checking of QoS property keys, because some QoS property keys are defined externally by back-end drivers. Therefore, Block Storage ignores any invalid or unsupported QoS property key.



IMPORTANT

Ensure that you spell the QoS property keys correctly. The volume performance limits that contain incorrectly spelled property keys are ignored.

For both the IOPS and data transfer rate measurements, you can configure the following performance limits:

Fixed limits

Typically, fixed limits should define the average usage of the volume performance measurement.

Burst limits

Typically, burst limits should define periods of intense activity of the volume performance measurement. A burst limit makes allowance for an increased rate of activity for a specific time, while keeping the fixed limits low for average usage.

**NOTE**

The burst limits all use a burst length of 1 second.

Total limits

Specify a global limit for both the read and write operations of the required performance limit, by using the **total_*** QoS property key.

**NOTE**

Instead of using a total limit you can apply separate limits to the read and write operations or choose to limit only the read or write operations.

Read limits

Specify a limit that only applies to the read operations of the required performance limit, by using the **read_*** QoS property key.

**NOTE**

This limit is ignored when you specify a total limit.

Write limits

Specify a limit that only applies to the write operations of the required performance limit, by using the **write_*** QoS property key.

**NOTE**

This limit is ignored when you specify a total limit.

You can use the following Block Storage QoS property keys to create volume performance limits for your deployment:

**NOTE**

The default value for all QoS property keys is **0**, which means that the limit is unrestricted.

Table 2.1. Block Storage QoS property keys

Performance limit	Measurement unit	QoS property keys
Fixed IOPS	IOPS	total_iops_sec read_iops_sec write_iops_sec

Performance limit	Measurement unit	QoS property keys
<p>Fixed IOPS calculated by the size of the volume.</p> <p>For more information about the usage restrictions of these limits, see QoS limits that scale according to volume size.</p>	IOPS per GB	total_iops_sec_per_gb read_iops_sec_per_gb write_iops_sec_per_gb
Burst IOPS	IOPS	total_iops_sec_max read_iops_sec_max write_iops_sec_max
Fixed data transfer rate	Bytes per second	total_bytes_sec read_bytes_sec write_bytes_sec
Burst data transfer rate	Bytes per second	total_bytes_sec_max read_bytes_sec_max write_bytes_sec_max
<p>Size of an IO request when calculating IOPS limits.</p> <p>For more information, see Set the IO request size for IOPS limits.</p>	Bytes	size_iops_sec

2.6.2.1. Set the IO request size for IOPS limits

If you implement IOPS volume performance limits, you should also specify the typical IO request size to prevent users from circumventing these limits. If you do not then users could submit several large IO requests instead of a lot of smaller ones.

Use the **size_iops_sec** QoS property key to specify the maximum size, in bytes, of a typical IO request. The Block Storage service uses this size to calculate the proportional number of typical IO requests for each IO request that is submitted, for example:

size_iops_sec=4096

- An 8 KB request is counted as two requests.

- A 6 KB request is counted as one and a half requests.
- Any request less than 4 KB is counted as one request.

The Block Storage service only uses this IO request size limit when calculating IOPS limits.



NOTE

The default value of **size_iops_sec** is **0**, which ignores the size of IO requests when applying IOPS limits.

2.6.2.2. IOPS limits that scale according to volume size

You can create IOPS volume performance limits that are determined by the capacity of the volumes that your users create. These Quality of Service (QoS) limits scale with the size of the provisioned volumes. For example, if the volume type has an IOPS limit of 500 per GB of volume size for read operations, then a provisioned 3 GB volume of this volume type would have a read IOPS limit of 1500.



IMPORTANT

The size of the volume is determined when the volume is attached to an instance. Therefore if the size of the volume is changed while it is attached to an instance, these limits are only recalculated for the new volume size when this volume is detached and then reattached to an instance.

You can use the following QoS property keys, specified in IOPS per GB, to create scalable volume performance limits:

- **total_iops_sec_per_gb**: Specify a global IOPS limit per GB of volume size for both the read and write operations.



NOTE

Instead of using a total limit you can apply separate limits to the read and write operations or choose to limit only the read or write operations.

- **read_iops_sec_per_gb**: Specify a IOPS limit per GB of volume size that only applies to the read operations.



NOTE


This limit is ignored when you specify a total limit.

- **write_iops_sec_per_gb**: Specify a IOPS limit per GB of volume size that only applies to the write operations.



NOTE

This limit is ignored when you specify a total limit.



IMPORTANT

The consumer of the QoS specification containing these QoS limits can either be **front-end** or **both**, but not **back-end**. For more information, see [Consumers of QoS specifications](#).

2.6.3. Creating and configuring a QoS specification with the Dashboard

A Quality of Service (QoS) specification is a list of volume performance QoS limits. You create each QoS limit by setting a QoS property key to your deployment specific value. To apply the QoS performance limits to a volume, you must associate the QoS specification with the required volume type.

Prerequisites

- You must be a project administrator to create, configure, associate, and disassociate QoS specifications.
- Access to the Red Hat OpenStack Platform (RHOSP) Dashboard (horizon). For more information, see [Introduction to the OpenStack Dashboard](#).

Procedure

1. Log into the dashboard as an admin user.
2. Select **Admin > Volumes > Volume Types**
3. On the **QoS Specs** table, click **Create QoS Spec**
4. Enter a name for the **QoS Spec**.
5. In the **Consumer** field, choose the consumer of this QoS specification. For more information, see [Consumers of QoS specifications](#).
6. Click **Create**. The new QoS specification is displayed in the **QoS Specs** table.
7. In the **QoS Specs** table, select the **Manage Specs** action of your new QoS specification to open the **Specs** window, where you add the QoS performance limits.
8. Click **Create** in the **Specs** window to open the **Create Extra Specs** window.
9. Specify the QoS property key for a QoS performance limit in the **Key** field, and set the performance limit value in the **Value** field. For more information on the available property keys, see [Block Storage QoS property keys](#).



IMPORTANT

Ensure that you spell the QoS property keys correctly. The volume performance limits that contain incorrectly spelled property keys are ignored.

10. Click **Create** to add the QoS limit to your QoS specification.
11. Repeat steps 7 to 10 for each QoS limit that you want to add to your QoS specification.

Next steps

- [Associating a QoS specification with a volume type by using the Dashboard](#)

2.6.4. Creating and configuring a QoS specification with the CLI

A Quality of Service (QoS) specification is a list of volume performance QoS limits. You create each QoS limit by setting a QoS property key to your deployment specific value. To apply the QoS performance limits to a volume, you must associate the QoS specification with the required volume type.

Prerequisites

- You must be a project administrator to create, configure, associate, and disassociate QoS specifications.

Procedure

1. Source the overcloud credentials file:

```
$ source ~/<credentials_file>
```

- Replace **<credentials_file>** with the name of your credentials file, for example, **overcloudrc**.

2. Create the QoS specification:


```
$ openstack volume qos create [--consumer <qos_spec_consumer>] <qos_spec_name>
```

- Optional: Replace **<qos_spec_consumer>** with the required consumer of this QoS specification. If not specified, the consumer defaults to **both**. For more information, see [Consumers of QoS specifications](#).
- Replace **<qos_spec_name>** with the name of your QoS specification.

3. Add the performance limits to the QoS specification, by specifying a separate **--property <key=value>** argument for each QoS limit that you want to add:

```
$ openstack volume qos set --property <key>=<value> <qos_spec_name>
```

- Replace **<key>** with the QoS property key of the required performance constraint. For more information, see [Block Storage QoS property keys](#).



IMPORTANT

Ensure that you spell the QoS property keys correctly. The volume performance limits that contain incorrectly spelled property keys are ignored.

- Replace **<value>** with your deployment-specific limit for this performance constraint, in the measurement unit required by the QoS property key.
- Replace **<qos_spec_name>** with the name or ID of your QoS specification.
Example:

```
$ openstack volume qos set \  
--property read_iops_sec=5000 \  

```

```
--property write_iops_sec=7000 \
myqoslimits
```

- Review the QoS specification:

```
$ openstack volume qos list
+-----+-----+-----+-----+
+-----+
| ID                | Name   | Consumer | Associations | Properties
|
+-----+-----+-----+-----+
+-----+
| 204c6ba2-c67c-4ac8-918a-03f101811235 | myqoslimits | front-end |           |
read_iops_sec='5000', write_iops_sec='7000' |
+-----+-----+-----+-----+
+-----+
```


This command provides a table of the configuration details of all the configured QoS specifications.

Next steps

- [Associating a QoS specification with a volume type by using the CLI](#)

2.6.5. Associating a QoS specification with a volume type by using the Dashboard

You must associate a Quality of Service (QoS) specification with an existing volume type to apply the QoS limits to volumes.



IMPORTANT

If a volume is already attached to an instance, then the QoS limits are only applied to this volume when the volume is detached and then reattached to this instance.

Prerequisites

- You must be a project administrator to create, configure, associate, and disassociate QoS specifications.
- The required volume type is created. For more information, see [Creating and configuring a volume type](#).
- The required QoS specification is created. For more information, see [Creating and configuring a QoS specification with the Dashboard](#).
- Access to the Red Hat OpenStack Platform (RHOSP) Dashboard (horizon). For more information, see [Introduction to the OpenStack Dashboard](#).

Procedure

- Log into the dashboard as an admin user.
- Select **Admin > Volumes > Volume Types**

3. In the **Volume Types** table, select the **Manage QoS Spec Association** action of the required volume type.
4. Select the required QoS specification from the **QoS Spec to be associated** list.
5. Click **Associate**. The QoS specification is added to the **Associated QoS Spec** column of the edited volume type.

2.6.6. Associating a QoS specification with a volume type by using the CLI

You must associate a Quality of Service (QoS) specification with an existing volume type to apply the QoS limits to volumes.



IMPORTANT

If a volume is already attached to an instance, then the QoS limits are only applied to this volume when the volume is detached and then reattached to this instance.

Prerequisites

- You must be a project administrator to create, configure, associate, and disassociate QoS specifications.
- The required volume type is created. For more information, see [Creating and configuring a volume type](#).
- The required QoS specification is created. For more information, see [Creating and configuring a QoS specification with the CLI](#).

Procedure

1. Source the overcloud credentials file:

```
$ source ~/<credentials_file>
```

- Replace **<credentials_file>** with the name of your credentials file, for example, **overcloudrc**.

2. Associate the required QoS specification with the required volume type:

```
$ openstack volume qos associate <qos_spec_name> <volume_type>
```

- Replace **<qos_spec_name>** with the name or ID of the QoS specification. You can run the **openstack volume qos list** command to list the name and ID of all the QoS specifications.
- Replace **<volume_type>** with the name or ID of the volume type. You can run the **cinder type-list** command to list the name and ID of all the volume types.

3. Verify that the QoS specification has been associated:

```
$ openstack volume qos list
```

The **Associations** column of the output table shows which volume types are associated with this QoS specification.

2.6.7. Disassociating a QoS specification from a volume type with the Dashboard

You can disassociate a Quality of Service (QoS) specification from a volume type when you no longer want the QoS limits to be applied to volumes of that volume type.



IMPORTANT

If a volume is already attached to an instance, then the QoS limits are only removed from this volume when the volume is detached and then reattached to this instance.

Prerequisites

- You must be a project administrator to create, configure, associate, and disassociate QoS specifications.
- Access to the Red Hat OpenStack Platform (RHOSP) Dashboard (horizon). For more information, see [Introduction to the OpenStack Dashboard](#).

Procedure

1. Log into the dashboard as an admin user.
2. Select **Admin > Volumes > Volume Types**
3. In the **Volume Types** table, select the **Manage QoS Spec Association** action of the required volume type.
4. Select **None** from the **QoS Spec to be associated** list.
5. Click **Associate**.
The QoS specification should be removed from the **Associated QoS Spec** column of the edited volume type.

2.6.8. Disassociating a QoS specification from volume types with the CLI

You can disassociate a Quality of Service (QoS) specification from a volume type when you no longer want the QoS limits to be applied to volumes of that volume type.



IMPORTANT

If a volume is already attached to an instance, then the QoS limits are only removed from this volume when the volume is detached and then reattached to this instance.

Prerequisites

- You must be a project administrator to create, configure, associate, and disassociate QoS specifications.

Procedure

1. Source the overcloud credentials file:

```
$ source ~/<credentials_file>
```

- Replace `<credentials_file>` with the name of your credentials file, for example, `overcloudrc`.
2. Disassociate the volume types associated with the QoS specification. You can either disassociate a specific volume type, or all volumes types when more than one volume type is associated to the same QoS specification:

- To disassociate a specific volume type associated with the QoS specification:

```
$ openstack volume qos disassociate <qos_spec_name> --volume-type <volume_type>
```

- Replace `<qos_spec_name>` with the name or ID of the QoS specification. You can run the `openstack volume qos list` command to list the name and ID of all the QoS specifications.
 - Replace `<volume_type>` with the name or ID of the volume type associated with this QoS specification. You can run the `cinder type-list` command to list the name and ID of all the volume types.
- To disassociate all volume types associated with the QoS specification:

```
$ openstack volume qos disassociate <qos_spec_name> --all
```

3. Verify that the QoS specification has been disassociated:

```
$ openstack volume qos list
```

The **Associations** column of this QoS specification should either not specify the volume type or be empty.

2.7. BLOCK STORAGE SERVICE (CINDER) VOLUME ENCRYPTION

Volume encryption helps provide basic data protection in case the volume back-end is either compromised or outright stolen. Both Compute and Block Storage services are integrated to allow instances to read access and use encrypted volumes. You must deploy Barbican to take advantage of volume encryption.



IMPORTANT

- Volume encryption is not supported on file-based volumes (such as NFS).
- Retyping an unencrypted volume to an encrypted volume of the same size is not supported, because encrypted volumes require additional space to store encryption data. For more information about encrypting unencrypted volumes, see [Encrypting unencrypted volumes](#).

Volume encryption is applied through volume type. For information on encrypted volume types, see [Configuring Block Storage service volume encryption with the Dashboard](#) or [Configuring Block Storage service volume encryption with the CLI](#).

For more information, on using the OpenStack Key Manager (barbican) to manage your Block Storage (cinder) encryption keys, see [Encrypting Block Storage \(cinder\) volumes](#).

2.7.1. Configuring Block Storage service volume encryption with the Dashboard

To create encrypted volumes, you first need an *encrypted volume type*. Encrypting a volume type involves setting what provider class, cipher, and key size it should use. You can also re-configure the encryption settings of an encrypted volume type.


You can invoke encrypted volume types to automatically create encrypted volumes.

Prerequisites

- You must be a project administrator to create encrypted volumes.
- Access to the Red Hat OpenStack Platform (RHOSP) Dashboard (horizon). For more information, see [Introduction to the OpenStack Dashboard](#).

Procedure

1. Log into the dashboard as an admin user.
2. Select **Admin > Volumes > Volume Types**
3. In the **Actions** column of the volume to be encrypted, select **Create Encryption** to launch the **Create Volume Type Encryption** wizard.
4. From there, configure the **Provider**, **Control Location**, **Cipher**, and **Key Size** settings of the volume type's encryption. The **Description** column describes each setting.



IMPORTANT

The values listed below are the only supported options for **Provider**, **Cipher**, and **Key Size**.

- a. Enter **luks** for **Provider**.
 - b. Enter **aes-xts-plain64** for **Cipher**.
 - c. Enter **256** for **Key Size**.
5. Click **Create Volume Type Encryption**

You can also re-configure the encryption settings of an encrypted volume type.

1. Select **Update Encryption** from the **Actions** column of the volume type to launch the **Update Volume Type Encryption** wizard.
2. In **Project > Compute > Volumes**, check the **Encrypted** column in the **Volumes** table to determine whether the volume is encrypted.
3. If the volume is encrypted, click **Yes** in that column to view the encryption settings.

Additional resources

- [Configuring Block Storage service volume encryption with the CLI](#)

2.7.2. Configuring Block Storage service volume encryption with the CLI

To create encrypted volumes, you first need an *encrypted volume type*. Encrypting a volume type involves setting what provider class, cipher, and key size it should use.

Prerequisites

- You must be a project administrator to create encrypted volumes.

Procedure

1. Source the overcloud credentials file:

```
$ source ~/<credentials_file>
```

- Replace **<credentials_file>** with the name of your credentials file, for example, **overcloudrc**.

2. Create a volume type:

```
$ cinder type-create myEncType
```

3. Configure the cipher, key size, control location, and provider settings:

```
$ cinder encryption-type-create --cipher aes-xts-plain64 --key-size 256 --control-location front-end myEncType luks
```

4. Create an encrypted volume:

```
$ cinder --debug create 1 --volume-type myEncType --name myEncVol
```

2.7.3. Automatic deletion of volume image encryption key

The Block Storage service (cinder) creates an encryption key in the Key Management service (barbican) when it uploads an encrypted volume to the Image service (glance). This creates a 1:1 relationship between an encryption key and a stored image.

Encryption key deletion prevents unlimited resource consumption of the Key Management service. The Block Storage, Key Management, and Image services automatically manage the key for an encrypted volume, including the deletion of the key.

The Block Storage service automatically adds two properties to a volume image:

- **cinder_encryption_key_id** - The identifier of the encryption key that the Key Management service stores for a specific image.
- **cinder_encryption_key_deletion_policy** - The policy that tells the Image service to tell the Key Management service whether to delete the key associated with this image.



IMPORTANT

The values of these properties are automatically assigned. **To avoid unintentional data loss, do not adjust these values.**

When you create a volume image, the Block Storage service sets the

cinder_encryption_key_deletion_policy property to **on_image_deletion**. When you delete a volume image, the Image service deletes the corresponding encryption key if the **cinder_encryption_key_deletion_policy** equals **on_image_deletion**.



IMPORTANT

Red Hat does not recommend manual manipulation of the **cinder_encryption_key_id** or **cinder_encryption_key_deletion_policy** properties. If you use the encryption key that is identified by the value of **cinder_encryption_key_id** for any other purpose, you risk data loss.

2.8. DEPLOYING AVAILABILITY ZONES FOR BLOCK STORAGE VOLUME BACK ENDS

An availability zone is a provider-specific method of grouping cloud instances and services. Director uses **CinderXXXAvailabilityZone** parameters (where **XXX** is associated with a specific back end) to configure different availability zones for Block Storage volume back ends.

Prerequisites

- The undercloud is installed. For more information, see [Installing director](#) in *Director Installation and Usage*.

Procedure

1. Log in to the undercloud host as the **stack** user.
2. Source the **stackrc** undercloud credentials file:

```
$ source ~/stackrc
```

3. Add the following parameters to the environment file to create two availability zones:

```
parameter_defaults:
  CinderXXXAvailabilityZone: zone1
  CinderYYYAvailabilityZone: zone2
```

- Replace **XXX** and **YYY** with supported back-end values, such as:

```
CinderISCSIAvailabilityZone
CinderNfsAvailabilityZone
CinderRbdAvailabilityZone
```



NOTE

Search the **/usr/share/openstack-tripleo-heat-templates/deployment/cinder/** directory for the heat template associated with your back end for the correct back end value.

The following example deploys two back ends where **rbd** is zone 1 and **iSCSI** is zone 2:


```
parameter_defaults:
  CinderRbdAvailabilityZone: zone1
  CinderISCSIAvailabilityZone: zone2
```

4. Save the updates to your environment file.
5. Add your environment file to the stack with your other environment files and deploy the overcloud.

2.9. BLOCK STORAGE SERVICE (CINDER) CONSISTENCY GROUPS

You can use the Block Storage (cinder) service to set consistency groups to group multiple volumes together as a single entity. This means that you can perform operations on multiple volumes at the same time instead of individually. You can use consistency groups to create snapshots for multiple volumes simultaneously. This also means that you can restore or clone those volumes simultaneously.

A volume can be a member of multiple consistency groups. However, you cannot delete, retype, or migrate volumes after you add them to a consistency group.

2.9.1. Configuring Block Storage service consistency groups

By default, Block Storage security policy disables consistency groups APIs. You must enable it here before you use the feature. The related consistency group entries in the `/etc/cinder/policy.json` file of the node that hosts the Block Storage API service, **openstack-cinder-api** list the default settings:

```
"consistencygroup:create" : "group:nobody",
"consistencygroup:delete": "group:nobody",
"consistencygroup:update": "group:nobody",
"consistencygroup:get": "group:nobody",
"consistencygroup:get_all": "group:nobody",
"consistencygroup:create_cgsnapshot" : "group:nobody",
"consistencygroup:delete_cgsnapshot": "group:nobody",
"consistencygroup:get_cgsnapshot": "group:nobody",
"consistencygroup:get_all_cgsnapshots": "group:nobody",
```

You must change these settings in an environment file and then deploy them to the overcloud by using the **openstack overcloud deploy** command. Do not edit the JSON file directly because the changes are overwritten next time the overcloud is deployed.

Prerequisites

- The undercloud is installed. For more information, see [Installing director](#) in *Director Installation and Usage*.

Procedure

1. Log in to the undercloud host as the **stack** user.
2. Source the **stackrc** undercloud credentials file:

```
$ source ~/stackrc
```

3. Edit an environment file and add a new entry to the **parameter_defaults** section. This ensures that the entries are updated in the containers and are retained whenever the environment is re-deployed by director with the **openstack overcloud deploy** command.
4. Add a new section to an environment file using **CinderApiPolicies** to set the consistency group settings. The equivalent **parameter_defaults** section with the default settings from the JSON file appear in the following way:

```
parameter_defaults:
  CinderApiPolicies: { \
    cinder-consistencygroup_create: { key: 'consistencygroup:create', value: 'group:nobody' }, \
    \
    cinder-consistencygroup_delete: { key: 'consistencygroup:delete', value: 'group:nobody' }, \
    \
    cinder-consistencygroup_update: { key: 'consistencygroup:update', value: 'group:nobody' \
  }, \
    cinder-consistencygroup_get: { key: 'consistencygroup:get', value: 'group:nobody' }, \
    cinder-consistencygroup_get_all: { key: 'consistencygroup:get_all', value: 'group:nobody' }, \
    \
    cinder-consistencygroup_create_cgsnapshot: { key: \
'consistencygroup:create_cgsnapshot', value: 'group:nobody' }, \
    cinder-consistencygroup_delete_cgsnapshot: { key: \
'consistencygroup:delete_cgsnapshot', value: 'group:nobody' }, \
    cinder-consistencygroup_get_cgsnapshot: { key: 'consistencygroup:get_cgsnapshot', \
value: 'group:nobody' }, \
    cinder-consistencygroup_get_all_cgsnapshots: { key: \
'consistencygroup:get_all_cgsnapshots', value: 'group:nobody' }, \
  }
```

5. The value **'group:nobody'** determines that no group can use this feature so it is effectively disabled. To enable it, change the group to another value.
6. For increased security, set the permissions for both consistency group API and volume type management API to be identical. The volume type management API is set to **"rule:admin_or_owner"** by default in the same **/etc/cinder/policy.json** file:

```
"volume_extension:types_manage": "rule:admin_or_owner",
```

7. To make the consistency groups feature available to all users, set the API policy entries to allow users to create, use, and manage their own consistency groups. To do so, use **rule:admin_or_owner**:

```
CinderApiPolicies: { \
  cinder-consistencygroup_create: { key: 'consistencygroup:create', value: \
'rule:admin_or_owner' }, \
  cinder-consistencygroup_delete: { key: 'consistencygroup:delete', value: \
'rule:admin_or_owner' }, \
  cinder-consistencygroup_update: { key: 'consistencygroup:update', value: \
'rule:admin_or_owner' }, \
  cinder-consistencygroup_get: { key: 'consistencygroup:get', value: 'rule:admin_or_owner' \
}, \
  cinder-consistencygroup_get_all: { key: 'consistencygroup:get_all', value: \
'rule:admin_or_owner' }, \
  cinder-consistencygroup_create_cgsnapshot: { key: \
'consistencygroup:create_cgsnapshot', value: 'rule:admin_or_owner' }, \
```

```

cinder-consistencygroup_delete_cgsnapshot: { key:
'consistencygroup:delete_cgsnapshot', value: 'rule:admin_or_owner' }, \
cinder-consistencygroup_get_cgsnapshot: { key: 'consistencygroup:get_cgsnapshot',
value: 'rule:admin_or_owner' }, \
cinder-consistencygroup_get_all_cgsnapshots: { key:
'consistencygroup:get_all_cgsnapshots', value: 'rule:admin_or_owner' }, \
}

```

8. Save the updates to your environment file.
9. Add your environment file to the stack with your other environment files and deploy the overcloud.

2.9.2. Creating Block Storage consistency groups with the Dashboard

After you enable the consistency groups API, you can start creating consistency groups.

Prerequisites

- You must be a project administrator or a volume owner to create consistency groups.
- Access to the Red Hat OpenStack Platform (RHOSP) Dashboard (horizon). For more information, see [Introduction to the OpenStack Dashboard](#).

Procedure

1. Log into the dashboard as an admin user or a volume owner.
2. Select **Project > Compute > Volumes > Volume Consistency Groups**
3. Click **Create Consistency Group**.
4. In the **Consistency Group Information** tab of the wizard, enter a name and description for your consistency group. Then, specify its **Availability Zone**.
5. You can also add volume types to your consistency group. When you create volumes within the consistency group, the Block Storage service will apply compatible settings from those volume types. To add a volume type, click its + button from the **All available volume types** list.
6. Click **Create Consistency Group**. It appears next in the **Volume Consistency Groups** table.

2.9.3. Managing Block Storage service consistency groups with the Dashboard

You can manage consistency groups for Block Storage volumes in the dashboard.

Prerequisites

- You must be a project administrator to manage consistency groups.
- Access to the Red Hat OpenStack Platform (RHOSP) Dashboard (horizon). For more information, see [Introduction to the OpenStack Dashboard](#).

Procedure

1. Log into the dashboard as an admin user.

2. Select **Project > Compute > Volumes > Volume Consistency Groups**
3. Optional: You can change the name or description of a consistency group by selecting **Edit Consistency Group** from its **Action** column.
4. To add or remove volumes from a consistency group directly, find the consistency group you want to configure. In the **Actions** column of that consistency group, select **Manage Volumes**. This launches the **Add/Remove Consistency Group Volumes** wizard.
 - a. To add a volume to the consistency group, click its + button from the **All available volumes** list.
 - b. To remove a volume from the consistency group, click its - button from the **Selected volumes** list.
5. Click **Edit Consistency Group**.

2.9.4. Creating and managing consistency group snapshots for the Block Storage service

After you add volumes to a consistency group, you can create snapshots from it.

Prerequisites

- You must be a project administrator to create and manage consistency group snapshots.

Procedure

1. Source the overcloud credentials file:

```
$ source ~/<credentials_file>
```

- Replace **<credentials_file>** with the name of your credentials file, for example, **overcloudrc**.

2. List all available consistency groups and their respective IDs:

```
$ cinder consisgroup-list
```

3. Create snapshots using the consistency group:

```
$ cinder cgsnapshot-create [--name <cgsnapname>] [--description "<description>"] <cgnameid>
```

- Replace **<cgsnapname>** with the name of the snapshot.
- Replace **<description>** with a description of the snapshot.
- Replace **<cgnameid>** with the name or ID of the consistency group.

4. Display a list of all available consistency group snapshots:

```
# cinder cgsnapshot-list
```

2.9.5. Cloning Block Storage service consistency groups

You can also use consistency groups to create a whole batch of pre-configured volumes simultaneously. You can do this by cloning an existing consistency group or restoring a consistency group snapshot. Both processes use the same command.

Prerequisites

- You must be a project administrator to clone consistency groups and restore consistency group snapshots.

Procedure

1. Source the overcloud credentials file:

```
$ source ~/<credentials_file>
```

- Replace **<credentials_file>** with the name of your credentials file, for example, **overcloudrc**.

2. To clone an existing consistency group:

```
$ cinder consisgroup-create-from-src --source-cg <cgnameid> [--name <cgname>] [--description "<description>"]
```

- Replace **<cgnameid>** with the name or ID of the consistency group you want to clone.
- Replace **<cgname>** with the name of your consistency group.
- Replace **<description>** with a description of your consistency group.

3. To create a consistency group from a consistency group snapshot:

```
$ cinder consisgroup-create-from-src --cgsnapshot <cgsnapname> [--name <cgname>] [--description "<description>"]
```

- Replace **<cgsnapname>** with the name or ID of the snapshot you are using to create the consistency group.

2.10. CONFIGURING THE DEFAULT BLOCK STORAGE SCHEDULER FILTERS

If the volume back end is not specified during volume creation, then the Block Storage scheduler uses filters to select suitable back ends. Ensure that you configure the following default filters:

AvailabilityZoneFilter

Filters out all back ends that do not meet the availability zone requirements of the requested volume.

CapacityFilter

Selects only back ends with enough space to accommodate the volume.

CapabilitiesFilter

Selects only back ends that can support any specified settings in the volume.

InstanceLocality

Configures clusters to use volumes local to the same node.

Prerequisites

- The undercloud is installed. For more information, see [Installing director](#) in *Director Installation and Usage*.

Procedure

1. Log in to the undercloud host as the **stack** user.
2. Source the **stackrc** undercloud credentials file:

```
$ source ~/stackrc
```

3. Add an environment file to your overcloud deployment command that contains the following parameters:

```
parameter_defaults:  
  ControllerExtraConfig: # 1  
    cinder::config::cinder_config:  
      DEFAULT/scheduler_default_filters:  
        value: 'AvailabilityZoneFilter,CapacityFilter,CapabilitiesFilter,InstanceLocality'
```

- 1 You can also add the **ControllerExtraConfig**: hook and its nested sections to the **parameter_defaults**: section of an existing environment file.

4. Save the updates to your environment file.
5. Add your environment file to the stack with your other environment files and deploy the overcloud.

2.11. ENABLING LVM2 FILTERING ON OVERCLOUD NODES

If you use LVM2 (Logical Volume Management) volumes with certain Block Storage service (cinder) back ends, the volumes that you create inside Red Hat OpenStack Platform (RHOSP) guests might become visible on the overcloud nodes that host **cinder-volume** or **nova-compute** containers. In this case, the LVM2 tools on the host scan the LVM2 volumes that the OpenStack guest creates, which can result in one or more of the following problems on Compute or Controller nodes:

- LVM appears to see volume groups from guests
- LVM reports duplicate volume group names
- Volume detachments fail because LVM is accessing the storage
- Guests fail to boot due to problems with LVM
- The LVM on the guest machine is in a partial state due to a missing disk that actually exists
- Block Storage service (cinder) actions fail on devices that have LVM
- Block Storage service (cinder) snapshots fail to remove correctly

- Errors during live migration: `/etc/multipath.conf` does not exist

To prevent this erroneous scanning, and to segregate guest LVM2 volumes from the host node, you can enable and configure a filter with the `LVMFilterEnabled` heat parameter when you deploy or update the overcloud. This filter is computed from the list of physical devices that host active LVM2 volumes. You can also allow and deny block devices explicitly with the `LVMFilterAllowlist` and `LVMFilterDenylist` parameters. You can apply this filtering globally, to specific node roles, or to specific devices.

Prerequisites

- The undercloud is installed. For more information, see [Installing director](#) in *Director Installation and Usage*.

Procedure

1. Log in to the undercloud host as the `stack` user.
2. Source the `stackrc` undercloud credentials file:

```
$ source ~/stackrc
```

3. Add an environment file to your overcloud deployment command that contains the following parameter:

```
parameter_defaults:
  LVMFilterEnabled: true
```

You can further customize the implementation of the LVM2 filter. For example, to enable filtering only on Compute nodes, use the following configuration:

```
parameter_defaults:
  ComputeParameters:
    LVMFilterEnabled: true
```

These parameters also support regular expression. To enable filtering only on Compute nodes, and ignore all devices that start with `/dev/sd`, use the following configuration:

```
parameter_defaults:
  ComputeParameters:
    LVMFilterEnabled: true
    LVMFilterDenylist:
      - /dev/sd.*
```

4. Save the updates to your environment file.
5. Add your environment file to the stack with your other environment files and deploy the overcloud.

2.12. MULTIPATH CONFIGURATION

Use multipath to configure multiple I/O paths between server nodes and storage arrays into a single device to create redundancy and improve performance.

2.12.1. Using director to configure multipath

You can configure multipath on a Red Hat OpenStack Platform (RHOSP) overcloud deployment for greater bandwidth and networking resiliency.



IMPORTANT

When you configure multipath on an existing deployment, the new workloads are multipath aware. If you have any pre-existing workloads, you must shelve and unshelve the instances to enable multipath on these instances.

Prerequisites

- The undercloud is installed. For more information, see [Installing director](#) in *Director Installation and Usage*.


Procedure

1. Log in to the undercloud host as the **stack** user.
2. Source the **stackrc** undercloud credentials file:

```
$ source ~/stackrc
```

3. Use an overrides environment file or create a new one, for example **multipath_overrides.yaml**. Add and set the following parameter:

```
parameter_defaults:
  ExtraConfig:
    cinder::config::cinder_config:
      backend_defaults/use_multipath_for_image_xfer:
        value: true
```



NOTE

The default settings will generate a basic multipath configuration that works for most environments. However, check with your storage vendor for recommendations, because some vendors have optimized configurations that are specific to their hardware. For more information about multipath, see [Configuring device mapper multipath](#).

4. Optional: If you have a multipath configuration file for your overcloud deployment, then you can use the **MultipathdCustomConfigFile** parameter to specify the location of this file:

```
parameter_defaults:
  MultipathdCustomConfigFile: <config_file_directory>/<config_file_name>
```

In the following example, **/home/stack** is the directory of the multipath configuration file and **multipath.conf** is the name of this file:

```
parameter_defaults:
  MultipathdCustomConfigFile: /home/stack/multipath.conf
```


**NOTE**

Other TripleO multipath parameters override any corresponding value in the local custom configuration file. For example, if **MultipathdEnableUserFriendlyNames** is **False**, the files on the overcloud nodes are updated to match, even if the setting is enabled in the local custom file.

For more information about multipath parameters, see [Multipath heat template parameters](#).

5. Save the updates to your overrides environment file.
6. Add your overrides environment file to the stack with your other environment files, such as:

```
----
/usr/share/openstack-tripleo-heat-templates/environments/multipathd.yaml
----
```

7. Deploy the overcloud.

Additional resources

- [Shelving an instance](#) in *Creating and Managing Instances*

2.12.1.1. Multipath heat template parameters

Use this to understand the following parameters that enable multipath.

Parameter	Description	Default value
MultipathdEnable	Defines whether to enable the multipath daemon. This parameter defaults to True through the configuration contained in the multipathd.yaml file	True
MultipathdEnableUserFriendlyNames	Defines whether to enable the assignment of a user friendly name to each path.	False
MultipathdEnableFindMultipaths	Defines whether to automatically create a multipath device for each path.	True
MultipathdSkipKpartx	Defines whether to skip automatically creating partitions on the device.	True

Parameter	Description	Default value
MultipathdCustomConfigFile	<p>Includes a local, custom multipath configuration file on the overcloud nodes. By default, a minimal multipath.conf file is installed.</p> <p>NOTE: Other TripleO multipath parameters override any corresponding value in any local, custom configuration file that you add. For example, if MultipathdEnableUserFriendlyNames is False, the files on the overcloud nodes are updated to match, even if the setting is enabled in your local, custom file.</p>	

2.12.2. Verifying multipath configuration

You can verify multipath configuration on new or existing overcloud deployments.

Procedure

1. Create an instance.
2. Attach a non-encrypted volume to the instance.
3. Get the name of the Compute node that contains the instance:

```
$ nova show <instance> | grep OS-EXT-SRV-ATTR:host
```

Replace **<instance>** with the name of the instance that you created.

4. Retrieve the virsh name of the instance:

```
$ nova show <instance> | grep instance_name
```

5. Get the IP address of the Compute node:

```
$ . stackrc
$ metalsmith list | grep <compute_name>
```

Replace **<compute_name>** with the name from the output of the **nova show <instance>** command to display two rows, from a table of six columns.

Find the row in which **<compute_name>** is in the fourth column. The IP address of **<compute_name>** is in the last column of this row.

In the following example, the IP address of compute-0 is 192.168.24.15 because compute-0 is in the fourth column of the second row:

```
$ . stackrc
```

```
$ metalsmith list | grep compute-0
| 3b1bf72e-c425-494c-9717-d0b89bb66580 | compute-0 | 95b21d3e-36be-470d-ba5c-
70d5dcd6d0b3 | compute-1 | ACTIVE | ctlplane=192.168.24.49 |
| 72a24883-25f9-435c-bf71-a20e66be172d | compute-1 | a59f79f7-006e-4f38-a9ad-
8164da47d58e | compute-0 | ACTIVE | ctlplane=192.168.24.15 |
```

- SSH into the Compute node that runs the instance:

```
$ ssh tripleo-admin@<compute_node_ip>
```

Replace **<compute_node_ip>** with the IP address of the Compute node.

- Log in to the container that runs virsh:

```
$ podman exec -it nova_libvirt /bin/bash
```

- Enter the following command on a Compute node instance to verify that it is using multipath in the cinder volume host location:

```
virsh domblklist <virsh_instance_name> | grep /dev/dm
```

Replace **<virsh_instance_name>** with the output of the **nova show <instance> | grep instance_name** command.

If the instance shows a value other than **/dev/dm-**, the connection is non-multipath and you must refresh the connection info with the **nova shelve** and **nova unshelve** commands:

```
$ nova shelve <instance>
$ nova unshelve <instance>
```



NOTE

If you have more than one type of back end, you must verify the instances and volumes on all back ends, because connection info that each back end returns might vary.

CHAPTER 3. PERFORMING BASIC OPERATIONS WITH THE BLOCK STORAGE SERVICE (CINDER)

Create and configure Block Storage volumes as the primary form of persistent storage for Compute instances in your overcloud. Create volumes, attach your volumes to instances, edit and resize your volumes, and modify volume ownership.

3.1. CREATING BLOCK STORAGE VOLUMES

Create volumes to provide persistent storage for instances that you launch with the Compute service (nova) in the overcloud.

To create an encrypted volume, you must first have a volume type configured specifically for volume encryption. In addition, you must configure both Compute and Block Storage services to use the same static key. For information about how to set up the requirements for volume encryption, see [Block Storage service \(cinder\) volume encryption](#).



IMPORTANT

The default maximum number of volumes you can create for a project is 10.


Prerequisites

- Access to the Red Hat OpenStack Platform (RHOSP) Dashboard (horizon). For more information, see [Introduction to the OpenStack Dashboard](#).

Procedure

1. Log into the dashboard.
2. Select **Project > Compute > Volumes**
3. Click **Create Volume**, and edit the following fields:

Field	Description
Volume name	Name of the volume.
Description	Optional, short description of the volume.

Field	Description
Type	<p>Optional volume type. For more information, see Group volume configuration with volume types.</p> <p>If you create a volume and do not specify a volume type, then Block Storage uses the default volume type. For more information on defining default volume types, see Defining a project-specific default volume type.</p> <p>If you do not specify a back end, the Block Storage scheduler will try to select a suitable back end for you. For more information, see Volume allocation on multiple back ends.</p> <p> NOTE</p> <p>If there is no suitable back end then the volume will not be created.</p> <p>You can also change the volume type after the volume has been created. For more information, see Block Storage volume retyping.</p>
Size (GB)	<p>Volume size (in gigabytes).</p> <p>If you want to create an encrypted volume from an unencrypted image, you must ensure that the volume size is larger than the image size so that the encryption data does not truncate the volume data.</p>
Availability Zone	<p>Availability zones (logical server groups), along with host aggregates, are a common method for segregating resources within OpenStack. Availability zones are defined during installation. For more information about availability zones and host aggregates, see Creating and managing host aggregates in the <i>Configuring the Compute Service for Instance Creation</i> guide.</p>

4. Specify a **Volume Source**:

Source	Description
No source, empty volume	The volume is empty and does not contain a file system or partition table.

Source	Description
Snapshot	Use an existing snapshot as a volume source. If you select this option, a new Use snapshot as a source list opens; you can then choose a snapshot from the list. If you want to create a new volume from a snapshot of an encrypted volume, you must ensure that the new volume is at least 1GB larger than the old volume. For more information about volume snapshots, see Creating new volumes from snapshots
Image	Use an existing image as a volume source. If you select this option, a new Use snapshot as a source list opens; you can then choose an image from the list.
Volume	Use an existing volume as a volume source. If you select this option, a new Use snapshot as a source list opens; you can then choose a volume from the list.

- Click **Create Volume**. After the volume is created, its name appears in the **Volumes** table.

3.2. EDITING A VOLUME NAME OR DESCRIPTION

You can change the names and descriptions of your volumes in the dashboard.

Prerequisites

- Access to the Red Hat OpenStack Platform (RHOSP) Dashboard (horizon). For more information, see [Introduction to the OpenStack Dashboard](#).

Procedure

- Log into the dashboard.
- Select **Project > Compute > Volumes**
- Select the volume's **Edit Volume** button.
- Edit the volume name or description as required.
- Click **Edit Volume** to save your changes.

3.3. RESIZING (EXTENDING) A BLOCK STORAGE SERVICE VOLUME

Resize volumes to increase the storage capacity of the volumes.

**NOTE**

The ability to resize a volume in use is supported but is driver dependent. RBD is supported. You cannot extend in-use multi-attach volumes. For more information about support for this feature, contact Red Hat Support.

Procedure

1. Source your credentials file.
2. List the volumes to retrieve the ID of the volume you want to extend:

```
$ cinder list
```

3. Increase the size of the volume:

```
$ cinder extend <volume_id> <size>
```

- Replace **<volume_id>** with the ID of the volume you want to extend.
- Replace **<size>** with the required size of this volume, in gigabytes.

**NOTE**

Ensure that the specified size is greater than the existing size of this volume.

For example:

```
$ cinder extend 573e024d-5235-49ce-8332-be1576d323f8 10
```

3.4. DELETING A BLOCK STORAGE SERVICE VOLUME

You can delete volumes that you no longer require.

**NOTE**

You cannot delete a volume if it has existing snapshots. For more information about deleting snapshots, see [Deleting volume snapshots](#).

Prerequisites


- Access to the Red Hat OpenStack Platform (RHOSP) Dashboard (horizon). For more information, see [Introduction to the OpenStack Dashboard](#).

Procedure

1. Log into the dashboard.
2. Select **Project > Compute > Volumes**
3. In the **Volumes** table, select the volume to delete.
4. Click **Delete Volumes**.

3.5. VOLUME ALLOCATION ON MULTIPLE BACK ENDS

When you create a volume, you can select the volume type for the required back end from the Type list. For more information, see [Creating Block Storage volumes](#).



NOTE

If the Block Storage service (cinder) is configured to use multiple back ends, then a volume type must be created for each back end.

If you do not specify a back end when creating the volume, the Block Storage scheduler will try to select a suitable back end for you.

The scheduler uses filters, for the following default associated settings of the volume, to select suitable back ends:

AvailabilityZoneFilter

Filters out all back ends that do not meet the availability zone requirements of the requested volume.

CapacityFilter

Selects only back ends with enough space to accommodate the volume.

CapabilitiesFilter

Selects only back ends that can support any specified settings in the volume.

InstanceLocality

Configures clusters to use volumes local to the same node.

If there is more than one suitable back end, then the scheduler uses a weighting method to pick the best back end. By default, the CapacityWeigher method is used, so that the filtered back end with the most available free space is selected.



NOTE

If there is no suitable back end then the volume will not be created.

Additional resources

- [Creating and configuring a volume type](#)
- [Block Storage volume retyping](#)
- [Configuring the default Block Storage scheduler filters](#)

3.6. ATTACHING A VOLUME TO AN INSTANCE

When you close an instance all the data is lost. You can attach a volume for persistent storage. You can attach a volume to only one instance at a time, unless it has a multi-attach volume type. For more information about creating multi-attach volumes, see [Volumes that can be attached to multiple instances](#).

Prerequisites

- Access to the Red Hat OpenStack Platform (RHOSP) Dashboard (horizon). For more information, see [Introduction to the OpenStack Dashboard](#).

Procedure

1. Log into the dashboard.
2. Select **Project > Compute > Volumes**
3. Select the **Edit Attachments** action. If the volume is not attached to an instance, the **Attach To Instance** drop-down list is visible.
4. From the **Attach To Instance** list, select the instance to which you want to attach the volume.
5. Click **Attach Volume**.

3.7. DETACHING A VOLUME FROM AN INSTANCE

You must detach a volume from an instance when you want to attach this volume to another instance, unless it has a multi-attach volume type. You must also detach a volume to change the access permissions to the volume or to delete the volume.

Prerequisites

- Access to the Red Hat OpenStack Platform (RHOSP) Dashboard (horizon). For more information, see [Introduction to the OpenStack Dashboard](#).

Procedure

1. Log into the dashboard.
2. Select **Project > Compute > Volumes**
3. Select the volume's **Manage Attachments** action. If the volume is attached to an instance, the instance's name is displayed in the **Attachments** table.
4. Click **Detach Volume** in this and the next dialog screen.

Next steps

- [Attaching a volume to an instance](#)

3.8. CONFIGURING THE ACCESS RIGHTS TO A VOLUME

The default state of a volume is read-write to allow data to be written to and read from it. You can mark a volume as read-only to protect its data from being accidentally overwritten or deleted.



NOTE

After changing a volume to be read-only you can change it back to read-write again.

Prerequisites

- If the volume is already attached to an instance, then detach this volume. For more information, see [Detaching a volume from an instance](#).

Procedure

1. Source your credentials file.
2. List the volumes to retrieve the ID of the volume you want to configure:

```
$ cinder list
```

3. Set the required access rights for this volume:

- To set the access rights of a volume to read-only:

```
$ cinder readonly-mode-update <volume_id> true
```

- Replace **<volume_id>** with the ID of the required volume.

- To set the access rights of a volume to read-write:

```
$ cinder readonly-mode-update <volume_id> false
```

4. If you detached this volume from an instance to change the access rights, then re-attach the volume. For more information, see [Attaching a volume to an instance](#).

3.9. CHANGING A VOLUME OWNER WITH THE DASHBOARD

To change a volume's owner, you will have to perform a volume transfer. A volume transfer is initiated by the volume's owner, and the volume's change in ownership is complete after the transfer is accepted by the volume's new owner.

Prerequisites

- Access to the Red Hat OpenStack Platform (RHOSP) Dashboard (horizon). For more information, see [Introduction to the OpenStack Dashboard](#).

Procedure

1. Log into the dashboard as the volume owner.
2. Select **Projects > Volumes**.
3. In the **Actions** column of the volume to transfer, select **Create Transfer**.
4. In the **Create Transfer** dialog box, enter a name for the transfer and click **Create Volume Transfer**.

The volume transfer is created, and in the **Volume Transfer** screen you can capture the **transfer ID** and the **authorization key** to send to the recipient project.

Click the **Download transfer credentials** button to download a **.txt** file containing the **transfer name**, **transfer ID**, and **authorization key**.



NOTE

The authorization key is available only in the **Volume Transfer** screen. If you lose the authorization key, you must cancel the transfer and create another transfer to generate a new authorization key.

5. Close the **Volume Transfer** screen to return to the volume list.
The volume status changes to **awaiting-transfer** until the recipient project accepts the transfer

Accept a volume transfer from the dashboard

1. Log into the dashboard as the recipient project owner.
2. Select **Projects > Volumes**.
3. Click **Accept Transfer**.
4. In the **Accept Volume Transfer** dialog box, enter the **transfer ID** and the **authorization key** that you received from the volume owner and click **Accept Volume Transfer**.
The volume now appears in the volume list for the active project.

3.10. CHANGING A VOLUME OWNER WITH THE CLI

To change a volume's owner, you will have to perform a volume transfer. A volume transfer is initiated by the volume's owner, and the volume's change in ownership is complete after the transfer is accepted by the volume's new owner.

Procedure

1. Log in as the volume's current owner.
2. List the available volumes:

```
$ cinder list
```

3. Initiate the volume transfer:

```
$ cinder transfer-create <volume>
```

Replace **<volume>** with the name or ID of the volume you wish to transfer. For example:

```
+-----+-----+
| Property |      Value      |
+-----+-----+
| auth_key | f03bf51ce7ead189 |
| created_at | 2014-12-08T03:46:31.884066 |
| id | 3f5dc551-c675-4205-a13a-d30f88527490 |
| name |      None      |
| volume_id | bcf7d015-4843-464c-880d-7376851ca728 |
+-----+-----+
```

The **cinder transfer-create** command clears the ownership of the volume and creates an **id** and **auth_key** for the transfer. These values can be given to, and used by, another user to accept the transfer and become the new owner of the volume.

4. The new user can now claim ownership of the volume. To do so, the user should first log in from the command line and run:

```
$ cinder transfer-accept <transfer_id> <transfer_key>
```

- Replace **<transfer_id>** with the **id** value returned by the **cinder transfer-create** command.
- Replace **<transfer_key>** with the **auth_key** value returned by the **cinder transfer-create** command.
For example:

```
$ cinder transfer-accept 3f5dc551-c675-4205-a13a-d30f88527490 f03bf51ce7ead189
```



NOTE

You can view all available volume transfers using:

```
$ cinder transfer-list
```

CHAPTER 4. PERFORMING ADVANCED OPERATIONS WITH THE BLOCK STORAGE SERVICE (CINDER)

Block Storage volumes form persistent storage for Compute instances in your overcloud. Configure advanced features of your volumes, for example, using volume snapshots, creating multi-attach volumes, retying volumes, and migrating volumes.

4.1. CREATING VOLUME SNAPSHOTS

You can preserve the state of a volume at a specific point in time by creating a volume snapshot. You can then use the snapshot to clone new volumes.



NOTE

Volume backups are different from snapshots. Backups preserve the data contained in the volume, whereas snapshots preserve the state of a volume at a specific point in time. You cannot delete a volume if it has existing snapshots. Volume backups prevent data loss, whereas snapshots facilitate cloning.

For this reason, snapshot back ends are typically colocated with volume back ends to minimize latency during cloning. By contrast, a backup repository is usually located in a different location, for example, on a different node, physical storage, or even geographical location in a typical enterprise deployment. This is to protect the backup repository from any damage that might occur to the volume back end.

For more information about volume backups, see the [Block Storage Backup Guide](#).

Prerequisites

- A volume that you want to snapshot. For more information about creating volumes, see [Creating Block Storage volumes](#).
- Access to the Red Hat OpenStack Platform (RHOSP) Dashboard (horizon). For more information, see [Introduction to the OpenStack Dashboard](#).

Procedure

1. Log into the dashboard.
2. Select **Project > Compute > Volumes**
3. Select the **Create Snapshot** action for the target volume.
4. Provide a **Snapshot Name** for the snapshot and click **Create a Volume Snapshot** The **Volume Snapshots** tab displays all snapshots.



NOTE

For RADOS block device (RBD) volumes for the Block Storage service (cinder) that are created from snapshots, you can use the **CinderRbdFlattenVolumeFromSnapshot** heat parameter to flatten and remove the dependency on the snapshot. When you set **CinderRbdFlattenVolumeFromSnapshot** to **true**, the Block Storage service flattens RBD volumes and removes the dependency on the snapshot and also flattens all future snapshots. The default value is **false**, which is also the default value for the cinder RBD driver.

Be aware that flattening a snapshot removes any potential block sharing with the parent and results in larger snapshot sizes on the back end and increases the time for snapshot creation.

Verification

- Verify that the new snapshot is present in the **Volume Snapshots** tab, or use the CLI to list volume snapshots and verify that the snapshot is created:

```
$ openstack volume snapshot list
```

4.2. CREATING NEW VOLUMES FROM SNAPSHOTS


You can create new volumes as clones of volume snapshots. These snapshots preserve the state of a volume at a specific point in time.

Prerequisites

- A volume snapshot that you want to clone and create a new volume from. For more information about creating volume snapshots, see [Creating volume snapshots](#)
- Access to the Red Hat OpenStack Platform (RHOSP) Dashboard (horizon). For more information, see [Introduction to the OpenStack Dashboard](#).

Procedure

1. Log into the dashboard.
2. Select **Project > Compute > Volumes**
3. In the **Volume Snapshots** table, select the **Create Volume** action for the snapshot that you want to create a new volume from. For more information about volume creation, see [Creating Block Storage volumes](#).



IMPORTANT

If you want to create a new volume from a snapshot of an encrypted volume, ensure that the new volume is at least 1GB larger than the old volume.

Verification

- Verify that the new volume is present in the **Volumes** tab, or use the CLI to list volumes and verify that the new volume is created:

```
$ openstack volume list
```

4.3. DELETING VOLUME SNAPSHOTS

Red Hat OpenStack Platform (RHOSP) 17.0 uses the RBD CloneV2 API, which means that you can delete volume snapshots even if they have dependencies. If your RHOSP deployment uses a director-deployed Ceph back end, the Ceph cluster is configured correctly by director.

If you use an external Ceph back end, you must configure the minimum client on the Ceph cluster. For more information about configuring an external Ceph cluster, see [Configuring the existing Red Hat Ceph Storage cluster](#) in *Integrating an Overcloud with an Existing Red Hat Ceph Storage Cluster*.

Prerequisites

- A volume snapshot that you want to delete.
- Access to the Red Hat OpenStack Platform (RHOSP) Dashboard (horizon). For more information, see [Introduction to the OpenStack Dashboard](#).

Procedure

1. Log into the dashboard.
2. Select **Project > Compute > Volumes**
3. In the **Volume Snapshots** table, select the **Delete Volume Snapshot** action for the snapshot that you want to delete.

If your OpenStack deployment uses a Red Hat Ceph back end, for more information about snapshot security and troubleshooting, see [Protected and unprotected snapshots in a Red Hat Ceph Storage back end](#).

Verification

- Verify that the snapshot is no longer present in the **Volume Snapshots** tab, or use the CLI to list volume snapshots and verify that the snapshot is deleted:

```
$ openstack volume snapshot list
```

4.4. RESTORING A VOLUME FROM A SNAPSHOT

You can recover the most recent snapshot of a volume. This means that you can perform an in-place revert of the volume data to its most recent snapshot.



WARNING

The ability to recover the most recent snapshot of a volume is supported but is driver dependent. The correct implementation of this feature is driver assisted. For more information about support for this feature, contact your driver vendor.

Limitations

- There might be limitations to using the revert-to-snapshot feature with multi-attach volumes. Check whether such limitations apply before you use this feature.
- You cannot revert a volume that you resize (extend) after you take a snapshot.
- You cannot use the revert-to-snapshot feature on an attached or in-use volume.

Prerequisites

- Block Storage (cinder) REST API microversion 3.40 or later.
- You must have created at least one snapshot for the volume.

Procedure

1. Source your credentials file.
2. Detach your volume:

```
$ nova volume-detach <instance_id> <vol_id>
```

Replace **<instance_id>** and **<vol_id>** with the IDs of the instance and volume that you want to revert.

3. Locate the ID or name of the snapshot that you want to revert. You can only revert the latest snapshot.

```
$ cinder snapshot-list
```

4. Revert the snapshot:

```
$ cinder --os-volume-api-version=3.40 revert-to-snapshot <snapshot_id or snapshot_name>
```

Replace **<snapshot_id or snapshot_name>** with the ID or the name of the snapshot.

5. Optional: You can use the **cinder snapshot-list** command to check that the volume you are reverting is in a reverting state.

```
$ cinder snapshot-list
```

6. Reattach the volume:

```
$ nova volume-attach <instance_id> <vol_id>
```

Replace **<instance_id>** and **<vol_id>** with the IDs of the instance and volume that you reverted.

Verification

- To check that the procedure is successful, you can use the **cinder list** command to verify that the volume you reverted is now in the available state.

\$ cinder list



NOTE

If you used Block Storage (cinder) as a bootable root volume, you cannot use the revert-to-snapshot feature on that volume because it is not in the available state. To use this feature, the instance must have been booted with the **delete_on_termination=false** (default) property to preserve the boot volume if the instance is terminated. When you want to revert to a snapshot, you must first delete the initial instance so that the volume is available. You can then revert it and create a new instance from the volume.

4.5. UPLOADING A VOLUME TO THE IMAGE SERVICE (GLANCE)

You can upload an existing volume as an image to the Image service directly.

Prerequisites

- Access to the Red Hat OpenStack Platform (RHOSP) Dashboard (horizon). For more information, see [Introduction to the OpenStack Dashboard](#).

Procedure

1. Log into the dashboard.
2. Select **Project > Compute > Volumes**
3. Select the target volume's **Upload to Image** action.
4. Provide an **Image Name** for the volume and select a **Disk Format** from the list.
5. Click **Upload**.

To view the uploaded image, select **Project > Compute > Images**. The new image appears in the **Images** table. For information on how to use and configure images, see [Creating and Managing Images](#).

4.6. VOLUMES THAT CAN BE ATTACHED TO MULTIPLE INSTANCES

You can create a multi-attach Block Storage volume that can be attached to multiple instances and these instances can simultaneously read and write to it. Multi-attach volumes require a multi-attach volume type.



WARNING

You must use a multi-attach or cluster-aware file system to manage write operations from multiple instances. Failure to do so causes data corruption.

Limitations of multi-attach volumes

- The Block Storage (cinder) back end must support multi-attach volumes. For information about supported back ends, contact Red Hat Support.
- Your Block Storage (cinder) driver must support multi-attach volumes. The Ceph RBD driver is supported. Contact Red Hat support to verify that multi-attach is supported for your vendor plugin. For more information about the certification for your vendor plugin, see the following locations:
 - <https://access.redhat.com/articles/1535373#Cinder>
 - <https://access.redhat.com/ecosystem/search/#/category/Software?sort=sortTitle%20asc&softwareCategories=Storage&ecosystem=Red%20Hat%20OpenStack%20Platform>
- Read-only multi-attach volumes are not supported.
- Live migration of multi-attach volumes is not available.
- Encryption of multi-attach volumes is not supported.
- Multi-attach volumes are not supported by the Bare Metal Provisioning service (ironic) virt driver. Multi-attach volumes are supported only by the libvirt virt driver.
- You cannot retype an attached volume from a multi-attach type to a non-multi-attach type, and you cannot retype a non-multi-attach type to a multi-attach type.
- You cannot use multi-attach volumes that have multiple read write attachments as the source or destination volume during an attached volume migration.
- You cannot attach multi-attach volumes to shelved offloaded instances.

4.6.1. Creating a multi-attach volume type

To attach a volume to multiple instances, set the **multiattach** flag to **<is> True** in the volume extra specs. When you create a multi-attach volume type, the volume inherits the flag and becomes a multi-attach volume.

Prerequisites

- You must be a project administrator to create a volume type.

Procedure

1. Source the overcloud credentials file:

```
$ source ~/<credentials_file>
```

- Replace **<credentials_file>** with the name of your credentials file, for example, **overcloudrc**.

2. Create a new volume type for multi-attach volumes:

```
$ cinder type-create multiattach
```

3. Enable the **multiattach** property for this multi-attach volume type:

```
$ cinder type-key multiattach set multiattach="<is> True"
```

- 4. Run the following command to specify the back end:

```
$ cinder type-key multiattach set volume_backend_name=<backend_name>
```

4.6.2. Multi-attach volume retyping

You can retype a volume to be multi-attach capable or retype a multi-attach capable volume to make it incapable of attaching to multiple instances. However, you can retype a volume only when it is not in use and its status is **available**.

When you attach a multi-attach volume, some hypervisors require special considerations, such as when you disable caching. Currently, there is no way to safely update an attached volume while keeping it attached the entire time. Retyping fails if you attempt to retype a volume that is attached to multiple instances.

4.6.3. Creating a multi-attach volume

You can create a Block Storage volume that can be attached to multiple instances and these instances can simultaneously read and write to it.



NOTE

This procedure creates a volume on any back end that supports **multiattach**. Therefore, if there are two back ends that support **multiattach**, the scheduler decides which back end to use. For more information, see [Volume allocation on multiple back ends](#).

Prerequisites

- A multi-attach volume type is available in your project.

Procedure

1. Source your credentials file.
2. Run the following command to create a multi-attach volume:

```
$ cinder create <volume_size> --name <volume_name> --volume-type multiattach
```

3. Run the following command to verify that a volume is multi-attach capable. If the volume is multi-attach capable, the **multiattach** field equals **True**.

```
$ cinder show <vol_id> | grep multiattach
```

```
| multiattach | True |
```

Next steps

- [Attaching a volume to an instance](#)

4.7. MOVING VOLUMES BETWEEN BACK ENDS

There are many reasons to move volumes from one storage back end to another, such as:

- To retire storage systems that are no longer supported.
- To change the storage class or tier of a volume.
- To change the availability zone of a volume.

With the Block Storage service (cinder), you can move volumes between back ends in the following ways:

- **Retype:** Only volume owners and administrators can retype volumes. The retype operation is the most common way to move volumes between back ends. For more information, see [Block Storage volume retyping](#).
- **Migrate:** Only administrators can migrate volumes. Volume migration is reserved for specific use cases, because it is restrictive and requires a clear understanding about how deployments work. For more information, see [Migrating a volume between back ends with the Dashboard](#) or [Migrating a volume between back ends with the CLI](#).

Restrictions

Red Hat supports moving volumes between back ends within and across availability zones (AZs), but with the following restrictions:

- Volumes must have either available or in-use status to move.
- Support for in-use volumes is driver dependent.
- Volumes cannot have snapshots.
- Volumes cannot belong to a group or consistency group.

4.7.1. Moving available volumes

You can move available volumes between all back ends, but performance depends on the back ends that you use. Many back ends support assisted migration. For more information about back-end support for assisted migration, contact the vendor.

Assisted migration works with both volume retype and volume migration. With assisted migration, the back end optimizes the movement of the data from the source back end to the destination back end, but both back ends must be from the same vendor.



NOTE

Red Hat supports back-end assisted migrations only with multi-pool back ends or when you use the cinder migrate operation for single-pool back ends, such as RBD.

When assisted migrations between back ends are not possible, the Block Storage service performs a generic volume migration.

Generic volume migration requires volumes on both back ends to be connected before the Block Storage (cinder) service moves data from the source volume to the Controller node and from the Controller node to the destination volume. The Block Storage service seamlessly performs the process regardless of the type of storage from the source and destination back ends.

**IMPORTANT**

Ensure that you have adequate bandwidth before you perform a generic volume migration. The duration of a generic volume migration is directly proportional to the size of the volume, which makes the operation slower than assisted migration.

4.7.2. Moving in-use volumes

There is no optimized or assisted option for moving in-use volumes. When you move in-use volumes, the Compute service (nova) must use the hypervisor to transfer data from a volume in the source back end to a volume in the destination back end. This requires coordination with the hypervisor that runs the instance where the volume is in use.

The Block Storage service (cinder) and the Compute service work together to perform this operation. The Compute service manages most of the work, because the data is copied from one volume to another through the Compute node.

**IMPORTANT**

Ensure that you have adequate bandwidth before you move in-use volumes. The duration of this operation is directly proportional to the size of the volume, which makes the operation slower than assisted migration.

Restrictions

- In-use multi-attach volumes cannot be moved while they are attached to more than one nova instance.
- Non block devices are not supported, which limits storage protocols on the target back end to be iSCSI, Fibre Channel (FC), and RBD.

4.8. BLOCK STORAGE VOLUME RETYPING

When you retype a volume, you apply a volume type and its settings to an existing volume. For more information about volume types, see [Group volume configuration with volume types](#) .

**NOTE**

Only volume owners and administrators can retype volumes.

You can retype a volume provided that the extra specs of the new volume type can be applied to the existing volume. You can retype a volume to apply predefined settings or storage attributes to an existing volume, such as:

- To move the volume to a different back end.
- To change the storage class or tier of a volume.
- To enable or disable features such as replication.

Volume retyping is the standard way to move volumes from one back end to another. But retyping a volume does not necessarily mean that you must move the volume from one back end to another. However, there are circumstances in which you must move a volume to complete a retype:

- The new volume type has a different **volume_backend_name** defined.
- The **volume_backend_name** of the current volume type is undefined, and the volume is stored in a different back end than the one specified by the **volume_backend_name** of the new volume type.

Moving a volume from one back end to another can require extensive time and resources. Therefore, when a retype requires moving data, the Block Storage service does not move data by default. The operation fails unless it is explicitly allowed by specifying a migration policy as part of the retype request. For more information, see [Retyping a volume from the Dashboard](#) or [Retyping a volume from the CLI](#).

Restrictions

- You cannot retype all volumes. For more information about moving volumes between back ends, see [Moving volumes between back ends](#).
- Unencrypted volumes cannot be retyped to encrypted volume types, but encrypted volumes can be retyped to unencrypted ones.
- Retyping an unencrypted volume to an encrypted volume of the same size is not supported, because encrypted volumes require additional space to store encryption data. For more information about encrypting unencrypted volumes, see [Encrypting unencrypted volumes](#).
- Users with no administrative privileges can only retype volumes that they own.

Additional resources

- [Creating and configuring a volume type](#)

4.8.1. Retyping a volume from the Dashboard

Retype a volume to apply a volume type and its settings to an existing volume.



IMPORTANT

Retyping an unencrypted volume to an encrypted volume of the same size is not supported, because encrypted volumes require additional space to store encryption data. For more information about encrypting unencrypted volumes, see [Encrypting unencrypted volumes](#).

Prerequisites

- Only volume owners and administrators can retype volumes.
- Access to the Red Hat OpenStack Platform (RHOSP) Dashboard (horizon). For more information, see [Introduction to the OpenStack Dashboard](#).

Procedure

1. Log into the dashboard as an admin user or volume owner.
2. Select **Project > Compute > Volumes**
3. In the **Actions** column of the volume you want to migrate, select **Change Volume Type**.

4. In the **Change Volume Type** dialog, select the target volume type and define the new back end from the **Type** list.
5. If you are migrating the volume to another back end, select **On Demand** from the **Migration Policy** list. For more information, see [Moving volumes between back ends](#).
6. Click **Change Volume Type** to start the migration.

4.8.2. Retyping a volume from the CLI

Retype a volume to apply a volume type and its settings to an existing volume.



IMPORTANT

Retyping an unencrypted volume to an encrypted volume of the same size is not supported, because encrypted volumes require additional space to store encryption data. For more information about encrypting unencrypted volumes, see [Encrypting unencrypted volumes](#).

Prerequisites

- Only volume owners and administrators can retype volumes.

Procedure

1. Source your credentials file.
2. Enter the following command to retype a volume:

```
$ cinder retype <volume name or id> <new volume type name>
```

3. If the retype operation requires moving the volume from one back end to another, the Block Storage service requires a specific flag:

```
$ cinder retype --migration-policy on-demand <volume name or id> <new volume type name>
```

4.9. MIGRATING A VOLUME BETWEEN BACK ENDS WITH THE DASHBOARD

With the Block Storage service (cinder) you can migrate volumes between back ends within and across availability zones (AZs). This is the least common way to move volumes from one back end to another.

In highly customized deployments or in situations in which you must retire a storage system, an administrator can migrate volumes. In both use cases, multiple storage systems share the same **volume_backend_name**, or it is undefined.

Restrictions

- The volume cannot be replicated.
- The destination back end must be different from the current back end of the volume.

- The existing volume type must be valid for the new back end, which means the following must be true:
 - Volume type must not have the **backend_volume_name** defined in its extra specs, or both Block Storage back ends must be configured with the same **backend_volume_name**.
 - Both back ends must support the same features configured in the volume type, such as support for thin provisioning, support for thick provisioning, or other feature configurations.



NOTE

Moving volumes from one back end to another might require extensive time and resources. For more information, see [Moving volumes between back ends](#).

Prerequisites

- You must be a project administrator to migrate volumes.
- Access to the Red Hat OpenStack Platform (RHOSP) Dashboard (horizon). For more information, see [Introduction to the OpenStack Dashboard](#).

Procedure

1. Log into the dashboard as an admin user.
2. Select **Admin > Volumes**.
3. In the **Actions** column of the volume you want to migrate, select **Migrate Volume**.
4. In the **Migrate Volume** dialog, select the target host from the **Destination Host** drop-down list.



NOTE

To bypass any driver optimizations for the host migration, select the **Force Host Copy** check box.

5. Click **Migrate** to start the migration.

4.10. MIGRATING A VOLUME BETWEEN BACK ENDS WITH THE CLI

With the Block Storage service (cinder) you can migrate volumes between back ends within and across availability zones (AZs). This is the least common way to move volumes from one back end to another.

In highly customized deployments or in situations in which you must retire a storage system, an administrator can migrate volumes. In both use cases, multiple storage systems share the same **volume_backend_name**, or it is undefined.

Restrictions

- The volume cannot be replicated.
- The destination back end must be different from the current back end of the volume.
- The existing volume type must be valid for the new back end, which means the following must be true:

- Volume type must not have the **backend_volume_name** defined in its extra specs, or both Block Storage back ends must be configured with the same **backend_volume_name**.
- Both back ends must support the same features configured in the volume type, such as support for thin provisioning, support for thick provisioning, or other feature configurations.



NOTE

Moving volumes from one back end to another might require extensive time and resources. For more information, see [Moving volumes between back ends](#).

Prerequisites

- You must be a project administrator to migrate volumes.

Procedure

1. Source the overcloud credentials file:

```
$ source ~/<credentials_file>
```

- Replace **<credentials_file>** with the name of your credentials file, for example, **overcloudrc**.

2. Enter the following command to retrieve the name of the destination back end:

```
$ cinder get-pools --detail

Property          | Value
...
| name             | localdomain@lvmdriver-1#lvmdriver-1
| pool_name        | lvmdriver-1
...
| volume_backend_name | lvmdriver-1
...
Property          | Value
...
| name             | localdomain@lvmdriver-2#lvmdriver-1
| pool_name        | lvmdriver-1
...
| volume_backend_name | lvmdriver-1
...
```

The destination back-end names use this syntax: **host@volume_backend_name#pool**.

In the example output, there are two LVM back ends exposed in the Block Storage service: **localdomain@lvmdriver-1#lvmdriver-1** and **localdomain@lvmdriver-2#lvmdriver-1**. Notice that both back ends share the same **volume_backend_name**, **lvmdriver-1**.

3. Enter the following command to migrate a volume from one back end to another:

```
$ cinder migrate <volume id or name> <new host>
```

4.11. ENCRYPTING UNENCRYPTED VOLUMES

You can encrypt an unencrypted volume.

If the **cinder-backup** service is available, then back up the unencrypted volume and then restore it to a new encrypted volume.

If the **cinder-backup** service is unavailable, then create a glance image from the unencrypted volume and create a new encrypted volume from this image.

Prerequisites

- You must be a project administrator to create encrypted volumes.
- An unencrypted volume that you want to encrypt.

Procedure

The **cinder-backup** service is available:

1. Source the overcloud credentials file:

```
$ source ~/<credentials_file>
```

- Replace **<credentials_file>** with the name of your credentials file, for example, **overcloudrc**.

2. Back up the current unencrypted volume:

```
cinder backup-create <unencrypted_volume>
```

- Replace **<unencrypted_volume>** with the name or ID of the unencrypted volume.

3. Create a new encrypted volume:

```
cinder create <encrypted_volume_size> --volume-type <encrypted_volume_type>
```

- Replace **<encrypted_volume_size>** with the size of the new volume in GB. This value must be larger than the size of the unencrypted volume by 1GB to accommodate the encryption metadata.
- Replace **<encrypted_volume_type>** with the encryption type that you require.

4. Restore the backup of the unencrypted volume to the new encrypted volume:

```
cinder backup-restore <backup> --volume <encrypted_volume>
```

- Replace **<backup>** with the name or ID of the unencrypted volume backup.
- Replace **<encrypted_volume>** with the ID of the new encrypted volume.

The **cinder-backup** service is unavailable:

1. Source the overcloud credentials file:

```
$ source ~/<credentials_file>
```

- Replace **<credentials_file>** with the name of your credentials file, for example, **overcloudrc**.

2. Create a glance image of the unencrypted volume:

```
cinder upload-to-image <unencrypted_volume> <new_image>
```

- Replace **<unencrypted_volume>** with the name or ID of the unencrypted volume.
- Replace **<new_image>** with a name for the new image.

3. Create a new volume from the image that is 1GB larger than the image:

```
cinder volume create --size <size> --volume-type luks --image <new_image>  
<encrypted_volume_name>
```

- Replace **<size>** with the size of the new volume. This value must be 1GB larger than the size of the old unencrypted volume.
- Replace **<new_image>** with the name of the image that you created from the unencrypted volume.
- Replace **<encrypted_volume_name>** with a name for the new encrypted volume.

4.12. PROTECTED AND UNPROTECTED SNAPSHOTS IN A RED HAT CEPH STORAGE BACK END

When you use Red Hat Ceph Storage (RHCS) as a back end for your Red Hat OpenStack Platform (RHOSP) deployment, you can set a snapshot to **protected** in the back end. Do not delete protected snapshots through the RHOSP dashboard or the **cinder snapshot-delete** command because deletion fails.

When this occurs, set the snapshot to **unprotected** in the RHCS back end first. You can then delete the snapshot through RHOSP as normal.

For more information about protecting snapshots, see [Protecting a block device snapshot](#) and [Unprotecting a block device snapshot](#) in the Red Hat Ceph Storage *Block Device Guide*.

CHAPTER 5. CONFIGURING THE OBJECT STORAGE SERVICE (SWIFT)

The Red Hat OpenStack Platform (RHOSP) Object Storage service (swift) stores its objects, or data, in containers. Containers are similar to directories in a file system although they cannot be nested. Containers provide an easy way for users to store any kind of unstructured data. For example, objects can include photos, text files, or images. Stored objects are not compressed.

5.1. OBJECT STORAGE RINGS

The Object Storage service (swift) uses a data structure called the ring to distribute partition space across the cluster. This partition space is core to the data durability engine in the Object Storage service. With rings, the Object Storage service can quickly and easily synchronize each partition across the cluster.

Rings contain information about Object Storage partitions and how partitions are distributed among the different nodes and disks. When any Object Storage component interacts with data, a quick lookup is performed locally in the ring to determine the possible partitions for each object.

The Object Storage service has three rings to store the following types of data:

- Account information
- Containers, to facilitate organizing objects under an account
- Object replicas

5.1.1. Checking cluster health

The Object Storage service (swift) runs many processes in the background to ensure long-term data availability, durability, and persistence. For example:

- Auditors constantly re-read database and object files and compare them by using checksums to make sure there is no silent bit-rot. Any database or object file that no longer matches its checksum is quarantined and becomes unreadable on that node. The replicators then copy one of the other replicas to make the local copy available again.
- Objects and files can disappear when you replace disks or nodes or when objects are quarantined. When this happens, replicators copy missing objects or database files to one of the other nodes.

The Object Storage service includes a tool called **swift-recon** that collects data from all nodes and checks the overall cluster health.

Procedure

1. Log in to one of the Controller nodes.
2. Run the following command:

```
[tripleo-admin@overcloud-controller-2 ~]$ sudo podman exec -it -u swift swift_object_server /usr/bin/swift-recon -arqIT --md5
```

```
-----  
-> Starting reconnaissance on 3 hosts (object)
```

```

=====
[2018-12-14 14:55:47] Checking async pendings
[async_pending] - No hosts returned valid data.
=====
[2018-12-14 14:55:47] Checking on replication
[replication_failure] low: 0, high: 0, avg: 0.0, total: 0, Failed: 0.0%, no_result: 0, reported: 3
[replication_success] low: 0, high: 0, avg: 0.0, total: 0, Failed: 0.0%, no_result: 0, reported: 3
[replication_time] low: 0, high: 0, avg: 0.0, total: 0, Failed: 0.0%, no_result: 0, reported: 3
[replication_attempted] low: 0, high: 0, avg: 0.0, total: 0, Failed: 0.0%, no_result: 0, reported:
3
Oldest completion was 2018-12-14 14:55:39 (7 seconds ago) by 172.16.3.186:6000.
Most recent completion was 2018-12-14 14:55:42 (4 seconds ago) by 172.16.3.174:6000.
=====
[2018-12-14 14:55:47] Checking load averages
[5m_load_avg] low: 1, high: 2, avg: 2.1, total: 6, Failed: 0.0%, no_result: 0, reported: 3
[15m_load_avg] low: 2, high: 2, avg: 2.6, total: 7, Failed: 0.0%, no_result: 0, reported: 3
[1m_load_avg] low: 0, high: 0, avg: 0.8, total: 2, Failed: 0.0%, no_result: 0, reported: 3
=====
[2018-12-14 14:55:47] Checking ring md5sums
3/3 hosts matched, 0 error[s] while checking hosts.
=====
[2018-12-14 14:55:47] Checking swift.conf md5sum
3/3 hosts matched, 0 error[s] while checking hosts.
=====
[2018-12-14 14:55:47] Checking quarantine
[quarantined_objects] low: 0, high: 0, avg: 0.0, total: 0, Failed: 0.0%, no_result: 0, reported: 3
[quarantined_accounts] low: 0, high: 0, avg: 0.0, total: 0, Failed: 0.0%, no_result: 0, reported:
3
[quarantined_containers] low: 0, high: 0, avg: 0.0, total: 0, Failed: 0.0%, no_result: 0,
reported: 3
=====
[2018-12-14 14:55:47] Checking time-sync
3/3 hosts matched, 0 error[s] while checking hosts.
=====

```

3. Optional: Use the **--all** option to return additional output.

This command queries all servers on the ring for the following data:

- Async pendings: If the cluster load is too high and processes cannot update database files fast enough, some updates occur asynchronously. These numbers decrease over time.
- Replication metrics: Review the replication timestamps; full replication passes happen frequently with few errors. An old entry, for example, an entry with a timestamp from six months ago, indicates that replication on the node has not completed in the last six months.
- Ring md5sums: This ensures that all ring files are consistent on all nodes.
- **swift.conf** md5sums: This ensures that all configuration files are consistent on all nodes.
- Quarantined files: There must be no, or very few, quarantined files for all nodes.
- Time-sync: All nodes must be synchronized.

5.1.2. Increasing ring partition power

The ring power determines the partition to which a resource, such as an account, container, or object, is

mapped. The partition is included in the path under which the resource is stored in a back-end file system. Therefore, changing the partition power requires relocating resources to new paths in the back-end file systems.

In a heavily populated cluster, a relocation process is time consuming. To avoid downtime, relocate resources while the cluster is still operating. You must do this without temporarily losing access to data or compromising the performance of processes, such as replication and auditing. For assistance with increasing ring partition power, contact Red Hat support.

5.1.3. Partition power recommendation for the Object Storage service

When using separate Red Hat OpenStack Platform (RHOSP) Object Storage service (swift) nodes, use a higher partition power value.

The Object Storage service distributes data across disks and nodes using modified *hash rings*. There are three rings by default: one for accounts, one for containers, and one for objects. Each ring uses a fixed parameter called *partition power*. This parameter sets the maximum number of partitions that can be created.

The partition power parameter is important and can only be changed for new containers and their objects. As such, it is important to set this value before *initial deployment*.

The default partition power value is **10** for environments that RHOSP director deploys. This is a reasonable value for smaller deployments, especially if you only plan to use disks on the Controller nodes for the Object Storage service.

The following table helps you to select an appropriate partition power if you use three replicas:

Table 5.1. Appropriate partition power values per number of available disks

Partition Power	Maximum number of disks
10	~ 35
11	~ 75
12	~ 150
13	~ 250
14	~ 500



IMPORTANT

Setting an excessively high partition power value (for example, **14** for only 40 disks) negatively impacts replication times.

To set the partition power, use the following resource:

```
parameter_defaults:
  SwiftPartPower: 11
```

TIP

You can also configure an additional object server ring for new containers. This is useful if you want to add more disks to an Object Storage service deployment that initially uses a low partition power.

Additional resources

- [The Rings](#) in swift upstream documentation
- [Modifying the overcloud environment](#) in *Director Installation and Usage*

5.1.4. Custom rings

As technology advances and demands for storage capacity increase, creating custom rings is a way to update existing Object Storage clusters.

When you add new nodes to a cluster, their characteristics might differ from those of the original nodes. Without custom adjustments, the larger capacity of the new nodes may be underused. Or, if weights change in the rings, data dispersion can become uneven, which reduces safety.

Automation might not keep pace with future technology trends. For example, some older Object Storage clusters in use today originated before SSDs were available.

The ring builder helps manage Object Storage as clusters grow and technologies evolve. For assistance with creating custom rings, contact Red Hat support.

5.2. CUSTOMIZING THE OBJECT STORAGE SERVICE

Depending on the requirements of your Red Hat OpenStack Platform (RHOSP) environment, you might want to customize some of the default settings of the Object Storage service (swift) to optimize your deployment performance.

5.2.1. Configuring fast-post

By default, the Object Storage service (swift) copies an object whole whenever any part of its metadata changes. You can prevent this by using the *fast-post* feature. The fast-post feature saves time when you change the content types of multiple large objects.

To enable the fast-post feature, disable the **object_post_as_copy** option on the Object Storage proxy service.

Procedure

1. Edit **swift_params.yaml**:

```
cat > swift_params.yaml << EOF
parameter_defaults:
  ExtraConfig:
    swift::proxy::copy::object_post_as_copy: False
EOF
```

2. Include the parameter file when you deploy or update the overcloud:

```
openstack overcloud deploy [... previous args ...] -e swift_params.yaml
```

5.2.2. Enabling at-rest encryption

By default, objects uploaded to the Object Storage service (swift) are unencrypted. Because of this, it is possible to access objects directly from the file system. This can present a security risk if disks are not properly erased before they are discarded. Object Storage objects. For more information, see [Encrypting Object Storage \(swift\) at-rest objects](#) in *Manage Secrets with OpenStack Key Manager*.

5.2.3. Deploying a standalone Object Storage service cluster

You can use the composable role concept to deploy a standalone Object Storage service (swift) cluster with the bare minimum of additional services, for example, OpenStack Identity service (keystone) or HAProxy.

Procedure

1. Copy the **roles_data.yaml** from **/usr/share/openstack-tripleo-heat-templates**.
2. Edit the new file.
3. Remove unneeded controller roles, for example Aodh*, Ceilometer*, Ceph*, Cinder*, Glance*, Heat*, Ironic*, Manila*, Nova*, Octavia*, Swift*.
4. Locate the ObjectStorage role within **roles_data.yaml**.
5. Copy this role to a new role within that same file and name it **ObjectProxy**.
6. Replace **SwiftStorage** with **SwiftProxy** in this role.

The example **roles_data.yaml** file below shows sample roles.

```
- name: Controller
  description: |
    Controller role that has all the controller services loaded and handles
    Database, Messaging and Network functions.
  CountDefault: 1
  tags:
  - primary
  - controller
  networks:
  - External
  - InternalApi
  - Storage
  - StorageMgmt
  - Tenant
  HostnameFormatDefault: '%stackname%-controller-%index%'
  ServicesDefault:
  - OS::TripleO::Services::AuditD
  - OS::TripleO::Services::CACerts
  - OS::TripleO::Services::CertmongerUser
  - OS::TripleO::Services::Clustercheck
  - OS::TripleO::Services::Docker
  - OS::TripleO::Services::Ec2Api
  - OS::TripleO::Services::Etcd
  - OS::TripleO::Services::HAproxy
  - OS::TripleO::Services::Keepalived
  - OS::TripleO::Services::Kernel
  - OS::TripleO::Services::Keystone
```



```

- OS::TripleO::Services::Memcached
- OS::TripleO::Services::MySQL
- OS::TripleO::Services::MySQLClient
- OS::TripleO::Services::Ntp
- OS::TripleO::Services::Pacemaker
- OS::TripleO::Services::RabbitMQ
- OS::TripleO::Services::Securetty
- OS::TripleO::Services::Snmp
- OS::TripleO::Services::Sshd
- OS::TripleO::Services::Timezone
- OS::TripleO::Services::TripleoFirewall
- OS::TripleO::Services::TripleoPackages
- OS::TripleO::Services::Vpp

- name: ObjectStorage
  CountDefault: 1
  description: |
    Swift Object Storage node role
  networks:
    - InternalApi
    - Storage
    - StorageMgmt
  disable_upgrade_deployment: True
  ServicesDefault:
    - OS::TripleO::Services::AuditD
    - OS::TripleO::Services::CACerts
    - OS::TripleO::Services::CertmongerUser
    - OS::TripleO::Services::Collectd
    - OS::TripleO::Services::Docker
    - OS::TripleO::Services::FluentdClient
    - OS::TripleO::Services::Kernel
    - OS::TripleO::Services::MySQLClient
    - OS::TripleO::Services::Ntp
    - OS::TripleO::Services::Securetty
    - OS::TripleO::Services::SensuClient
    - OS::TripleO::Services::Snmp
    - OS::TripleO::Services::Sshd
    - OS::TripleO::Services::SwiftRingBuilder
    - OS::TripleO::Services::SwiftStorage
    - OS::TripleO::Services::Timezone
    - OS::TripleO::Services::TripleoFirewall
    - OS::TripleO::Services::TripleoPackages

- name: ObjectProxy
  CountDefault: 1
  description: |
    Swift Object proxy node role
  networks:
    - InternalApi
    - Storage
    - StorageMgmt
  disable_upgrade_deployment: True
  ServicesDefault:
    - OS::TripleO::Services::AuditD
    - OS::TripleO::Services::CACerts
    - OS::TripleO::Services::CertmongerUser

```

```

- OS::TripleO::Services::Collectd
- OS::TripleO::Services::Docker
- OS::TripleO::Services::FluentdClient
- OS::TripleO::Services::Kernel
- OS::TripleO::Services::MySQLClient
- OS::TripleO::Services::Ntp
- OS::TripleO::Services::Securetty
- OS::TripleO::Services::SensuClient
- OS::TripleO::Services::Snmp
- OS::TripleO::Services::Sshd
- OS::TripleO::Services::SwiftRingBuilder
- OS::TripleO::Services::SwiftProxy
- OS::TripleO::Services::Timezone
- OS::TripleO::Services::TripleoFirewall
- OS::TripleO::Services::TripleoPackages

```

7. Deploy the overcloud with your regular **openstack deploy** command, including the new roles.

```
$ openstack overcloud deploy --templates -r roles_data.yaml -e [...]
```

5.2.4. Using external SAN disks

By default, when Red Hat OpenStack Platform (RHOSP) director deploys the Object Storage service (swift), Object Storage is configured and optimized to use independent local disks. This configuration ensures that the workload is distributed across all disks, which helps minimize performance impact during node failure or other system issues.

In performance-impacting events, an environment that uses a single SAN can experience decreased performance across all LUNs. The Object Storage service cannot mitigate performance issues in an environment that uses SAN disks. Therefore, use additional local disks for Object Storage to meet performance and disk space requirements. For more information, see [Disk recommendation for the Object Storage service](#).

Using an external SAN for Object Storage requires evaluation on a per-case basis. For more information, contact Red Hat Support.

IMPORTANT

If you choose to use an external SAN for Object Storage, be aware of the following conditions:

- Red Hat does not provide support for issues related to performance that result from using an external SAN for Object Storage.
- Red Hat does not provide support for issues that arise outside of the core Object Storage service offering. For support with high availability and performance, contact your storage vendor.
- Red Hat does not test SAN solutions with the Object Storage service. For more information about compatibility, guidance, and support for third-party products, contact your storage vendor.
- Red Hat recommends that you evaluate and test performance demands with your deployment. To confirm that your SAN deployment is tested, supported, and meets your performance requirements, contact your storage vendor.

Procedure

- This template is an example of how to use two devices (`/dev/mapper/vdb` and `/dev/mapper/vdc`) for Object Storage:

```
parameter_defaults:
  SwiftMountCheck: true
  SwiftUseLocalDir: false
  SwiftRawDisks: {"vdb": {"base_dir": "/dev/mapper/"}, "vdc": {"base_dir": "/dev/mapper/"}}
```

5.2.5. Disk recommendation for the Object Storage service

Use one or more separate, local disks for the Red Hat OpenStack Platform (RHOSP) Object Storage service (swift).

By default, RHOSP director uses the directory `/srv/node/d1` on the system disk for the Object Storage service. On the Controller, this disk is also used by other services, and the disk can become a performance bottleneck.

The following example is an excerpt from a RHOSP Orchestration service (heat) custom environment file. On each Controller node, the Object Storage service uses two separate disks. The entirety of both disks contains an XFS file system:

```
parameter_defaults:
  SwiftRawDisks: {"sdb": {}, "sdc": {}}
```

SwiftRawDisks defines each storage disk on the node. This example defines both **sdb** and **sdc** disks on each Controller node.



IMPORTANT

When configuring multiple disks, ensure that the Bare Metal service (ironic) uses the intended root disk.

Additional resources

- [Defining the root disk for multi-disk clusters](#) in the *Director Installation and Usage* guide.

5.3. ADDING OR REMOVING OBJECT STORAGE NODES

To add new Object Storage (swift) nodes to your cluster, you must increase the node count, update the rings, and synchronize the changes. You can increase the node count by adding nodes to the overcloud or by scaling up bare-metal nodes.

To remove Object Storage nodes from your cluster, you can perform a simple removal or an incremental removal, depending on the quantities of data in the cluster.

5.3.1. Adding nodes to the overcloud

You can add more nodes to your overcloud.



NOTE

A fresh installation of Red Hat OpenStack Platform does not include certain updates, such as security errata and bug fixes. As a result, if you are scaling up a connected environment that uses the Red Hat Customer Portal or Red Hat Satellite Server, RPM updates are not applied to new nodes. To apply the latest updates to the overcloud nodes, you must do one of the following:

- Complete an overcloud update of the nodes after the scale-out operation.
- Use the **virt-customize** tool to modify the packages to the base overcloud image before the scale-out operation. For more information, see the Red Hat Knowledgebase solution [Modifying the Red Hat Linux OpenStack Platform Overcloud Image with virt-customize](#).

Procedure

1. Create a new JSON file called **newnodes.json** that contains details of the new node that you want to register:

```
{
  "nodes":[
    {
      "mac":[
        "dd:dd:dd:dd:dd:dd"
      ],
      "cpu":"4",
      "memory":"6144",
      "disk":"40",
      "arch":"x86_64",
      "pm_type":"ipmi",
      "pm_user":"admin",
      "pm_password":"p@55w0rd!",
      "pm_addr":"192.168.24.207"
    },
    {
      "mac":[
        "ee:ee:ee:ee:ee:ee"
      ],
      "cpu":"4",
      "memory":"6144",
      "disk":"40",
      "arch":"x86_64",
      "pm_type":"ipmi",
      "pm_user":"admin",
      "pm_password":"p@55w0rd!",
      "pm_addr":"192.168.24.208"
    }
  ]
}
```

2. Register the new nodes:

```
$ source ~/stackrc
(undercloud)$ openstack overcloud node import newnodes.json
```

3. Launch the introspection process for each new node:

```
(undercloud)$ openstack overcloud node introspect \
--provide <node_1> [node_2] [node_n]
```

- Use the **--provide** option to reset all the specified nodes to an **available** state after introspection.
- Replace **<node_1>**, **[node_2]**, and all nodes up to **[node_n]** with the UUID of each node that you want to introspect.

4. Configure the image properties for each new node:

```
(undercloud)$ openstack overcloud node configure <node>
```

5.3.2. Scaling up bare-metal nodes

To increase the count of bare-metal nodes in an existing overcloud, increment the node count in the **overcloud-baremetal-deploy.yaml** file and redeploy the overcloud.

Prerequisites

- The new bare-metal nodes are registered, introspected, and available for provisioning and deployment. For more information, see [Registering nodes for the overcloud](#) and [Creating an inventory of the bare-metal node hardware](#).

Procedure

1. Source the **stackrc** undercloud credential file:

```
$ source ~/stackrc
```

2. Open the **overcloud-baremetal-deploy.yaml** node definition file that you use to provision your bare-metal nodes.
3. Increment the **count** parameter for the roles that you want to scale up. For example, the following configuration increases the Object Storage node count to 4:

```
- name: Controller
  count: 3
- name: Compute
  count: 10
- name: ObjectStorage
  count: 4
```

4. Optional: Configure predictive node placement for the new nodes. For example, use the following configuration to provision a new Object Storage node on **node03**:

```
- name: ObjectStorage
  count: 4
  instances:
    - hostname: overcloud-objectstorage-0
      name: node00
    - hostname: overcloud-objectstorage-1
```

```

name: node01
- hostname: overcloud-objectstorage-2
  name: node02
- hostname: overcloud-objectstorage-3
  name: node03

```

- Optional: Define any other attributes that you want to assign to your new nodes. For more information about the properties you can use to configure node attributes in your node definition file, see [Bare-metal node provisioning attributes](#).
- If you use the Object Storage service (swift) and the whole disk overcloud image, **overcloud-hardened-uefi-full**, configure the size of the `/srv` partition based on the size of your disk and your storage requirements for `/var` and `/srv`. For more information, see [Configuring whole disk partitions for the Object Storage service](#).
- Provision the overcloud nodes:

```

(undercloud)$ openstack overcloud node provision \
--stack <stack> \
--output <deployment_file> \
/home/stack/templates/overcloud-baremetal-deploy.yaml

```

- Replace `<stack>` with the name of the stack for which the bare-metal nodes are provisioned. If not specified, the default is **overcloud**.
 - Replace `<deployment_file>` with the name of the heat environment file to generate for inclusion in the deployment command, for example `/home/stack/templates/overcloud-baremetal-deployed.yaml`.
- Monitor the provisioning progress in a separate terminal. When provisioning is successful, the node state changes from **available** to **active**:

```

(undercloud)$ watch openstack baremetal node list

```

- Add the generated **overcloud-baremetal-deployed.yaml** file to the stack with your other environment files and deploy the overcloud:

```

(undercloud)$ openstack overcloud deploy --templates \
-e [your environment files] \
-e /home/stack/templates/overcloud-baremetal-deployed.yaml \
--deployed-server \
--disable-validations \
...

```

5.3.3. Defining dedicated Object Storage nodes

Dedicate additional nodes to the Red Hat OpenStack Platform (RHOSP) Object Storage service to improve performance.

If you are dedicating additional nodes to the Object Storage service, edit the custom `roles_data.yaml` file to remove the Object Storage service entry from the Controller node. Specifically, remove the following line from the `ServicesDefault` list of the `Controller` role:

```

- OS::TripleO::Services::SwiftStorage

```

5.3.4. Updating and rebalancing the Object Storage rings

The Object Storage service (swift) requires the same ring files on all Controller and Object Storage nodes. If a Controller node or Object Storage node is replaced, added or removed, these must be synced after an overcloud update to ensure proper functionality.

Procedure

1. Log in to the undercloud as the **stack** user and create a temporary directory:

```
$ mkdir temp && cd temp/
```

2. Download the overcloud ring files from one of the previously existing nodes (Controller 0 in this example) to the new directory:

```
$ ssh tripleo-admin@overcloud-controller-0.ctlplane 'sudo tar -czvf - /var/lib/config-  
data/puppet-generated/swift_ringbuilder/etc/swift/{*.builder,*.ring.gz,backups/*.builder}' >  
swift-rings.tar.gz
```

3. Extract the rings and change into the ring subdirectory:

```
$ tar xzvf swift-rings.tar.gz && cd var/lib/config-data/puppet-  
generated/swift_ringbuilder/etc/swift/
```

4. Collect the values for the following variables according to your device details:

- **<device_name>**:

```
$ openstack baremetal introspection data save <node_name> | jq ".inventory.disks"
```

- **<node_ip>**:

```
$ metalsmith <node_name> show
```

- **<port>**: The default port is **600x**. If you altered the default, use the applicable port.

- **<builder_file>**: The builder file name from step 3.

- **<weight>** and **<zone>** variables are user-defined.

5. Use **swift-ring-builder** to add and update the new node to the existing rings. Replace the variables according to the device details.



NOTE

You must install the **python3-swift** RPM to use the **swift-ring-builder** command.

```
$ swift-ring-builder etc/swift/<builder_file> add <zone>-<node_ip>:<port>/<device_name>  
<weight>
```

6. Rebalance the ring to ensure that the new device is used:

```
$ swift-ring-builder etc/swift/<builder_file> rebalance
```

-
- 7. Upload the modified ring files to the Controller nodes and ensure that these ring files are used. Use a script, similar to the following example, to distribute ring files:

```
#!/bin/sh
set -xe

ALL="tripleo-admin@overcloud-controller-0.ctlplane tripleo-admin@overcloud-controller-1.ctlplane tripleo-admin@overcloud-controller-2.ctlplane"
```

- Upload the rings to all nodes and restart Object Storage services:

```
for DST in ${ALL}; do
  cat swift-rings.tar.gz | ssh "${DST}" 'sudo tar -C / -xvzf -'
  ssh "${DST}" 'sudo podman restart swift_copy_rings'
  ssh "${DST}" 'sudo systemctl restart tripleo_swift*'
done
```

5.3.5. Syncing node changes and migrating data

You must deliver the changed ring files to the Object Storage (swift) containers after you copy new ring files to their correct folders.

Important

- Do not migrate all of the data at the same time. Migrate the data in 10% increments. For example, configure the weight of the source device to equal 90.0 and the target device to equal 10.0. Then configure the weight of the source device to equal 80.0 and 20.0. Continue to incrementally migrate the data until you complete the process. During the migration, if you move all of the data at the same time, the old data is on the source device but the ring points to the new target device for all replicas. Until the replicators have moved all of the data to the target device, the data is inaccessible.
- During migration, the Object Storage rings reassign the location of data, and then the replicator moves the data to the new location. As cluster activity increases, the process of replication slows down due to increased load. The larger the cluster, the longer a replication pass takes to complete. This is the expected behavior, but it can result in 404 errors in the log files if a client accesses the data that is currently being relocated. When a proxy attempts to retrieve data from a new location, but the data is not yet in the new location, **swift-proxy** reports a 404 error in the log files.
When the migration is gradual, the proxy accesses replicas that are not being moved and no error occurs. When the proxy attempts to retrieve the data from an alternative replica, 404 errors in log files are resolved. To confirm that the replication process is progressing, refer to the replication logs. The Object Storage service (swift) issues replication logs every five minutes.

Procedure

1. Use a script, similar to the following example, to distribute ring files from a previously existing Controller node to all Controller nodes and restart the Object Storage service containers on those nodes:

```
#!/bin/sh
set -xe
```



```
SRC="tripleo-admin@overcloud-controller-0.ctlplane"
ALL="tripleo-admin@overcloud-controller-0.ctlplane tripleo-admin@overcloud-controller-1.ctlplane tripleo-admin@overcloud-controller-2.ctlplane"
```

- Fetch the current set of ring files:

```
ssh "${SRC}" 'sudo tar -czvf - /var/lib/config-data/puppet-generated/swift_ringbuilder/etc/swift/{*.builder,*.ring.gz,backups/*.builder}' > swift-rings.tar.gz
```

- Upload the rings to all nodes and restart Object Storage services:

```
for DST in ${ALL}; do
  cat swift-rings.tar.gz | ssh "${DST}" 'sudo tar -C / -xvzf -'
  ssh "${DST}" 'sudo podman restart swift_copy_rings'
  ssh "${DST}" 'sudo systemctl restart tripleo_swift*'
done
```

2. To confirm that the data is being moved to the new disk, run the following command on the new storage node:

```
$ sudo grep -i replication /var/log/container/swift/swift.log
```

5.3.6. Removing Object Storage nodes

There are two methods to remove an Object Storage (swift) node:

- Simple removal: This method removes the node in one action and is appropriate for an efficiently-powered cluster with smaller quantities of data.
- Incremental removal: Alter the rings to decrease the weight of the disks on the node that you want to remove. This method is appropriate if you want to minimize impact on storage network usage or if your cluster contains larger quantities of data.

For both methods, you follow the **Scaling down bare-metal nodes** procedure. However, for incremental removal, complete these prerequisites to alter the storage rings to decrease the weight of the disks on the node that you want to remove:

Prerequisites

- Object Storage rings are updated and rebalanced. For more information, see [Updating and rebalancing the Object Storage rings](#).
- Changes in the Object Storage rings are synchronized. For more information, see [Syncing node changes and migrating data](#).

For information about replacing an Object Storage node, see the prerequisites at the beginning of the **Scaling down bare-metal nodes** procedure .

5.3.7. Scaling down bare-metal nodes

To scale down the number of bare-metal nodes in your overcloud, tag the nodes that you want to delete from the stack in the node definition file, redeploy the overcloud, and then delete the bare-metal node from the overcloud.

Prerequisites

- A successful undercloud installation. For more information, see [Installing director on the undercloud](#).
- A successful overcloud deployment. For more information, see [Configuring a basic overcloud with pre-provisioned nodes](#).
- If you are replacing an Object Storage node, replicate data from the node you are removing to the new replacement node. Wait for a replication pass to finish on the new node. Check the replication pass progress in the `/var/log/swift/swift.log` file. When the pass finishes, the Object Storage service (swift) adds entries to the log similar to the following example:

```
Mar 29 08:49:05 localhost object-server: Object replication complete.  
Mar 29 08:49:11 localhost container-server: Replication run OVER  
Mar 29 08:49:13 localhost account-server: Replication run OVER
```

Procedure

1. Source the **stackrc** undercloud credential file:

```
$ source ~/stackrc
```

2. Decrement the **count** parameter in the **overcloud-baremetal-deploy.yaml** file, for the roles that you want to scale down.
3. Define the **hostname** and **name** of each node that you want to remove from the stack, if they are not already defined in the **instances** attribute for the role.
4. Add the attribute **provisioned: false** to the node that you want to remove. For example, to remove the node **overcloud-objectstorage-1** from the stack, include the following snippet in your **overcloud-baremetal-deploy.yaml** file:

```
- name: ObjectStorage  
  count: 3  
  instances:  
    - hostname: overcloud-objectstorage-0  
      name: node00  
    - hostname: overcloud-objectstorage-1  
      name: node01  
      # Removed from cluster due to disk failure  
      provisioned: false  
    - hostname: overcloud-objectstorage-2  
      name: node02  
    - hostname: overcloud-objectstorage-3  
      name: node03
```

After you redeploy the overcloud, the nodes that you define with the **provisioned: false** attribute are no longer present in the stack. However, these nodes are still running in a provisioned state.

**NOTE**

To remove a node from the stack temporarily, deploy the overcloud with the attribute **provisioned: false** and then redeploy the overcloud with the attribute **provisioned: true** to return the node to the stack.

5. Delete the node from the overcloud:

```
(undercloud)$ openstack overcloud node delete \
--stack <stack> \
--baremetal-deployment /home/stack/templates/overcloud-baremetal-deploy.yaml
```

- Replace **<stack>** with the name of the stack for which the bare-metal nodes are provisioned. If not specified, the default is **overcloud**.

**NOTE**

Do not include the nodes that you want to remove from the stack as command arguments in the **openstack overcloud node delete** command.

6. Provision the overcloud nodes to generate an updated heat environment file for inclusion in the deployment command:

```
(undercloud)$ openstack overcloud node provision \
--stack <stack> \
--output <deployment_file> \
/home/stack/templates/overcloud-baremetal-deploy.yaml
```

- Replace **<deployment_file>** with the name of the heat environment file to generate for inclusion in the deployment command, for example **/home/stack/templates/overcloud-baremetal-deployed.yaml**.

7. Add the **overcloud-baremetal-deployed.yaml** file generated by the provisioning command to the stack with your other environment files, and deploy the overcloud:

```
(undercloud)$ openstack overcloud deploy \
...
-e /usr/share/openstack-tripleo-heat-templates/environments/deployed-server-
environment.yaml \
-e /home/stack/templates/overcloud-baremetal-deployed.yaml \
--deployed-server \
--disable-validations \
...
```

5.4. CONTAINER MANAGEMENT IN THE OBJECT STORAGE SERVICE

To help with organization in the Object Storage service (swift), you can use pseudo folders. These folders are logical devices that can contain objects and be nested. For example, you might create an *Images* folder in which to store pictures and a *Media* folder in which to store videos.

You can create one or more containers in each project, and one or more objects or pseudo folders in each container.

5.4.1. Creating private and public containers

Use the dashboard to create a container in the Object Storage service (swift).

Procedure

1. In the dashboard, select **Project > Object Store > Containers**
2. Click **Create Container**.
3. Specify the **Container Name**, and select one of the following in the **Container Access** field.

Type	Description
Private	Limits access to a user in the current project.
Public	Permits API access to anyone with the public URL. However, in the dashboard, project users cannot see public containers and data from other projects.

4. Click **Create Container**.
5. Optional: New containers use the default storage policy. If you have multiple storage policies defined, for example, a default policy and another policy that enables erasure coding, you can configure a container to use a non-default storage policy:

```
$ swift post -H "X-Storage-Policy:<policy>" <container_name>
```

- Replace **<policy>** with the name or alias of the policy that you want the container to use.
- Replace **<container_name>** with the name of the container.

5.4.2. Creating pseudo folders for containers

Use the dashboard to create a pseudo folder for a container in the Object Storage service (swift).

Procedure

1. In the dashboard, select **Project > Object Store > Containers**
2. Click the name of the container to which you want to add the pseudo folder.
3. Click **Create Pseudo-folder**.
4. Specify the name in the **Pseudo-folder Name** field, and click **Create**.

5.4.3. Deleting containers from the Object Storage service

Use the dashboard to delete a container from the Object Storage service (swift).

Procedure

1. In the dashboard, select **Project > Object Store > Containers**

2. Browse for the container in the **Containers** section, and ensure that all objects are deleted. For more information, see [Deleting objects from the Object Storage service](#) .
3. Select **Delete Container** in the container arrow menu.
4. Click **Delete Container** to confirm the container removal.

5.4.4. Uploading objects to containers

If you do not upload an actual file to the Object Storage service (swift), the object is still created as a placeholder that you can use later to upload the file.

Procedure

1. In the dashboard, select **Project > Object Store > Containers**
2. Click the name of the container in which you want to place the uploaded object. If a pseudo folder already exists in the container, you can click its name.
3. Browse for your file, and click **Upload Object**.
4. Specify a name in the **Object Name** field:
 - You can specify pseudo folders in the name by using a / character, for example, *Images/myImage.jpg*. If the specified folder does not already exist, it is created when the object is uploaded.
 - A name that is not unique to the location, that is, the object already exists, overwrites the contents of the object.
5. Click **Upload Object**.

5.4.5. Copying objects between containers

Use the dashboard to copy an object in the Object Storage service (swift).

Procedure

1. In the dashboard, select **Project > Object Store > Containers**
2. Click the name of the object's container or folder (to display the object).
3. Click **Upload Object**.
4. Browse for the file to be copied, and select **Copy** in its arrow menu.
5. Specify the following:

Field	Description
Destination container	Target container for the new object.
Path	Pseudo-folder in the destination container; if the folder does not already exist, it is created.

Field	Description
Destination object name	New object's name. If you use a name that is not unique to the location (that is, the object already exists), it overwrites the object's previous contents.

6. Click **Copy Object**.

5.4.6. Deleting objects from the Object Storage service

Use the dashboard to delete an object from the Object Storage service (swift).

Procedure

1. In the dashboard, select **Project > Object Store > Containers**
2. Browse for the object, and select **Delete Object** in its arrow menu.
3. Click **Delete Object** to confirm the object is removed.

CHAPTER 6. CONFIGURING THE SHARED FILE SYSTEMS SERVICE (MANILA)

With the Shared File Systems service (manila), you can provision shared file systems that multiple compute instances, bare-metal nodes, or containers can consume. Cloud administrators create share types to prepare the share service and enable end users to create and manage shares.

Prerequisites

- An end user requires at least one share type to use the Shared File Systems service.
- For back ends where **driver_handles_share_servers=False**, a cloud administrator configures the requisite networking in advance rather than dynamically in the shared file system back end.
- For a CephFS through NFS back end, a cloud administrator deploys Red Hat OpenStack Platform (RHOSP) director with isolated networks and environment arguments and a custom **network_data** file to create an isolated StorageNFS network for NFS exports. After deployment, before the overcloud is used, the administrator creates a corresponding Networking service (neutron) StorageNFS shared provider network that maps to the isolated StorageNFS network of the data center.
- For a Compute instance to connect to this shared provider network, the user must add an additional neutron port.

6.1. CONFIGURING SHARED FILE SYSTEMS SERVICE BACK ENDS

When cloud administrators use Red Hat OpenStack Platform (RHOSP) director to deploy the Shared File Systems service (manila), they can choose one or more of the following supported back ends:

- Native CephFS and CephFS NFS. For more information, see [Deploying Red Hat Ceph Storage and Red Hat OpenStack Platform together with director](#).
- NetApp. For more information, see the NetApp documentation: [Shared File Systems Service \(Manila\)](#).
- Dell EMC Unity, Dell VNX, or Dell PowerMax. For more information, see the Dell documentation: [Dell EMC Manila Backend Deployment Guide for Red Hat OpenStack Platform 16](#).

For a complete list of supported back-end appliances and drivers, see [Component, Plug-In, and Driver Support in RHEL OpenStack Platform](#).

6.1.1. Configuring multiple back ends

A back end is a storage system or technology that is paired with the Shared File Systems service (manila) driver to export file systems. The Shared File Systems service requires at least one back end to operate. In many cases, one back end is sufficient. However, you can also use multiple back ends in a single Shared File Systems service installation.



IMPORTANT

Currently, Red Hat OpenStack Platform (RHOSP) does not support multiple instances of the same back end to a Shared File Systems service deployment. For example, you cannot add two Ceph Storage clusters as back ends within the same deployment. The back ends that you choose must be heterogeneous. CephFS native and CephFS through NFS are considered one back end with different protocols.

The scheduler for the Shared File Systems service determines the destination back end for share creation requests. A single back end in the Shared File Systems service can expose multiple storage pools.

When you configure multiple back ends, the scheduler chooses one storage pool to create a resource from all the pools exposed by all configured back ends. This is abstracted from the end user. End users see only the capabilities that are exposed by the cloud administrator.

6.1.2. Deploying multiple back ends

By default, a standard Shared File Systems deployment environment file has a single back end. Use the following example procedure to add multiple back ends to the Shared File Systems service (manila) and deploy an environment with a CephFS-Ganesha and a NetApp back end.

Prerequisites

- At least two back ends.
- If a back end requires a custom container, you must use one from the [Red Hat Ecosystem Catalog](#) instead of the standard Shared File Systems service container. For example, if you want to use a Dell EMC Unity storage system back end with Ceph, choose the Dell EMC Unity container from the catalog.

Procedure

1. Create a storage customization YAML file. You can use this file to provide any values or overrides that suit your environment:

```
$ vi /home/stack/templates/storage_customizations.yaml
```

2. Configure the storage customization YAML file to include any overrides, including enabling multiple back ends:

```
parameter_defaults:
  ManilaEnabledShareProtocols:
    - NFS
  ManilaNetappLogin: '<login_name>'
  ManilaNetappPassword: '<password>'
  ManilaNetappServerHostname: '<netapp-hostname>'
  ManilaNetappVserver: '<netapp-vserver>'
  ManilaNetappDriverHandlesShareServers: 'false'
```

3. Specify the back-end templates by using the **openstack overcloud deploy** command. The example configuration enables the Shared File Systems service with a NetApp back end and a CephFS through NFS back end.

**NOTE**

Execute **source ~/stackrc** before issuing the **openstack overcloud deploy** command.

```
[stack@undercloud ~]$ source ~/stackrc
$ openstack overcloud deploy \
  --timeout 100 \
  --stack overcloud \
  --templates /usr/share/openstack-tripleo-heat-templates \
  -n /usr/share/openstack-tripleo-heat-templates/network_data_ganesha.yaml \
  -e /usr/share/openstack-tripleo-heat-templates/environments/ceph-ansible/ceph-mds.yaml \
  -e /usr/share/openstack-tripleo-heat-templates/environments/ceph-ansible/ceph-
ansible.yaml \
  -r /home/stack/templates/roles/roles_data.yaml \
  -e /usr/share/openstack-tripleo-heat-templates/environments/manila-cephfs-ganesha-
config.yaml \
  -e /usr/share/openstack-tripleo-heat-templates/environments/manila-netapp-config.yaml \
  -e /home/stack/templates/storage_customizations.yaml \
  ...
```

Additional resources

- For more information about the **ManilaEnabledShareProtocols** parameter, see [Section 6.1.4, “Overriding allowed NAS protocols”](#).
- For more information about the deployment command, see [Director Installation and Usage](#).

6.1.3. Confirming deployment of multiple back ends

Use the **manila service-list** command to verify that your back ends deployed successfully. If you use a health check on multiple back ends, a ping test returns a response even if one of the back ends is unresponsive, so this is not a reliable way to verify your deployment.

Procedure

1. Log in to the undercloud host as the **stack** user.
2. Source the **overcloudrc** credentials file:

```
$ source ~/overcloudrc
```

3. Confirm the list of Shared File Systems service back ends:

```
$ manila service-list
+-----+-----+-----+-----+-----+-----+-----+
| Id | Binary | Host | Zone | Status | State | Updated_at |
+-----+-----+-----+-----+-----+-----+-----+
| 2 | manila-scheduler | hostgroup | nova | enabled | up | 2021-03-24T16:49:09.000000 |
| 5 | manila-share | hostgroup@cephfs | nova | enabled | up | 2021-03-24T16:49:12.000000 |
| 8 | manila-share | hostgroup@tripleo_netapp | nova | enabled | up | 2021-03-
24T16:49:06.000000 |
```

The status of each successfully deployed back end shows **enabled** and the state shows **up**.

6.1.4. Overriding allowed NAS protocols

The Shared File Systems service can export shares in one of many network attached storage (NAS) protocols, such as NFS, CIFS, or CEPHFS. By default, the Shared File Systems service enables all of the NAS protocols supported by the back ends in a deployment.

As a Red Hat OpenStack Platform (RHOSP) administrator, you can override the **ManilaEnabledShareProtocols** parameter and list only the protocols that you want to allow in your cloud.

For example, if back ends in your deployment support both NFS and CIFS, you can override the default value and enable only one protocol.

Procedure

1. Log in to the undercloud host as the **stack** user.
2. Source the **overcloudrc** credentials file:

```
$ source ~/overcloudrc
```

3. Create a storage customization YAML file. This file can be used to provide any values or overrides that suit your environment:

```
$ vi /home/stack/templates/storage_customizations.yaml
```

4. Configure the **ManilaEnabledShareProtocols** parameter with the values that you want:

```
parameter_defaults:
  ManilaEnabledShareProtocols:
    - NFS
    - CEPHFS
```

5. Include the environment file that contains your new content in the **openstack overcloud deploy** command by using the **-e** option. Ensure that you include all other environment files that are relevant to your deployment.

```
$ openstack overcloud deploy \
...
-e /home/stack/templates/storage_customizations.yaml
```



NOTE

The deployment does not validate the settings. The NAS protocols that you assign must be supported by the back ends in your Shared File Systems service deployment.

6.1.5. Viewing back-end capabilities

The scheduler component of the Shared File Systems service (manila) makes intelligent placement decisions based on several factors such as capacity, provisioning configuration, placement hints, and the capabilities that the back-end storage system driver detects and exposes.

Procedure

- Run the following command to view the available capabilities:

```
$ manila pool-list --detail
+-----+
| Property          | Value          |
+-----+
| name              | hostgroup@cephfs#cephfs |
| pool_name         | cephfs         |
| total_capacity_gb | 1978           |
| free_capacity_gb  | 1812           |
| ...              |                |
| driver_handles_share_servers | False          |
| snapshot_support  | True           |
| create_share_from_snapshot_support | False          |
| revert_to_snapshot_support | False          |
| mount_snapshot_support | False          |
| ...              |                |
+-----+
+-----+
| Property          | Value          |
+-----+
| name              | hostgroup@tripleo_netapp#aggr1_n1 |
| pool_name         | aggr1_n1       |
| total_capacity_gb | 6342.1         |
| free_capacity_gb  | 6161.99        |
| ...              |                |
| driver_handles_share_servers | False          |
| mount_snapshot_support | False          |
| replication_type   | None           |
| replication_domain | None           |
| sg_consistent_snapshot_support | host          |
| ipv4_support       | True           |
| ipv6_support       | False          |
+-----+
+-----+
| Property          | Value          |
+-----+
| name              | hostgroup@tripleo_netapp#aggr1_n2 |
| pool_name         | aggr1_n2       |
| total_capacity_gb | 6342.1         |
| free_capacity_gb  | 6209.26        |
| ...              |                |
| snapshot_support  | True           |
| create_share_from_snapshot_support | True          |
| revert_to_snapshot_support | True          |
| driver_handles_share_servers | False          |
| mount_snapshot_support | False          |
| replication_type   | None           |
| replication_domain | None           |
| sg_consistent_snapshot_support | host          |
```

ipv4_support	True	
ipv6_support	False	
+-----+-----+		

Related information

To influence placement decisions, as an administrator, you can use share types and extra specifications. For more information about share types, see [Creating share types](#).

6.2. CREATING SHARE TYPES

Share types serve as hints to the Shared File Systems service scheduler to perform placement decisions. Red Hat OpenStack Platform (RHOSP) director configures the Shared File Systems service with a default share type named default, but does not create the share type.



IMPORTANT

An end user requires at least one share type to use the Shared File Systems service.

Procedure


1. After you deploy the overcloud, run the following command as the cloud administrator to create a share type:

```
$ manila type-create default <spec_driver_handles_share_servers>
```

The **<spec_driver_handles_share_servers>** parameter is a Boolean value:

- For CephFS through NFS or native CephFS, the value is false.
 - For other back ends, the value can be true or false. Set **<spec_driver_handles_share_servers>** to match the value of the **Manila<backend>DriverHandlesShareServers** parameter. For example, if you use a NetApp back end, the parameter is called **ManilaNetappDriverHandlesShareServers**.
2. Add specifications to the default share type or create additional share types to use with multiple configured back ends. For example, configure the default share type to select a CephFS back end and an additional share type that uses a NetApp **driver_handles_share_servers=True** back end:

```
(overcloud) [stack@undercloud-0 ~]$ manila type-create default false --extra-specs
share_backend_name='cephfs'
(overcloud) [stack@undercloud-0 ~]$ manila type-create netapp true --extra-specs
share_backend_name='tripleo_netapp'
```



NOTE

By default, share types are public, which means they are visible to and usable by all cloud projects. However, you can create private share types for use within specific projects.

Additional resources

- For more information about how to make private share types or set additional share-type options, see the [Security and Hardening Guide](#).

6.3. COMPARING COMMON CAPABILITIES OF SHARE TYPES

Share types define the common capabilities of shares. Review the common capabilities of share types to understand what you can do with your shares.

Table 6.1. Capabilities of share types

Capability	Values	Description
driver_handles_share_servers	true or false	Grants permission to use share networks to create shares.
snapshot_support	true or false	Grants permission to create snapshots of shares.
create_share_from_snapshot_support	true or false	Grants permission to create clones of share snapshots.
revert_to_snapshot_support	true or false	Grants permission to revert your shares to the most recent snapshot.
mount_snapshot_support	true or false	Grants permission to export and mount your snapshots.
replication_type	dr	Grants permission to create replicas for disaster recovery. Only one active export is allowed at a time.
	readable	Grants permission to create read-only replicas. Only one writable, active export is allowed at a time.
	writable	Grants permission to create read/write replicas. Any number of active exports are allowed at a time per share.
availability_zones	a list of one or more availability zones	Grants permission to create shares only on the availability zones listed.

6.4. PLANNING NETWORKING FOR SHARED FILE SYSTEMS

Shared file systems are accessed over a network. It is important to plan the networking on your cloud to ensure that end user clients can connect their shares to workloads that run on Red Hat OpenStack Platform (RHOSP) virtual machines, bare-metal servers, and containers.

Depending on the level of security and isolation required for end users, as an administrator, you can set the **driver_handles_share_servers** parameter to true or false.

If you set the **driver_handles_share_servers** parameter to true, this enables the service to export shares to end user-defined share networks with the help of isolated share servers.

When the **driver_handles_share_servers** parameter equals true, users can provision their workloads on self-service share networks. This ensures that their shares are exported by completely isolated NAS file servers on dedicated network segments.

The share networks used by end users can be the same as the private project networks that they can create. As an administrator, you must ensure that the physical network to which you map these isolated networks extends to your storage infrastructure.

You must also ensure that the network segmentation style by project networks is supported by the storage system used. Storage systems, such as NetApp ONTAP and Dell EMC PowerMax, Unity, and VNX, do not support virtual overlay segmentation styles such as GENEVE or VXLAN.

As an alternative, you can terminate the overlay networking at top-of-rack switches and use a more primitive form of networking for your project networks, such as VLAN. Another alternative is to allow VLAN segments on shared provider networks or provide access to a pre-existing segmented network that is already connected to your storage system.

If you set the **driver_handles_share_servers** parameter to false, users cannot create shares on their own share networks. Instead, they must connect their clients to the network configured by the cloud administrator.

When the **driver_handles_share_servers** parameter equals false, director can create a dedicated shared storage network for you. For example, when you deploy the native CephFS back end with standard director templates, director creates a shared provider network called **Storage**. When you deploy CephFS through the NFS back end, the shared provider network is called **StorageNFS**. Your end users must connect their clients to the shared storage network to access their shares.

Not all shared file system storage drivers support both modes of operation. Regardless of which mode you choose, the service ensures hard data path multi-tenancy isolation guarantees.

If you want to offer hard network path multi-tenancy isolation guarantees to tenant workloads as part of a self-service model, you must deploy with back ends that support the **driver_handles_share_servers** driver mode.

For information about network connectivity to the share, see [Section 6.5, “Ensuring network connectivity to the share”](#)

6.5. ENSURING NETWORK CONNECTIVITY TO THE SHARE

Clients that need to connect to a file share must have network connectivity to one or more of the export locations for that share.

There are many ways to configure networking with the Shared File Systems service, including using network plugins.

When the **driver_handles_share_servers** parameter for a share type equals true, a cloud user can create a share network with the details of a network to which the compute instance attaches and then reference it when creating shares.

When the **driver_handles_share_servers** parameter for a share type equals false, a cloud user must connect their compute instance to the shared storage network.

For more information about how to configure and validate network connectivity to a shared network, see [Section 7.5, “Connecting to a shared network to access shares”](#).

6.6. CHANGING THE DEFAULT QUOTAS IN THE SHARED FILE SYSTEMS SERVICE

To prevent system capacities from being exhausted without notification, cloud administrators can configure quotas. Quotas are operational limits. The Shared File Systems service (manila) enforces some sensible limits by default. These limits are called default quotas. Cloud administrators can override default quotas so that individual projects have different consumption limits.

6.6.1. Updating quotas for projects, users, and share types

As a cloud administrator, you can list the quotas for a project or user by using the **manila quota-show** command.

You can update quotas for all users in a project, or a specific project user, or a share type used by the project users. You can update the following quotas for the target you choose:

- **shares**: Number of shares that you can create.
- **snapshots**: Number of snapshots that you can create.
- **gigabytes**: Total size in GB that you can allocate for all shares.
- **snapshot-gigabytes**: Total size in GB that you can allocate for all snapshots of shares.
- **share-networks**: Total number of share networks that you can create.
- **share_groups**: Total number of share groups that you can create.
- **share_group_snapshots**: Total number of share group snapshots that you can create.
- **share-replicas**: Total number of share replicas that you can create.
- **replica-gigabytes**: Total size in GB that you can allocate across all share replicas.



NOTE

You can only specify **share-type** quotas at the project level. You cannot set **share-type** quotas for specific project users.



IMPORTANT

In the following procedures, enter the values carefully. The Shared File Systems service does not detect or report incorrect values.

Procedure

1. You can use the following commands to view quotas. If you include the **--user** option, you can view the quota for a specific user in the specified project. If you omit the **--user** option, you can view the quotas that apply to all users for the specified project. Similarly, if you include the optional **--share-type**, you can view the quota for a specific share type as it relates to the project. The **--user** and **--share-type** options are mutually exclusive.

■

\$ manila quota-show

- Example for a project:

```
$ manila quota-show --project af2838436f3f4cf6896399dd97c4c050
+-----+-----+
| Property      | Value                |
+-----+-----+
| gigabytes     | 1000                 |
| id            | af2838436f3f4cf6896399dd97c4c050 |
| replica_gigabytes | 1000                 |
| share_group_snapshots | 50                   |
| share_groups  | 49                   |
| share_networks | 10                   |
| share_replicas | 100                  |
| shares        | 50                   |
| snapshot_gigabytes | 1000                 |
| snapshots     | 50                   |
+-----+-----+
```

- Example for a project user:

```
$ manila quota-show --project af2838436f3f4cf6896399dd97c4c050 --user
81ebb491dd0e4c2aae0775dd564e76d1
+-----+-----+
| Property      | Value                |
+-----+-----+
| gigabytes     | 500                  |
| id            | af2838436f3f4cf6896399dd97c4c050 |
| replica_gigabytes | 1000                 |
| share_group_snapshots | 50                   |
| share_groups  | 49                   |
| share_networks | 10                   |
| share_replicas | 100                  |
| shares        | 25                   |
| snapshot_gigabytes | 1000                 |
| snapshots     | 50                   |
+-----+-----+
```

- Example for a project for a specific share type:

```
$ manila quota-show --project af2838436f3f4cf6896399dd97c4c050 --share-type
dhss_false
+-----+-----+
| Property      | Value                |
+-----+-----+
| gigabytes     | 1000                 |
| id            | af2838436f3f4cf6896399dd97c4c050 |
| replica_gigabytes | 1000                 |
| share_replicas | 100                  |
| shares        | 15                   |
| snapshot_gigabytes | 1000                 |
| snapshots     | 50                   |
+-----+-----+
```


2. Use the **manila quota-update** command to update the quotas. You can update quotas for all project users, a specific project user, or a share type in a project:

- Update quotas for all users in a project:

```
$ manila quota-update <id> [--shares <share_quota> --gigabytes
<share_gigabytes_quota> ...]
```

Replace **<id>** with the project ID. This value must be the project ID, not the project name.

- Update quotas for a specific user in a project:

```
$ manila quota-update <id> --user <user_id> [--shares <new_share_quota> --gigabytes
<new_share_gigabytes_quota> ...]
```

- Replace **<id>** with the project ID. This value must be the project ID, not the project name.
- Replace **<user_id>** with the user ID. The value must be the user ID, not the user name.

- Update quotas for all users who use a specific share type:

```
$ manila quota-update <id> --share-type <share_type> [--shares <new_share_quota>30
--gigabytes <new-share_gigabytes_quota> ...]
```

- Replace **<id>** with the project ID. This value must be the project ID, not the project name.
- Replace **<share_type>** with the name or ID of the share type that you want to apply the quota to.

Verification

- The **quota-update** command does not produce any output. Use the **quota-show** command to verify that a quota was successfully updated.

6.6.2. Resetting quotas for projects, users, and share types

You can remove quota overrides to return quotas to their default values. The target entity is restricted by the default quota that applies to all projects with no overrides.

Procedure

- Use the **manila quota-delete** command to return quotas to default values. You can return quotas to default values for all project users, a specific project user, or a share type in a project:

- Reset project quotas:

```
$ manila quota-delete --project <id>
```

Replace **<id>** with the project ID. This value must be the project ID, not the project name.

- Reset quotas for a specific user:

```
$ manila quota-delete --project <id> --user <user_id>
```

- Replace **<id>** with the project ID. This value must be the project ID, not the project name.
- Replace **<user_id>** with the user ID. The value must be the user ID, not the user name.
- Reset quotas for a share type used by project users:

```
$ manila quota-delete --project <id> --share-type <share_type>
```

- Replace **<id>** with the project ID. This value must be the project ID, not the project name.
- Replace **<share_type>** with the name or ID of the share type the quota must be reset for.

Verification

1. The **quota-delete** command does not produce any output. Use the **quota-show** command to verify whether a quota was successfully reset.
2. List the default quotas for all projects. Default quotas apply to projects that have no overrides.

```
$ manila quota-class-show default
```

6.6.3. Updating the default quota values for Shared File Systems service projects

As a cloud administrator, you can update default quotas that apply to all projects that do not already have quota overrides.

Procedure

1. View the usage statement of the **manila quota-class-update** command:

```
$ manila help quota-class-update
usage: manila quota-class-update [--shares <shares>] [--snapshots <snapshots>]
      [--gigabytes <gigabytes>]
      [--snapshot-gigabytes <snapshot_gigabytes>]
      [--share-networks <share_networks>]
      [--share-replicas <share_replicas>]
      [--replica-gigabytes <replica_gigabytes>]
      <class_name>
```



NOTE

The parameter **<class_name>** is a positional argument. It identifies the quota class for which the quotas are set. Set the value of this parameter to **default**. No other quota classes are supported.

You can update the values for any of the following optional parameters:

- **--shares <shares>** Adds a new value for the **shares** quota.
- **--snapshots <snapshots>** Adds a new value for the **snapshots** quota.

- **--gigabytes <gigabytes>** Adds a new value for the **gigabytes** quota.
 - **--snapshot-gigabytes <snapshot_gigabytes>** or **--snapshot_gigabytes <snapshot_gigabytes>** Adds a new value for the **snapshot_gigabytes** quota.
 - **--share-networks <share_networks>** or **--share_networks <share_networks>** Adds a new value for the **share_networks** quota.
 - **--share-replicas <share_replicas>**, **--share_replicas <share_replicas>**, or **--replicas <share_replicas>** Adds a new value for the **share_replicas** quota.
 - **--replica-gigabytes <replica_gigabytes>** or **--replica_gigabytes <replica_gigabytes>** Adds a new value for the **replica_gigabytes** quota.
2. Use the information from the usage statement to update the default quotas. The following example updates the default quotas for **shares** and **gigabytes**:

```
$ manila quota-class-update default --shares 30 --gigabytes 512
$ manila quota-class-show default
+-----+-----+
| Property      | Value |
+-----+-----+
| gigabytes     | 512   |
| id            | default |
| replica_gigabytes | 1000 |
| share_group_snapshots | 50   |
| share_groups  | 50    |
| share_networks | 10    |
| share_replicas | 100   |
| shares        | 30    |
| snapshot_gigabytes | 1000 |
| snapshots     | 50    |
+-----+-----+
```

CHAPTER 7. PERFORMING OPERATIONS WITH THE SHARED FILE SYSTEMS SERVICE (MANILA)

Cloud users can create and manage shares from the available share types in the Shared File Systems service (manila).

7.1. DISCOVERING SHARE TYPES

As a cloud user, you must specify a share type when you create a share.

Procedure

- Discover the available share types:

```
$ manila type-list
```

The command output lists the name and ID of the share type.

7.2. CREATING NFS OR NATIVE CEPHFS SHARES

Create an NFS or native CephFS share to read and write data.

To create an NFS or native CephFS share, use a command similar to the following:

```
$ manila create [--share-type <sharetype>] [--name <sharename>] proto GB
```

Replace the following values:

- **sharetype** applies settings associated with the specified share type.
 - Optional: If not supplied, the **default** share type is used.
- **sharename** is the name of the share.
 - Optional: Shares are not required to have a name, nor is the name guaranteed to be unique.
- **proto** is the share protocol you want to use.
 - For CephFS with NFS, **proto** is **nfs**.
 - For native CephFS, **proto** is **cephfs**.
 - For NetApp and Dell EMC storage back ends, **proto** is **nfs** or **cifs**.
- **GB** is the size of the share in gigabytes.

For example, in [Section 6.2, “Creating share types”](#), the cloud administrator created a **default** share type with a CephFS back end and another share type named **netapp** with a NetApp back end.

Procedure

1. Create a 10 GB NFS share named **share-01**. This example does not specify the optional share type because it uses the available **default** share type created by the cloud administrator. It also uses the shared storage network configured by the cloud administrator:

```
(user) $ manila create --name share-01 nfs 10
```

2. Create a 15 GB native CephFS share named **share-02**:

```
(user) $ manila create --name share-02 cephfs 15
```

3. Create a 20 GB NFS share named **share-03** and specify a custom share type and share network:

```
(user) $ manila create --name share-03 --share-type netapp --share-network mynet nfs 20
```

7.3. LISTING SHARES AND EXPORTING INFORMATION

To verify that you successfully created the shares, complete the following steps.

Procedure

1. List the shares:

```
(user) $ manila list
```

Name	...	Status	...	ID
8c3bedd8-bc82-4100-a65d-53ec51b5fe81		share-01	...	available ...


2. View the export locations of the share:

```
(user) $ manila share-export-location-list share-01
```

Path
172.17.5.13:/volumes/_nogroup/e840b4ae-6a04-49ee-9d6e-67d4999fbc01

3. View the parameters for the share:

```
(user) $ manila share-export-location-show <id>
```



NOTE

You use the export location to mount the share in [Section 7.8.2, “Mounting NFS or native CephFS”](#).

7.4. CREATING A SNAPSHOT OF DATA ON A SHARED FILE SYSTEM

A snapshot is a read-only, point-in-time copy of data on a share. You can use a snapshot to recover data lost through accidental data deletion or file system corruption. Snapshots are more space efficient than backups, and they do not impact the performance of the Shared File Systems service (manila).

Prerequisites

- The **snapshot_support** parameter must equal **true** on the parent share. You can run the following command to verify:

```
$ manila show | grep snapshot_support
```

Procedure

1. As a cloud user, create a snapshot of a share:

```
$ manila snapshot-create [--name <snapshot_name>] <share>
```

- Replace **<share>** with the name or ID of the share for which you want to create a snapshot.
- Optional: Replace **<snapshot_name>** with the name of the snapshot.

Example output

```
+-----+-----+
| Property | Value |
+-----+-----+
| id       | dbdcb91b-82ba-407e-a23d-44ffca4da04c |
| share_id | ee7059aa-5887-4b87-b03e-d4f0c27ed735 |
| share_size | 1 |
| created_at | 2022-01-07T14:20:55.541084 |
| status    | creating |
| name     | snapshot_name |
| description | None |
| size     | 1 |
| share_proto | NFS |
| provider_location | None |
| user_id  | 6d414c62237841dcbe63d3707c1cdd90 |
| project_id | 041ff9e24eba469491d770ad8666682d |
+-----+-----+
```

2. Confirm that you created the snapshot:

```
$ manila snapshot-list --share-id <share>
```

Replace **<share>** with the name or ID of the share from which you created the snapshot.

7.4.1. Creating a share from a snapshot

You can create a share from a snapshot. If the parent share the snapshot was created from has a share type with **driver_handles_share_servers** set to **true**, the new share is created on the same share network as the parent.



NOTE

If the share type of the parent share has **driver_handles_share_servers** set to **true**, you cannot change the share network for the share you create from the snapshot.

Prerequisites

- The **create_share_from_snapshot_support** share attribute is set to **true**.
For more information about share types, see [Comparing common capabilities of share types](#).
- The **status** attribute of the snapshot is set to **available**.

Procedure

1. Retrieve the ID of the share snapshot that contains the data that you require for your new share:

```
$ manila snapshot-list
```

2. A share created from a snapshot can be larger, but not smaller, than the snapshot. Retrieve the size of the snapshot:

```
$ manila snapshot-show <snapshot-id>
```

3. Create a share from a snapshot:

```
$ manila create <share_protocol> <size> \  
--snapshot-id <snapshot_id> \  
--name <name>
```

- Replace **<share_protocol>** with the protocol, such as NFS.
 - Replace **<size>** with the size of the share to be created, in GiB.
 - Replace **<snapshot_id>** with the ID of the snapshot.
 - Replace **<name>** with the name of the new share.
4. List the shares to confirm that the share was created successfully:

```
$ manila list
```

5. View the properties of the new share:

```
$ manila show <name>
```

Verification

After you create a snapshot, confirm that the snapshot is available.

- List the snapshots to confirm that they are available:

```
$ manila snapshot-list
```

7.4.2. Deleting a snapshot

When you create a snapshot of a share, you cannot delete the share until you delete all of the snapshots created from that share.

Procedure

1. Identify the snapshot you want to delete and retrieve its ID:

```
$ manila snapshot-list
```

2. Delete the snapshot:

```
$ manila snapshot-delete <snapshot>
```



NOTE

Repeat this step for each snapshot that you want to delete.

3. After you delete the snapshot, run the following command to confirm that you deleted the snapshot:

```
$ manila snapshot-list
```

7.5. CONNECTING TO A SHARED NETWORK TO ACCESS SHARES

When the **driver_handles_share_servers** parameter equals false, shares are exported to the shared provider network that the administrator made available. As an end user, you must connect your client, such as a Compute instance, to the shared provider network to access your shares.

In this example procedure, the shared provider network is called StorageNFS. StorageNFS is configured when director deploys the Shared File Systems service with the CephFS through NFS back end. Follow similar steps to connect to the network made available by your cloud administrator.



NOTE

In the example procedure, the IP address family version of the client is not important. The steps in this procedure use IPv4 addressing, but the steps are identical for IPv6.

Procedure

1. Create a security group for the StorageNFS port that allows packets to egress the port, but which does not allow ingress packets from unestablished connections:

```
(user) $ openstack security group create no-ingress -f yaml
created_at: '2018-09-19T08:19:58Z'
description: no-ingress
id: 66f67c24-cd8b-45e2-b60f-9eaedc79e3c5
name: no-ingress
project_id: 1e021e8b322a40968484e1af538b8b63
revision_number: 2
rules: 'created_at="2018-09-19T08:19:58Z", direction="egress", ethertype="IPv4",
id="6c7f643f-3715-4df5-9fef-0850fb6eaaf2", updated_at="2018-09-19T08:19:58Z"

created_at="2018-09-19T08:19:58Z", direction="egress", ethertype="IPv6",
id="a8ca1ac2-fbe5-40e9-ab67-3e55b7a8632a", updated_at="2018-09-19T08:19:58Z"
updated_at: '2018-09-19T08:19:58Z'
```

2. Create a port on the StorageNFS network with security enforced by the **no-ingress** security group.


```
(user) $ openstack port create nfs-port0 --network StorageNFS --security-group no-ingress -f
yaml

admin_state_up: UP
allowed_address_pairs: ""
binding_host_id: null
binding_profile: null
binding_vif_details: null
binding_vif_type: null
binding_vnic_type: normal
created_at: '2018-09-19T08:03:02Z'
data_plane_status: null
description: ""
device_id: ""
device_owner: ""
dns_assignment: null
dns_name: null
extra_dhcp_opts: ""
fixed_ips: ip_address='172.17.5.160', subnet_id='7bc188ae-aab3-425b-a894-863e4b664192'
id: 7a91cbbc-8821-4d20-a24c-99c07178e5f7
ip_address: null
mac_address: fa:16:3e:be:41:6f
name: nfs-port0
network_id: cb2cbc5f-ea92-4c2d-beb8-d9b10e10efae
option_name: null
option_value: null
port_security_enabled: true
project_id: 1e021e8b322a40968484e1af538b8b63
qos_policy_id: null
revision_number: 6
security_group_ids: 66f67c24-cd8b-45e2-b60f-9eaedc79e3c5
status: DOWN
subnet_id: null
tags: ""
trunk_details: null
updated_at: '2018-09-19T08:03:03Z'
```


**NOTE**

StorageNFSSubnet assigned IP address 172.17.5.160 to **nfs-port0**.

3. Add **nfs-port0** to a Compute instance.

```
(user) $ openstack server add port instance0 nfs-port0
(user) $ openstack server list -f yaml
- Flavor: m1.micro
  ID: 0b878c11-e791-434b-ab63-274ecfc957e8
  Image: manila-test
  Name: demo-instance0
  Networks: demo-network=172.20.0.4, 10.0.0.53; StorageNFS=172.17.5.160
  Status: ACTIVE
```

In addition to its private and floating addresses, the Compute instance is assigned a port with the IP address 172.17.5.160 on the StorageNFS network that you can use to mount NFS shares when access is granted to that address for the share in question.



NOTE

You might need to adjust the networking configuration on the Compute instance and restart the services for the Compute instance to activate an interface with this address.

7.6. CONFIGURING AN IPV6 INTERFACE BETWEEN THE NETWORK AND AN INSTANCE

When the shared network to which shares are exported uses IPv6 addressing, you might experience an issue with DHCPv6 on the secondary interface. If this issue occurs, configure an IPv6 interface manually on the instance.

Prerequisites

- Connection to a shared network to access shares

Procedure

1. Log in to the instance.
2. Configure the IPv6 interface address:

```
$ sudo ip address add fd00:fd00:fd00:7000::c/64 dev eth1
```

3. Activate the interface:

```
$ sudo ip link set dev eth1 up
```

4. Ping the IPv6 address in the export location of the share to test interface connectivity:

```
$ ping -6 fd00:fd00:fd00:7000::21
```

5. Alternatively, verify that you can reach the NFS server through Telnet:

```
$ sudo dnf install -y telnet
$ telnet fd00:fd00:fd00:7000::21 2049
```

7.7. GRANTING SHARE ACCESS FOR END-USER CLIENTS

You must grant end-user clients access to the share so that users can read data from and write data to the share.

You grant a client compute instance access to an NFS share through the IP address of the instance. The **user** rules for CIFS shares and **cephx** rules for CephFS shares follow a similar pattern. With **user** and **cephx** access types, you can use the same **clientidentifier** across multiple clients, if required.

Before you can mount a share on a client, such as a compute instance, you must grant the client access to the share by using a command similar to the following command:

```
$ manila access-allow <share> <accesstype> --access-level <accesslevel> <clientidentifier>
```

Replace the following values:

- **share**: The share name or ID of the share created in [Section 7.2, “Creating NFS or native CephFS shares”](#).
- **accesstype**: The type of access to be requested on the share. Some types include:
 - **user**: Use to authenticate by user or group name.
 - **ip**: Use to authenticate an instance through its IP address.
 - **cephx**: Use to authenticate by native CephFS client username.



NOTE

The type of access depends on the protocol of the share. For CIFS, you can use **user**. For NFS shares, you must use **ip**. For native CephFS shares, you must use **cephx**.

- **accesslevel**: Optional; the default is **rw**.
 - **rw**: Read-write access to shares.
 - **ro**: Read-only access to shares.
- **clientidentifier**: Varies depending on **accesstype**.
 - Use an IP address for **ip accesstype**.
 - Use a CIFS user or group for **user accesstype**.
 - Use a username string for **cephx accesstype**.

7.7.1. Granting access to an NFS share

You provide access to NFS shares through IP addresses.



NOTE

In the example procedure, the IP address family version of the client is not important. The steps in this procedure use IPv4 addressing, but the steps are identical for IPv6.

Procedure

1. Retrieve the IP address of the client compute instance where you plan to mount the share. Make sure that you pick the IP address that corresponds to the network that can reach the shares. In this example, it is the IP address of the StorageNFS network:

```
(user) $ openstack server list -f yaml
- Flavor: m1.micro
```

```
ID: 0b878c11-e791-434b-ab63-274ecfc957e8
Image: manila-test
Name: demo-instance0
Networks: demo-network=172.20.0.4, 10.0.0.53;
StorageNFS=172.17.5.160
Status: ACTIVE
```

```
(user) $ manila access-allow share-01 ip 172.17.5.160
```

**NOTE**

Access to the share has its own ID (**accessid**).

```
+-----+-----+
| Property | Value |
+-----+-----+
| access_key | None |
| share_id | db3bedd8-bc82-4100-a65d-53ec51b5cba3 |
| created_at | 2018-09-17T21:57:42.000000 |
| updated_at | None |
| access_type | ip |
| access_to | 172.17.5.160 |
| access_level | rw |
| state | queued_to_apply |
| id | 875c6251-c17e-4c45-8516-fe0928004fff |
+-----+-----+
```

2. Verify that the access configuration was successful:

```
(user) $ manila access-list share-01

+-----+-----+-----+-----+-----+ ...
| id      | access_type | access_to | access_level | state | ...
+-----+-----+-----+-----+-----+
| 875c6251-... | ip      | 172.17.5.160 | rw      | active | ...
+-----+-----+-----+-----+-----+ ...
```

7.7.2. Granting access to a native CephFS share

You provide access to native CephFS shares through Ceph client usernames. The Shared File Systems service (manila) prevents the use of pre-existing Ceph users so you must create unique Ceph client usernames.

To mount a share, you need a Ceph client username and an access key. You can retrieve access keys by using the Shared File Systems service API. By default, access keys are visible to all users in a project namespace. You can provide the same user with access to different shares in the project namespace. Users can then access the shares by using the CephFS kernel client on the client machine.

**IMPORTANT**

Use the native CephFS driver with trusted clients only. For information about native CephFS back-end security, see [Native CephFS back-end security](#) in *Deploying Red Hat Ceph Storage and Red Hat OpenStack Platform together with director*.

Procedure

1. Grant users access to the native CephFS share:

```
$ manila access-allow <share-02> cephx <user-01>
```

- Replace **<share-02>** with either the share name or the share ID.
- Replace **<user-01>** with the cephx user.

2. Collect the access key for the user:

```
$ manila access-list <share-02>
```

7.7.3. Revoking access to a share

The owner of a share can revoke access to the share for any reason. Complete the following steps to revoke previously-granted access to a share.

Procedure

- Revoke access to a share:

```
$ manila access-deny <share> <accessid>
```

- Replace **<share>** with either the share name or the share ID.
For example:

```
(user) $ manila access-list share-01
+-----+-----+-----+-----+-----+
| id      | access_type | access_to | access_level | state | ...
+-----+-----+-----+-----+-----+ ...
| 875c6251-... | ip      | 172.17.5.160 | rw      | active | ...
+-----+-----+-----+-----+-----+

(user) $ manila access-deny share-01 875c6251-c17e-4c45-8516-fe0928004fff

(user) $ manila access-list share-01

+-----+-----+-----+-----+ ...
| id      | access_type | access_to | access_level | state | ...
+-----+-----+-----+-----+ ...
+-----+-----+-----+-----+ ...
```



NOTE

If you have an existing client that has read-write permissions, you must revoke their access to a share and add a read-only rule if you want the client to have read-only permissions.

7.8. MOUNTING SHARES ON COMPUTE INSTANCES

After you grant share access to clients, the clients can mount and use the shares. Any type of client can access shares as long as there is network connectivity to the client.

The steps used to mount an NFS share on a virtual compute instance are similar to the steps to mount an NFS share on a bare-metal compute instance. For more information about how to mount shares on OpenShift containers, see [Product Documentation for OpenShift Container Platform](#).



NOTE

Client packages for the different protocols must be installed on the Compute instance that mounts the shares. For example, for the Shared File Systems service with CephFS through NFS, the NFS client packages must support NFS 4.1.

7.8.1. Listing share export locations

Retrieve the export locations of shares so that you can mount a share.

Procedure

- Retrieve the export locations of a share:

```
(user) $ manila share-export-location-list share-01
```

When multiple export locations exist, choose one for which the value of the **preferred** metadata field equals **True**. If no preferred locations exist, you can use any export location.

7.8.2. Mounting NFS or native CephFS

After you create NFS or native CephFS shares and grant share access to end-user clients, users can mount the shares on the client to enable access to data. Any type of client can access shares as long as there is network connectivity to the client.

Prerequisites

- To mount NFS shares, the **nfs-utils** package must be installed on the client machine.
- To mount native CephFS shares, the **ceph-common** package must be installed on the client machine. Users access native CephFS shares by using the CephFS kernel client on the client machine.

Procedure

1. Log in to the instance:

```
(user) $ openstack server ssh demo-instance0 --login user
```

2. To mount an NFS share, refer to the following example for sample syntax:

```
$ mount -t nfs -v <172.17.5.13:/volumes/_nogroup/e840b4ae-6a04-49ee-9d6e-67d4999fbc01> /mnt
```

- Replace **<172.17.5.13:/volumes/_nogroup/e840b4ae-6a04-49ee-9d6e-67d4999fbc01>** with the export location of the share.

- Retrieve the export location as described in [Section 7.8.1, “Listing share export locations”](#).
3. To mount a native CephFS share, refer to the following example for sample syntax:

```
$ mount -t ceph \
<192.168.1.7:6789,192.168.1.8:6789,192.168.1.9:6789:/volumes/_nogroup/4c55ad20-9c55-4a5e-9233-8ac64566b98c> \
-o name=<user-01>,secret='<AQA8+ANW/<4ZWNRAAOtWJMFPEihBA1unFImJczA==>'
```

- Replace **<192.168.1.7:6789,192.168.1.8:6789,192.168.1.9:6789:/volumes/_nogroup/4c55ad20-9c55-4a5e-9233-8ac64566b98c>** with the export location of the share.
- Retrieve the export location as described in [Section 7.8.1, “Listing share export locations”](#).
- Replace **<user-01>** with the cephx user who has access to the share.
- Replace the **secret** value with the access key that you collected in [Section 7.7.2, “Granting access to a native CephFS share”](#).

Verification

- Verify that the mount command succeeded:

```
$ df -k
```

7.9. DELETING SHARES

The Shared File Systems service (manila) provides no protections to prevent you from deleting your data. The Shared File Systems service does not check whether clients are connected or workloads are running. When you delete a share, you cannot retrieve it.



WARNING

Back up your data before you delete a share.

Prerequisites

- If you created snapshots from a share, you must delete all of the snapshots and replicas before you can delete the share. For more information, see [Deleting a snapshot](#).

Procedure

- Delete a share:

```
$ manila delete <share>
```

- Replace **<share>** with either the share name or the share ID.

7.10. LISTING RESOURCE LIMITS OF THE SHARED FILE SYSTEMS SERVICE

As a cloud user, you can list the current resource limits. This can help you plan workloads and prepare for any action based on your resource consumption.

Procedure

- List the resource limits and current resource consumption for the project:

```
$ manila absolute-limits
+-----+-----+
| Name                | Value |
+-----+-----+
| maxTotalReplicaGigabytes | 1000 |
| maxTotalShareGigabytes   | 1000 |
| maxTotalShareGroupSnapshots | 50  |
| maxTotalShareGroups      | 49  |
| maxTotalShareNetworks    | 10  |
| maxTotalShareReplicas    | 100 |
| maxTotalShareSnapshots   | 50  |
| maxTotalShares           | 50  |
| maxTotalSnapshotGigabytes | 1000 |
| totalReplicaGigabytesUsed | 22  |
| totalShareGigabytesUsed   | 25  |
| totalShareGroupSnapshotsUsed | 0   |
| totalShareGroupsUsed      | 9   |
| totalShareNetworksUsed    | 2   |
| totalShareReplicasUsed    | 9   |
| totalShareSnapshotsUsed   | 4   |
| totalSharesUsed           | 12  |
| totalSnapshotGigabytesUsed | 4   |
+-----+-----+
```

7.11. TROUBLESHOOTING OPERATION FAILURES

In the event that Shared File Systems (manila) operations, such as create share or create share group, fail asynchronously, as an end user, you can run queries from the command line for more information about the errors.

7.11.1. Fixing create share or create share group failures

In this example, the goal of the end user is to create a share to host software libraries on several virtual machines. The example deliberately introduces two share creation failures to illustrate how to use the command line to retrieve user support messages.

Procedure

- To create the share, you can use a share type that specifies some capabilities that you want the share to have. Cloud administrators can create share types. View the available share types:

```
clouduser1@client:~$ manila type-list
+-----+-----+-----+-----+
-----+-----+-----+-----+
```



```

| ID | Name | visibility | is_default | required_extra_specs |
optional_extra_specs | Description |
+-----+-----+-----+-----+-----+
-----+
| 1cf5d45a-61b3-44d1-8ec7-89a21f51a4d4 | dhss_false | public | YES |
driver_handles_share_servers : False | create_share_from_snapshot_support : True | None
|
| | | | |
mount_snapshot_support : False | | |
| | | | |
revert_to_snapshot_support : False | | |
| | | | |
| | | | | snapshot_support :
True | | |
| 277c1089-127f-426e-9b12-711845991ea1 | dhss_true | public | - |
driver_handles_share_servers : True | create_share_from_snapshot_support : True | None
|
| | | | |
mount_snapshot_support : False | | |
| | | | |
revert_to_snapshot_support : False | | |
| | | | |
| | | | | snapshot_support :
True | | |
+-----+-----+-----+-----+-----+
-----+

```

In this example, two share types are available.

- To use a share type that specifies the **driver_handles_share_servers=True** capability, you must create a share network on which to export the share. Create a share network from a private project network.

```

clouduser1@client:~$ openstack subnet list
+-----+-----+-----+-----+-----+
-----+
| ID | Name | Network | Subnet |
+-----+-----+-----+-----+-----+
-----+
| 78c6ac57-bba7-4922-ab81-16cde31c2d06 | private-subnet | 74d5cfb3-5dd0-43f7-b1b2-5b544cb16212 | 10.0.0.0/26 |
| a344682c-718d-4825-a87a-3622b4d3a771 | ipv6-private-subnet | 74d5cfb3-5dd0-43f7-b1b2-5b544cb16212 | fd36:18fc:a8e9::/64 |
+-----+-----+-----+-----+-----+
-----+

clouduser1@client:~$ manila share-network-create --name mynet --neutron-net-id 74d5cfb3-5dd0-43f7-b1b2-5b544cb16212 --neutron-subnet-id 78c6ac57-bba7-4922-ab81-16cde31c2d06
+-----+-----+-----+-----+-----+
| Property | Value |
+-----+-----+-----+-----+-----+
| network_type | None |
| name | mynet |
| segmentation_id | None |
| created_at | 2018-10-09T21:32:22.485399 |
| neutron_subnet_id | 78c6ac57-bba7-4922-ab81-16cde31c2d06 |
| updated_at | None |

```

```

| mtu          | None          |
| gateway      | None          |
| neutron_net_id | 74d5cfb3-5dd0-43f7-b1b2-5b544cb16212 |
| ip_version   | None          |
| cidr         | None          |
| project_id   | cadd7139bc3148b8973df097c0911016 |
| id           | 0b0fc320-d4b5-44a1-a1ae-800c56de550c |
| description  | None          |
+-----+-----+

```

```
clouduser1@client:~$ manila share-network-list
```

```

+-----+-----+
| id           | name         |
+-----+-----+
| 6c7ef9ef-3591-48b6-b18a-71a03059edd5 | mynet |
+-----+-----+

```

3. Create the share:

```
clouduser1@client:~$ manila create nfs 1 --name software_share --share-network mynet --share-type dhss_true
```

```

+-----+-----+
| Property          | Value          |
+-----+-----+
| status            | creating      |
| share_type_name   | dhss_true     |
| description       | None          |
| availability_zone | None          |
| share_network_id  | 6c7ef9ef-3591-48b6-b18a-71a03059edd5 |
| share_server_id   | None          |
| share_group_id    | None          |
| host              |               |
| revert_to_snapshot_support | False        |
| access_rules_status | active        |
| snapshot_id       | None          |
| create_share_from_snapshot_support | False        |
| is_public         | False         |
| task_state        | None          |
| snapshot_support  | False         |
| id                | 243f3a51-0624-4bdd-950e-7ed190b53b67 |
| size              | 1             |
| source_share_group_snapshot_member_id | None        |
| user_id           | 61aef4895b0b41619e67ae83fba6defe |
| name              | software_share |
| share_type        | 277c1089-127f-426e-9b12-711845991ea1 |
| has_replicas      | False         |
| replication_type  | None          |
| created_at        | 2018-10-09T21:12:21.000000 |
| share_proto       | NFS           |
| mount_snapshot_support | False        |
| project_id        | cadd7139bc3148b8973df097c0911016 |
| metadata          | {}            |
+-----+-----+

```

4. View the status of the share:

■

```

clouduser1@client:~$ manila list
+-----+-----+-----+-----+-----+-----+-----+
| ID          | Name          | Size | Share Proto | Status | Is Public | Share Type |
| Name | Host | Availability Zone |
+-----+-----+-----+-----+-----+-----+-----+
| 243f3a51-0624-4bdd-950e-7ed190b53b67 | software_share | 1 | NFS | error | False |
| dhss_true | | None |
+-----+-----+-----+-----+-----+-----+

```

In this example, an error occurred during the share creation.

- To view the user support message, run the **message-list** command. Use the **--resource-id** to filter to the specific share you want to find out about.

```

clouduser1@client:~$ manila message-list
+-----+-----+-----+-----+-----+-----+
| ID          | Resource Type | Resource ID          | Action ID | User |
| Message     |               |                     | Detail ID | Created At |
+-----+-----+-----+-----+-----+-----+
| 7d411c3c-46d9-433f-9e21-c04ca30b209c | SHARE | 243f3a51-0624-4bdd-950e-7ed190b53b67 | 001 | allocate host: No storage could be allocated for this share request, Capabilities filter didn't succeed. | 008 | 2018-10-09T21:12:21.000000 |
+-----+-----+-----+-----+-----+-----+

```

In the **User Message** column, notice that the Shared File Systems service failed to create the share because of a capabilities mismatch.

- To view more message information, run the **message-show** command, followed by the ID of the message from the **message-list** command:

```

clouduser1@client:~$ manila message-show 7d411c3c-46d9-433f-9e21-c04ca30b209c
+-----+-----+-----+-----+-----+
| Property | Value |
+-----+-----+-----+-----+-----+
| request_id | req-0a875292-6c52-458b-87d4-1f945556feac |
| detail_id | 008 |
| expires_at | 2018-11-08T21:12:21.000000 |
| resource_id | 243f3a51-0624-4bdd-950e-7ed190b53b67 |
| user_message | allocate host: No storage could be allocated for this share request, Capabilities filter didn't succeed. |
| created_at | 2018-10-09T21:12:21.000000 |

```

```

|
| message_level | ERROR
| id           | 7d411c3c-46d9-433f-9e21-c04ca30b209c
|
| resource_type | SHARE
| action_id    | 001
+-----+
-----+

```

7. As the cloud user, you can check capabilities through the share type so you can review the share types available. The difference between the two share types is the value of **driver_handles_share_servers**:

```

clouduser1@client:~$ manila type-list
+-----+-----+-----+-----+-----+
-----+-----+-----+-----+-----+
| ID                | Name      | visibility | is_default | required_extra_specs |
| optional_extra_specs | Description |
+-----+-----+-----+-----+-----+
-----+-----+-----+-----+-----+
| 1cf5d45a-61b3-44d1-8ec7-89a21f51a4d4 | dhss_false | public    | YES       |
| driver_handles_share_servers : False | create_share_from_snapshot_support : True | None
|
| mount_snapshot_support : False
|
| revert_to_snapshot_support : False
|
| True | snapshot_support :
|
| 277c1089-127f-426e-9b12-711845991ea1 | dhss_true | public    | -        |
| driver_handles_share_servers : True | create_share_from_snapshot_support : True | None
|
| mount_snapshot_support : False
|
| revert_to_snapshot_support : False
|
| True | snapshot_support :
+-----+-----+-----+-----+-----+
-----+-----+-----+-----+-----+

```

8. Create a share with the other available share type:

```

clouduser1@client:~$ manila create nfs 1 --name software_share --share-network mynet --
share-type dhss_false
+-----+-----+-----+-----+-----+
| Property          | Value
+-----+-----+-----+-----+-----+
| status            | creating
| share_type_name   | dhss_false
| description       | None
| availability_zone | None
| share_network_id  | 6c7ef9ef-3591-48b6-b18a-71a03059edd5 |
| share_group_id    | None
| revert_to_snapshot_support | False

```

```

| access_rules_status      | active          |
| snapshot_id             | None           |
| create_share_from_snapshot_support | True          |
| is_public               | False         |
| task_state              | None          |
| snapshot_support        | True          |
| id                      | 2d03d480-7cba-4122-ac9d-edc59c8df698 |
| size                    | 1             |
| source_share_group_snapshot_member_id | None         |
| user_id                 | 5c7bdb6eb0504d54a619acf8375c08ce |
| name                    | software_share |
| share_type              | 1cf5d45a-61b3-44d1-8ec7-89a21f51a4d4 |
| has_replicas            | False         |
| replication_type        | None          |
| created_at              | 2018-10-09T21:24:40.000000 |
| share_proto             | NFS           |
| mount_snapshot_support | False         |
| project_id              | cadd7139bc3148b8973df097c0911016 |
| metadata                 | {}            |
+-----+-----+

```

In this example, the second share creation attempt fails.

9. View the user support message:

```

clouduser1@client:~$ manila list
+-----+-----+-----+-----+-----+-----+-----+
-----+-----+-----+
| ID          | Name          | Size | Share Proto | Status | Is Public | Share Type
Name | Host | Availability Zone |
+-----+-----+-----+-----+-----+-----+-----+
-----+-----+-----+
| 2d03d480-7cba-4122-ac9d-edc59c8df698 | software_share | 1 | NFS | error | False
| dhss_false | nova |
| 243f3a51-0624-4bdd-950e-7ed190b53b67 | software_share | 1 | NFS | error | False
| dhss_true | None |
+-----+-----+-----+-----+-----+-----+-----+
-----+-----+-----+

clouduser1@client:~$ manila message-list
+-----+-----+-----+-----+-----+-----+-----+
-----+-----+-----+
---+-----+-----+
| ID          | Resource Type | Resource ID          | Action ID | User
Message          | Detail ID | Created At
|
+-----+-----+-----+-----+-----+-----+-----+
-----+-----+-----+
| ed7e02a2-0cdb-4ff9-b64f-e4d2ec1ef069 | SHARE | 2d03d480-7cba-4122-ac9d-
edc59c8df698 | 002 | create: Driver does not expect share-network to be provided with
current configuration. | 003 | 2018-10-09T21:24:40.000000 |
| 7d411c3c-46d9-433f-9e21-c04ca30b209c | SHARE | 243f3a51-0624-4bdd-950e-
7ed190b53b67 | 001 | allocate host: No storage could be allocated for this share request,
Capabilities filter didn't succeed. | 008 | 2018-10-09T21:12:21.000000 |

```

```

+-----+-----+-----+-----+
-----+-----
---+-----+

```

The service does not expect a share network for the share type that you used.

- Without consulting the administrator, you can discover that the administrator has not made available a storage back end that supports exporting shares directly on to your private neutron network. Create the share without the **share-network** parameter:

```
clouduser1@client:~$ manila create nfs 1 --name software_share --share-type dhss_false
```

```

+-----+-----+-----+-----+
| Property          | Value          |
+-----+-----+-----+-----+
| status            | creating       |
| share_type_name   | dhss_false     |
| description       | None           |
| availability_zone | None           |
| share_network_id  | None           |
| share_group_id    | None           |
| revert_to_snapshot_support | False         |
| access_rules_status | active         |
| snapshot_id       | None           |
| create_share_from_snapshot_support | True         |
| is_public         | False          |
| task_state        | None           |
| snapshot_support  | True           |
| id                | 4d3d7fcf-5fb7-4209-90eb-9e064659f46d |
| size              | 1              |
| source_share_group_snapshot_member_id | None         |
| user_id           | 5c7bdb6eb0504d54a619acf8375c08ce |
| name              | software_share |
| share_type        | 1cf5d45a-61b3-44d1-8ec7-89a21f51a4d4 |
| has_replicas      | False          |
| replication_type  | None           |
| created_at        | 2018-10-09T21:25:40.000000 |
| share_proto       | NFS            |
| mount_snapshot_support | False         |
| project_id        | cadd7139bc3148b8973df097c0911016 |
| metadata          | {}             |
+-----+-----+-----+-----+

```

- Ensure that the share was created successfully:

```
clouduser1@client:~$ manila list
```

```

+-----+-----+-----+-----+-----+-----+-----+-----+
-----+-----+-----+-----+-----+-----+-----+-----+
| ID              | Name          | Size | Share Proto | Status | Is Public | Share Type |
| Name | Host | Availability Zone |
+-----+-----+-----+-----+-----+-----+-----+-----+
-----+-----+-----+-----+-----+-----+-----+-----+
| 4d3d7fcf-5fb7-4209-90eb-9e064659f46d | software_share | 1 | NFS | available |
False | dhss_false | | nova |
| 2d03d480-7cba-4122-ac9d-edc59c8df698 | software_share | 1 | NFS | error | False |
| dhss_false | | nova |
| 243f3a51-0624-4bdd-950e-7ed190b53b67 | software_share | 1 | NFS | error |

```

```
False | dhss_true | | None |
+-----+-----+-----+-----+-----+-----+-----+
-----+-----+-----+-----+-----+-----+-----+
```

12. Delete the shares and support messages:

```
clouduser1@client:~$ manila message-list
+-----+-----+-----+-----+-----+-----+-----+
-----+-----+-----+-----+-----+-----+-----+
---+-----+
| ID | Resource Type | Resource ID | Action ID | User
Message | Detail ID | Created At
|
+-----+-----+-----+-----+-----+-----+-----+
-----+-----+-----+-----+-----+-----+-----+
---+-----+
| ed7e02a2-0cdb-4ff9-b64f-e4d2ec1ef069 | SHARE | 2d03d480-7cba-4122-ac9d-
edc59c8df698 | 002 | create: Driver does not expect share-network to be provided with
current configuration. | 003 | 2018-10-09T21:24:40.000000 |
| 7d411c3c-46d9-433f-9e21-c04ca30b209c | SHARE | 243f3a51-0624-4bdd-950e-
7ed190b53b67 | 001 | allocate host: No storage could be allocated for this share request,
Capabilities filter didn't succeed. | 008 | 2018-10-09T21:12:21.000000 |
+-----+-----+-----+-----+-----+-----+-----+
-----+-----+-----+-----+-----+-----+-----+
---+-----+

clouduser1@client:~$ manila delete 2d03d480-7cba-4122-ac9d-edc59c8df698 243f3a51-
0624-4bdd-950e-7ed190b53b67
clouduser1@client:~$ manila message-delete ed7e02a2-0cdb-4ff9-b64f-e4d2ec1ef069
7d411c3c-46d9-433f-9e21-c04ca30b209c

clouduser1@client:~$ manila message-list
+-----+-----+-----+-----+-----+-----+-----+
| ID | Resource Type | Resource ID | Action ID | User Message | Detail ID | Created At |
+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+
```

7.11.2. Debugging share mounting failures

If you experience an issue when you mount shares, use these verification steps to identify the root cause.

Procedure

1. Verify the access control list of the share to ensure that the rule that corresponds to your client is correct and has been successfully applied.

```
$ manila access-list share-01
```

In a successful rule, the **state** attribute equals **active**.

2. If the share type parameter is configured to **driver_handles_share_servers=False**, copy the hostname or IP address from the export location and ping it to confirm connectivity to the NAS server:

■

```
$ ping -c 1 172.17.5.13
PING 172.17.5.13 (172.17.5.13) 56(84) bytes of data.
64 bytes from 172.17.5.13: icmp_seq=1 ttl=64 time=0.048 ms--- 172.17.5.13 ping statistics --
-
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 7.851/7.851/7.851/0.000 ms
If using the NFS protocol, you may verify that the NFS server is ready to respond to NFS rpcs
on the proper port:
$ rpcinfo -T tcp -a 172.17.5.13.8.1 100003 4
program 100003 version 4 ready and waiting
```

**NOTE**

The IP address is written in universal address format (uaddr), which adds two extra octets (8.1) to represent the NFS service port, 2049.

If these verification steps fail, there might be a network connectivity issue, or your shared file system back-end storage has failed. Collect the log files and contact Red Hat Support.