



Red Hat OpenStack Platform 16.1

Storage Guide

Understanding, using, and managing persistent storage in OpenStack

Red Hat OpenStack Platform 16.1 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	7
1.1. SCALABILITY AND BACK END	8
1.2. ACCESSIBILITY AND ADMINISTRATION	8
1.3. SECURITY	9
1.4. REDUNDANCY AND DISASTER RECOVERY	9
CHAPTER 2. BLOCK STORAGE AND VOLUMES	11
2.1. BACK ENDS	11
2.1.1. Third-Party Storage Providers	11
2.2. BLOCK STORAGE SERVICE ADMINISTRATION	11
2.2.1. Active-active deployment for high availability	11
2.2.1.1. Enabling active-active configuration for high availability	12
2.2.1.2. Maintenance commands for active-active configurations	12
2.2.1.3. Volume manage and unmanage	13
2.2.1.4. Volume migration on a clustered service	13
2.2.1.5. Initiating server maintenance	13
2.2.2. Group Volume Settings with Volume Types	14
2.2.2.1. List the Capabilities of a Host Driver	15
2.2.2.2. Create and Configure a Volume Type	16
2.2.2.3. Edit a Volume Type	16
2.2.2.4. Delete a Volume Type	17
2.2.2.5. Create and Configure Private Volume Types	17
2.2.3. Create and Configure an Internal Tenant for the Block Storage service	17
2.2.4. Configure and Enable the Image-Volume Cache	18
2.2.5. Use Quality-of-Service Specifications	19
2.2.5.1. Basic volume Quality of Service	20
2.2.5.2. Create and Configure a QOS Spec	20
2.2.5.3. Set Capacity-Derived QoS Limits	21
2.2.5.4. Associate a QOS Spec with a Volume Type	21
2.2.5.5. Disassociate a QOS Spec from a Volume Type	22
2.2.6. Configure volume encryption	22
2.2.6.1. Configure Volume Type Encryption Through the Dashboard	22
2.2.6.2. Configure Volume Type Encryption Through the CLI	23
2.2.6.3. Automatic deletion of volume image encryption key	23
2.2.7. Configure How Volumes are Allocated to Multiple Back Ends	24
2.2.8. Deploying availability zones	25
2.2.9. Configure and use consistency groups	25
2.2.9.1. Configure consistency groups	26
2.2.9.2. Creating consistency groups	28
2.2.9.3. Managing consistency groups	28
2.2.9.4. Create and manage consistency group snapshots	28
2.2.9.5. Clone consistency groups	29
2.3. BASIC VOLUME USAGE AND CONFIGURATION	29
2.3.1. Create a volume	30
2.3.2. Specify back end for volume creation	31
2.3.3. Edit a volume name or description	31
2.3.4. Resize (extend) a volume	32
2.3.5. Delete a volume	32

2.3.6. Attach and detach a volume to an instance	32
2.3.6.1. Attaching a volume to an instance	32
2.3.6.2. Detaching a volume from an instance	33
2.3.7. Attach a volume to multiple instances	33
2.3.7.1. Creating a multi-attach volume type	34
2.3.7.2. Volume retyping	34
2.3.7.3. Creating a multi-attach volume	34
2.3.7.4. Supported back ends	35
2.3.8. Read-only volumes	35
2.3.9. Change a volume owner	35
2.3.9.1. Transfer a volume from the command line	35
2.3.9.2. Transfer a volume by using the dashboard	36
2.3.10. Create, use, or delete volume snapshots	37
2.3.10.1. Protected and unprotected snapshots in a Red Hat Ceph Storage back end	38
2.3.11. Use a snapshot to restore to the last state of a volume	38
2.3.11.1. Verifying that your revert is successful	39
2.3.12. Upload a volume to the Image Service	40
2.3.13. Moving volumes between back ends	40
2.3.14. Moving available volumes	40
2.3.14.1. Generic volume migration	41
2.3.15. Moving in-use volumes	41
2.3.16. Volume retyping	41
2.3.16.1. Retyping a volume from the dashboard UI	42
2.3.16.2. Retyping a volume from the command line	43
2.3.17. Enabling LVM2 filtering on overcloud nodes	44
2.4. ADVANCED VOLUME CONFIGURATION	45
2.4.1. Migrate a Volume	45
2.4.1.1. Migrate between back ends	46
2.4.1.2. Migrating between back ends from the command line	46
2.4.1.3. Verifying volume migration	47
2.4.2. Encrypting unencrypted volumes	48
2.5. MULTIPATH CONFIGURATION	49
2.5.1. Configuring multipath on new deployments	50
2.5.2. Configuring multipath on existing deployments	52
2.5.3. Verifying multipath configuration	54
CHAPTER 3. OBJECT STORAGE SERVICE	57
3.1. OBJECT STORAGE RINGS	57
3.1.1. Rebalancing rings	57
3.1.2. Checking cluster health	57
3.1.3. Increasing ring partition power	59
3.1.4. Creating custom rings	59
3.2. OBJECT STORAGE SERVICE ADMINISTRATION	59
3.2.1. Configuring fast-post	59
3.2.2. Enabling at-rest encryption	59
3.2.3. Deploying a standalone Object Storage cluster	60
3.2.3.1. Creating the roles_data.yaml File	60
3.2.3.2. Deploying the New Roles	62
3.2.4. Using external SAN disks	62
3.2.4.1. SAN disk deployment configuration	63
3.3. INSTALL AND CONFIGURE STORAGE NODES FOR RED HAT ENTERPRISE LINUX	63
3.3.1. Preparing storage devices	64
3.3.2. Configuring components	65

3.4. BASIC CONTAINER MANAGEMENT	67
3.4.1. Creating a container	67
3.4.2. Creating a pseudo folder for a container	68
3.4.3. Deleting a container	68
3.4.4. Uploading an object	68
3.4.5. Copying an object	69
3.4.6. Deleting an object	69
CHAPTER 4. SHARED FILE SYSTEMS SERVICE	70
4.1. SHARED FILE SYSTEMS SERVICE (MANILA) BACK ENDS	71
4.1.1. Networking for shared file systems	71
4.1.2. Creating a share type	72
4.1.3. Common capabilities of share types	73
4.1.4. Discovering share types	74
4.1.5. Creating a share	74
4.1.6. Listing shares and exporting information	75
4.2. ENSURING NETWORK CONNECTIVITY TO THE SHARE	75
4.2.1. Connecting to a shared network to access shares	76
4.2.2. Configuring an IPv6 interface between the network and an instance	77
4.3. GRANT SHARE ACCESS	78
4.3.1. Granting access to a share	79
4.3.2. Revoking access to a share	80
4.4. MOUNT SHARE ON COMPUTE INSTANCES	81
4.4.1. Listing shares export locations	81
4.4.2. Mounting the share	81
4.5. DELETING A SHARE	82
4.6. CHANGE THE DEFAULT QUOTAS IN THE SHARED FILE SYSTEMS SERVICE	82
4.6.1. Listing quotas	82
4.7. TROUBLESHOOTING FAILURES	82
4.7.1. Fixing create share or create share group failures	83
4.7.2. Debugging share mounting failures	90

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

This guide discusses procedures for creating and managing persistent storage. Within Red Hat OpenStack Platform (RHOSP), this 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 tenants 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

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. 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 can be used for cloning or restoring a volume to a previous state (see [Section 1.4, "Redundancy and Disaster Recovery"](#)). The Block Storage service also allows you to create *volume types*, which aggregate volume settings (for example, size and

back end) that can be easily invoked by users when creating new volumes. These types can be further associated with *Quality-of-Service* specifications, which allow you to create different storage tiers for users.

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. 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 2.2.6, "Configure 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 [Image service](#) in the *Creating and Managing Images* guide.

1.4. REDUNDANCY AND DISASTER RECOVERY

The Block Storage service features volume backup and restoration, providing basic disaster recovery for user storage. Backups allow you to protect volume contents. On top of this, the service also supports snapshots; aside from cloning, snapshots are also useful in restoring a volume to a previous state.

In a multi-backend 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, however, with snapshots, as snapshots are dependent on their source volumes.

The Block Storage service also supports the creation of *consistency groups*, which allow you to group volumes together for simultaneous snapshot creation. This, in turn, allows for a greater level of data consistency across multiple volumes. See [Section 2.2.9, "Configure and use consistency groups"](#) for more details.

The Object Storage service provides no built-in backup features. As such, all backups must be performed at the file system or node level. The service, however, 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. BLOCK STORAGE AND VOLUMES

The Block Storage service (**openstack-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, refer to the [Block Storage Backup Guide](#).

2.1. BACK ENDS

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

Red Hat OpenStack Platform supports [Red Hat Ceph](#) and NFS as Block Storage 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 Ceph with OpenStack, see [Deploying an Overcloud with Containerized Red Hat Ceph](#).

For instructions on how to set up NFS storage in the overcloud, see [Configuring NFS Storage](#) in the [Advanced Overcloud Customization Guide](#).

2.1.1. Third-Party Storage Providers

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

For a complete list of supported back end appliances and drivers, see [Component, Plug-In, and Driver Support in RHEL 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. For more information about deploying a third-party storage appliance plugin, see [Deploying a vendor plugin](#) in the [Advanced Overcloud Customization](#) guide.

2.2. BLOCK STORAGE SERVICE ADMINISTRATION

The following procedures explain how to configure the Block Storage service to suit your needs. All of these procedures require administrator privileges.



IMPORTANT

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

2.2.1. Active-active deployment for high availability

Important

Active-active mode is supported only in distributed compute node (DCN) architecture at edge sites.

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.

2.2.1.1. Enabling active-active configuration for high availability

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.

Procedure

To enable active-active Block Storage service volumes, include the following environment file in your overcloud deployment:

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

2.2.1.2. Maintenance commands for active-active configurations

After deploying an active-active configuration, there are several commands you can use to interact with the environment when using API version 3.17 and later.

User goal	Command
-----------	---------

See the service listing, including details such as cluster name, host, zone, status, state, disabled reason, and back end state. NOTE: When deployed by director for the Ceph back end, the default cluster name is tripleo@tripleo_ceph .	cinder service-list
See detailed and summary information about clusters as a whole as opposed to individual services.	cinder cluster-list NOTE: This command requires a cinder API microversion of 3.7 or later.
See detailed information about a specific cluster.	cinder cluster-show <cluster_name> NOTE: This command requires a cinder API microversion of 3.7 or later.
Enable a disabled service.	cinder cluster-enable <cluster_name> NOTE: This command requires a cinder API microversion of 3.7 or later.
Disable a clustered service.	cinder cluster-disable <cluster_name> NOTE: This command requires a cinder API microversion of 3.7 or later.

2.2.1.3. 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 API version 3.17 or later, you can see lists of volumes and snapshots that are available for management in Block Storage clusters. To see these lists, use the **--cluster** argument with **cinder manageable-list** or **cinder snapshot-manageable-list**.

In API version 3.16 and later, the **cinder manage** command also accepts the optional **--cluster** argument so that you can add previously unmanaged volumes to a Block Storage cluster.

2.2.1.4. Volume migration on a clustered service

With API version 3.16 and later, the **cinder migrate** and **cinder-manage** commands accept the **--cluster** argument to define the destination for active-active deployments.

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

2.2.1.5. Initiating server 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.



NOTE

The **work-cleanup** command works only on servers running API version 3.24 or later.

Procedure

1. 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.

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

```
$ cinder service-list
```

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

```
$ cinder 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 work-cleanup --cluster tripleo@tripleo_ceph
```

2.2.2. Group Volume Settings with Volume Types

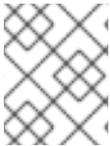
With Red Hat OpenStack Platform you can create volume types so that you can apply associated settings to the volume type. You can apply settings during volume creation, see [Create a Volume](#). You can also apply settings after you create a volume, see [Changing the Type of a Volume \(Volume Re-typing\)](#). The following list shows some of the associated setting that you can apply to a volume type:

- The encryption of a volume. For more information, see [Configure Volume Type Encryption](#).
- The back end that a volume uses. For more information, see [Specify Back End for Volume Creation](#) and [Migrate a Volume](#).
- Quality-of-Service (QoS) Specs

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.

Volume types provide the capability to provide different users with storage tiers. By associating specific performance, resilience, and other settings as key-value pairs to a volume type, you can map tier-specific settings to different volume types. You can then apply tier settings when creating a volume by specifying the corresponding volume type.

2.2.2.1. List the Capabilities of a Host Driver



NOTE

Available and supported Extra Specs vary per back end driver. Consult the driver documentation for a list of valid Extra Specs.

Alternatively, you can query the Block Storage host directly to determine which well-defined standard Extra Specs are supported by its driver. Start by logging in (through the command line) to the node hosting the Block Storage service. Then:

```
# 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 ...
+-----+-----+-----+-----+
```

To display the driver capabilities (and, in turn, determine the supported Extra Specs) of a Block Storage service, run:

```
# cinder get-capabilities _VOLSVCHOST_
```

Where *VOLSVCHOST* is the complete name of the **cinder-volume**'s host. 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.

2.2.2.2. Create and Configure a Volume Type

1. As an admin user in the dashboard, select **Admin > Volumes > Volume Types**
2. Click **Create Volume Type**.
3. Enter the volume type name in the **Name** field.
4. Click **Create Volume Type**. The new type appears in the **Volume Types** table.
5. Select the volume type's **View Extra Specs** action.
6. 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.
7. 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 tenants. If you need to create volume types with restricted access, you will need to do so through the CLI. For instructions, see [Section 2.2.2.5, "Create and Configure Private Volume Types"](#).



NOTE

You can also associate a QoS Spec to the volume type. For more information, see [Section 2.2.5.4, "Associate a QOS Spec with a Volume Type"](#).

2.2.2.3. Edit a Volume Type

1. As an admin user in the dashboard, select **Admin > Volumes > Volume Types**
2. In the **Volume Types** table, select the volume type's **View Extra Specs** action.
3. 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.

2.2.2.4. Delete a Volume Type

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

2.2.2.5. Create and Configure Private Volume Types

By default, all volume types are available to all tenants. You can create a restricted volume type by marking it **private**. To do so, set the type's **is-public** flag to **false**.

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

To create a private volume type, run:

```
$ cinder type-create --is-public false <TYPE-NAME>
```

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

```
$ cinder type-list --all
```

This command lists both public and private volume types, and it also includes the name and ID of each one. You need the volume type's ID to provide access to it.

Access to a private volume type is granted at the tenant level. To grant a tenant access to a private volume type, run:

```
$ cinder type-access-add --volume-type <TYPE-ID> --project-id <TENANT-ID>
```

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

```
$ cinder type-access-list --volume-type <TYPE-ID>
```

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

```
$ cinder type-access-remove --volume-type <TYPE-ID> --project-id <TENANT-ID>
```



NOTE

By default, only users with administrative privileges can create, view, or configure access for private volume types.

2.2.3. Create and Configure an Internal Tenant for the Block Storage service

Some Block Storage features (for example, the Image-Volume cache) require the configuration of an internal tenant. The Block Storage service uses this tenant (or 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.

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 Tenant" 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      |
+-----+-----+
```

The procedure for adding Extra Config options creates an internal tenant. For more information, see [Section 2.2.4, "Configure and Enable the Image-Volume Cache"](#).

2.2.4. Configure and Enable 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 [Section 2.3.1, "Create a volume"](#).

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 Image-Volume cache's limit 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.



NOTE

The Image-Volume cache requires that an *internal tenant* be configured for the Block Storage service. For instructions, see [Section 2.2.3, "Create and Configure an Internal Tenant for the Block Storage service"](#).

To enable and configure the Image-Volume cache on a back end (*BACKEND*), add the values to an **ExtraConfig** section of an environment file on the undercloud. For example:

```
parameter_defaults:
  ExtraConfig:
    cinder::config::cinder_config:
      DEFAULT/cinder_internal_tenant_project_id:
        value: TENANTID
      DEFAULT/cinder_internal_tenant_user_id:
```

```

value: USERID
BACKEND/image_volume_cache_enabled: ❶
value: True
BACKEND/image_volume_cache_max_size_gb:
value: MAXSIZE ❷
BACKEND/image_volume_cache_max_count:
value: MAXNUMBER ❸

```

- ❶ Replace *BACKEND* with the name of the target back end (specifically, its **volume_backend_name** value).
- ❷ By default, the Image-Volume cache size is only limited by the back end. Change *MAXSIZE* to a number in GB.
- ❸ You can also set a maximum number of images using *MAXNUMBER*.

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.

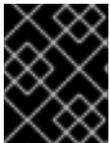
After you create the environment file in **/home/stack/templates/**, log in as the stack user and deploy the configuration by running:

```

$ openstack overcloud deploy --templates \
-e /home/stack/templates/<ENV_FILE>.yaml

```

Where **ENV_FILE.yaml** is the name of the file with the **ExtraConfig** settings added earlier.



IMPORTANT

If you passed any extra environment files when you created the overcloud, pass them again here using the **-e** option to avoid making undesired changes to the overcloud.

For additional information on the **openstack overcloud deploy** command, see [Deployment command](#) in the *Director Installation and Usage Guide*.

2.2.5. Use Quality-of-Service Specifications

You can map multiple performance settings to a single Quality-of-Service specification (QOS Specs). Doing so allows you to provide performance tiers for different user types.

Performance settings are mapped as key-value pairs to QOS Specs, similar to the way volume settings are associated to a volume type. However, QOS Specs are different from volume types in the following respects:

- QOS Specs are used to apply performance settings, which include limiting read/write operations to disks. Available and supported performance settings vary per storage driver. To determine which QOS Specs are supported by your back end, consult the documentation of your back end device's volume driver.
- Volume types are directly applied to volumes, whereas QOS Specs are not. Rather, QOS Specs are associated to volume types. During volume creation, specifying a volume type also applies the performance settings mapped to the volume type's associated QOS Specs.

2.2.5.1. Basic volume Quality of Service

You can define performance limits for volumes on a per-volume basis using basic volume QOS values. The Block Storage service supports the following options:

- **read_iops_sec**
- **write_iops_sec**
- **total_iops_sec**
- **read_bytes_sec**
- **write_bytes_sec**
- **total_bytes_sec**
- **read_iops_sec_max**
- **write_iops_sec_max**
- **total_iops_sec_max**
- **read_bytes_sec_max**
- **write_bytes_sec_max**
- **total_bytes_sec_max**
- **size_iops_sec**

2.2.5.2. Create and Configure a QOS Spec

As an administrator, you can create and configure a QOS Spec through the QOS Specs table. You can associate more than one key/value pair to the same QOS Spec.

1. As an admin user in the dashboard, select **Admin > Volumes > Volume Types**
2. On the **QOS Specs** table, click **Create QOS Spec**.
3. Enter a name for the **QOS Spec**.
4. In the **Consumer** field, specify where the QOS policy should be enforced:

Table 2.1. Consumer Types

Type	Description
back-end	QOS policy will be applied to the Block Storage back end.
front-end	QOS policy will be applied to Compute.
both	QOS policy will be applied to both Block Storage and Compute.

5. Click **Create**. The new QOS Spec should now appear in the **QOS Specs** table.

6. In the **QOS Specs** table, select the new spec's **Manage Specs** action.
7. Click **Create**, and specify the **Key** and **Value**. The key-value pair must be valid; otherwise, specifying a volume type associated with this QOS Spec during volume creation will fail. For example, to set read limit IOPS to **500**, use the following Key/Value pair:

```
read_iops_sec=500
```

8. Click **Create**. The associated setting (key-value pair) now appears in the **Key-Value Pairs** table.

2.2.5.3. Set Capacity-Derived QoS Limits

You can use volume types to implement capacity-derived Quality-of-Service (QoS) limits on volumes. This will allow you to set a deterministic IOPS throughput based on the size of provisioned volumes. Doing this simplifies how storage resources are provided to users – namely, providing a user with pre-determined (and, ultimately, highly predictable) throughput rates based on the volume size they provision.

In particular, the Block Storage service allows you to set how much IOPS to allocate to a volume based on the actual provisioned size. This throughput is set on an IOPS per GB basis through the following QoS keys:

```
read_iops_sec_per_gb
write_iops_sec_per_gb
total_iops_sec_per_gb
```

These keys allow you to set read, write, or total IOPS to scale with the size of provisioned volumes. For example, if the volume type uses **read_iops_sec_per_gb=500**, then a provisioned 3GB volume would automatically have a read IOPS of 1500.

Capacity-derived QoS limits are set per volume type, and configured like any normal QoS spec. In addition, these limits are supported by the underlying Block Storage service directly, and is not dependent on any particular driver.

For more information about volume types, see [Section 2.2.2, "Group Volume Settings with Volume Types"](#) and [Section 2.2.2.2, "Create and Configure a Volume Type"](#). For instructions on how to set QoS specs, [Section 2.2.5, "Use Quality-of-Service Specifications"](#).



WARNING

When you apply a volume type (or perform a volume re-type) with capacity-derived QoS limits to an attached volume, the limits will not be applied. The limits will only be applied once you detach the volume from its instance.

See [Section 2.3.16, "Volume re-typing"](#) for information about volume re-typing.

2.2.5.4. Associate a QOS Spec with a Volume Type

As an administrator, you can associate a QOS Spec to an existing volume type using the **Volume Types** table.

1. As an administrator in the dashboard, select **Admin > Volumes > Volume Types**
2. In the **Volume Types** table, select the type's **Manage QOS Spec Association** action.
3. Select a QOS Spec from the **QOS Spec to be associated** list.
4. Click **Associate**. The selected QOS Spec now appears in the **Associated QOS Spec** column of the edited volume type.

2.2.5.5. Disassociate a QOS Spec from a Volume Type

1. As an administrator in the dashboard, select **Admin > Volumes > Volume Types**
2. In the **Volume Types** table, select the type's **Manage QOS Spec Association** action.
3. Select **None** from the QOS Spec to be associated list.
4. Click **Associate**. The selected QOS Spec is no longer in the **Associated QOS Spec** column of the edited volume type.

2.2.6. Configure volume encryption

Volume encryption provides basic data protection in case the volume back end is either compromised or stolen. Both Compute and Block Storage services are integrated to allow instances to read access and use encrypted volumes. You must deploy the Key Manager service (barbican) to use 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 by using volume type. See [Section 2.2.6.1, "Configure Volume Type Encryption Through the Dashboard"](#) for information about encrypted volume types.

2.2.6.1. Configure Volume Type Encryption Through 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:

1. As an admin user in the dashboard, select **Admin > Volumes > Volume Types**
2. In the **Actions** column of the volume to be encrypted, select **Create Encryption** to launch the **Create Volume Type Encryption** wizard.
3. 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**.
4. Click **Create Volume Type Encryption**

Once you have an encrypted volume type, you can invoke it to automatically create encrypted volumes. For more information on creating a volume type, see [Section 2.2.2.2, "Create and Configure a Volume Type"](#). Specifically, select the encrypted volume type from the Type drop-down list in the **Create Volume** window (see [Section 2.3, "Basic volume usage and configuration"](#)).

To configure an encrypted volume type through the CLI, see [Section 2.2.6.2, "Configure Volume Type Encryption Through the CLI"](#).

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.

2.2.6.2. Configure Volume Type Encryption Through the CLI

To configure Block Storage volume encryption, do the following:

1. Create a volume type:

```
$ cinder type-create encrypt-type
```

2. 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 encrypt-type luks
```

3. Create an encrypted volume:

```
$ cinder --debug create 1 --volume-type encrypt-type --name DemoEncVol
```

2.2.6.3. 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.

For additional information, refer to the [Manage secrets with the OpenStack Key Manager](#) guide.

2.2.7. Configure How Volumes are Allocated to Multiple Back Ends

If the Block Storage service is configured to use multiple back ends, you can use configured volume types to specify where a volume should be created. For details, see [Section 2.3.2, "Specify back end for volume creation"](#).

The Block Storage service will automatically choose a back end if you do not specify one during volume creation. Block Storage sets the first defined back end as a default; this back end will be used until it runs out of space. At that point, Block Storage will set the second defined back end as a default, and so on.

If this is not suitable for your needs, you can use the filter scheduler to control how Block Storage should select back ends. This scheduler can use different filters to triage suitable back ends, such as:

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.

To configure the filter scheduler, add an environment file to your deployment containing:

```
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.

2.2.8. Deploying availability zones

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.

Procedure

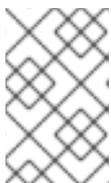
To deploy different availability zones for Block Storage volume back ends:

1. 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
```

2. Deploy the overcloud and include the updated environment file.

2.2.9. Configure and use 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.

2.2.9.1. Configure 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.

Procedure

1. 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.
2. 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' }, \
}
```

- 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.
- 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",
```

- 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' }, \
}
```

- When you have created the environment file in `/home/stack/templates/`, log in as the stack user and deploy the configuration:

```
$ openstack overcloud deploy --templates \
-e /home/stack/templates/<ENV_FILE>.yaml
```

Replace `<ENV_FILE.yaml>` with the name of the file with the **ExtraConfig** settings you added.



IMPORTANT

If you passed any extra environment files when you created the overcloud, pass them again here by using the **-e** option to avoid making undesired changes to the overcloud.

For more information about the **openstack overcloud deploy** command, see [Creating the Overcloud with the CLI Tools](#) in the *Director Installation and Usage* guide.

2.2.9.2. Creating consistency groups

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

Procedure

1. As an admin user in the dashboard, select **Project > Compute > Volumes > Volume Consistency Groups**.
2. Click **Create Consistency Group**.
3. In the **Consistency Group Information** tab of the wizard, enter a name and description for your consistency group. Then, specify its **Availability Zone**.
4. 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.
5. Click **Create Consistency Group**. It appears next in the **Volume Consistency Groups** table.

2.2.9.3. Managing consistency groups

Procedure

1. Optional: You can change the name or description of a consistency group by selecting **Edit Consistency Group** from its **Action** column.
2. To add or remove volumes from a consistency group directly, as an admin user in the dashboard, select **Project > Compute > Volumes > Volume Consistency Groups**
3. 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.
4. Click **Edit Consistency Group**.

2.2.9.4. Create and manage consistency group snapshots

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

Procedure

1. Log in as **admin** user from the command line on the node that hosts the **openstack-cinder-api** and enter:

```
# export OS_VOLUME_API_VERSION=2
```

This configures the client to use version **2** of the **openstack-cinder-api**.

- List all available consistency groups and their respective IDs:

```
# cinder consisgroup-list
```

- Create snapshots using the consistency group:

```
# cinder cgsnapshot-create --name <CGSNAPNAME> --description "<DESCRIPTION>"
<CGNAMEID>
```

Replace:

- **<CGSNAPNAME>** with the name of the snapshot (optional).
- **<DESCRIPTION>** with a description of the snapshot (optional).
- **<CGNAMEID>** with the name or ID of the consistency group.

- Display a list of all available consistency group snapshots:

```
# cinder cgsnapshot-list
```

2.2.9.5. Clone 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.

Procedure

- To clone an existing consistency group:

```
# cinder consisgroup-create-from-src --source-cg <CGNAMEID> --name <CGNAME> --
description "<DESCRIPTION>"
```

Replace:

- **<CGNAMEID>** is the name or ID of the consistency group you want to clone.
- **<CGNAME>** is the name of your consistency group (optional).
- **<DESCRIPTION>** is a description of your consistency group (optional).

- 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.

2.3. BASIC VOLUME USAGE AND CONFIGURATION

The following procedures describe how to perform basic end-user volume management. These procedures do not require administrative privileges.

**IMPORTANT**

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

2.3.1. Create a volume**IMPORTANT**

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

Procedure

1. In the dashboard, select **Project > Compute > Volumes**
2. Click **Create Volume**, and edit the following fields:

Field	Description
Volume name	Name of the volume.
Description	Optional, short description of the volume.
Type	Optional volume type (see Section 2.2.2, "Group Volume Settings with Volume Types"). If you have multiple Block Storage back ends, you can use this to select a specific back end. See Section 2.3.2, "Specify back end for volume creation" .
Size (GB)	Volume size (in gigabytes). If you want to create an encrypted volume from an unencrypted image, you must ensure that the volume size is at least 1GB larger than the image size so that the encryption data does not truncate the volume data.
Availability Zone	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 .

3. 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 Section 2.3.10, "Create, use, or delete volume 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.

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

You can also change the volume type later on. For more information, see [Section 2.3.16, "Volume retyping"](#).

2.3.2. Specify back end for volume creation

Whenever multiple Block Storage (cinder) back ends are configured, you must also create a volume type for each back end. You can then use the type to specify which back end to use for a created volume. For more information about volume types, see [Section 2.2.2, "Group Volume Settings with Volume Types"](#).

To specify a back end when creating a volume, select its corresponding volume type from the Type list (see [Section 2.3.1, "Create a volume"](#)).

If you do not specify a back end during volume creation, the Block Storage service automatically chooses one for you. By default, the service chooses the back end with the most available free space. You can also configure the Block Storage service to choose randomly among all available back ends instead. For more information, see [Section 2.2.7, "Configure How Volumes are Allocated to Multiple Back Ends"](#).

2.3.3. Edit a volume name or description

1. In the dashboard, select **Project > Compute > Volumes**
2. Select the volume's **Edit Volume** button.
3. Edit the volume name or description as required.
4. Click **Edit Volume** to save your changes.

**NOTE**

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 [Section 2.2.6, "Configure volume encryption"](#).

2.3.4. Resize (extend) a volume

**NOTE**

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

1. List the volumes to retrieve the ID of the volume you want to extend:

```
$ cinder list
```

2. To resize the volume, run the following commands to specify the correct API microversion, then pass the volume ID and the new size (a value greater than the old size) as parameters:

```
$ OS_VOLUME_API_VERSION=<API microversion>
$ cinder extend <volume ID> <size>
```

Replace <API_microversion>, <volume_ID>, and <size> with appropriate values, for example:

```
$ OS_VOLUME_API_VERSION=3.42
$ cinder extend 573e024d-5235-49ce-8332-be1576d323f8 10
```

2.3.5. Delete a volume

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

**NOTE**

A volume cannot be deleted if it has existing snapshots. For instructions on how to delete snapshots, see [Section 2.3.10, "Create, use, or delete volume snapshots"](#).

2.3.6. Attach and detach a volume to an instance

Instances can use a volume for persistent storage. A volume can only be attached to one instance at a time. For more information, see [Attaching a volume to an instance](#) in the *Creating and Managing Instances* guide.

2.3.6.1. Attaching a volume to an instance

1. In the dashboard, select **Project > Compute > Volumes**

2. Select the **Edit Attachments** action. If the volume is not attached to an instance, the **Attach To Instance** drop-down list is visible.
3. From the **Attach To Instance** list, select the instance to which you want to attach the volume.
4. Click **Attach Volume**.

2.3.6.2. Detaching a volume from an instance

1. In the dashboard, select **Project > Compute > Volumes**
2. 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.
3. Click **Detach Volume** in this and the next dialog screen.

2.3.7. Attach a volume to multiple instances

Volume multi-attach gives multiple instances simultaneous read/write access to a Block Storage volume. The Ceph RBD driver is supported.



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.

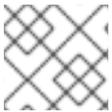
Limitations of multi-attach volumes

- 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. 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:
 - [Component, Plug-In, and Driver Support in Red Hat OpenStack Platform](#)
 - [Red Hat Ecosystem Catalog](#)
- 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.

2.3.7.1. 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.



NOTE

By default, creating a new volume type is an admin-only operation.

Procedure

1. Run the following commands to create a multi-attach volume type:

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



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 based on the available space at the time of creation.

2. Run the following command to specify the back end:

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

2.3.7.2. 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.

2.3.7.3. Creating a multi-attach volume

After you create a multi-attach volume type, create a multi-attach volume.

Procedure

1. Run the following command to create a multi-attach volume:

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

- 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 |
```

You can now attach the volume to multiple instances. For information about how to attach a volume to an instance, see [Attach a volume to an instance](#).

2.3.7.4. Supported back ends

The Block Storage back end must support multi-attach. For information about supported back ends, contact Red Hat Support.

2.3.8. Read-only volumes

A volume can be marked read-only to protect its data from being accidentally overwritten or deleted. To do so, set the volume to read-only by using the following command:

```
# cinder readonly-mode-update <VOLUME-ID> true
```

To set a read-only volume back to read-write, run:

```
# cinder readonly-mode-update <VOLUME-ID> false
```

2.3.9. Change a volume owner

To change the owner of a volume, you must perform a volume transfer. A volume transfer is initiated by the volume owner, and the change in ownership is complete after the transfer is accepted by the new owner.

2.3.9.1. Transfer a volume from the command line

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

```
# cinder list
```

- Initiate the volume transfer:

```
# cinder transfer-create VOLUME
```

Where **VOLUME** is 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 TRANSFERID TRANSFERKEY
```

Where **TRANSFERID** and **TRANSFERKEY** are the **id** and **auth_key** values returned by the **cinder transfer-create** command, respectively. For example,

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



NOTE

You can view all available volume transfers using:

```
# cinder transfer-list
```

2.3.9.2. Transfer a volume by using the dashboard

Create a volume transfer from the dashboard

1. As the volume owner in the dashboard, select **Projects > Volumes**.
2. In the **Actions** column of the volume to transfer, select **Create Transfer**.
3. 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.

4. 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. As the recipient project owner in the dashboard, select **Projects > Volumes**.

2. Click **Accept Transfer**.
3. 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.

2.3.10. Create, use, or delete 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](#).

To create a volume snapshot:

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

You can clone new volumes from a snapshot when it appears in the **Volume Snapshots** table. Select the **Create Volume** action of the snapshot. For more information about volume creation, see [Section 2.3.1, "Create a volume"](#).



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.

To delete a snapshot, select its **Delete Volume Snapshot** action.

If your OpenStack deployment uses a Red Hat Ceph back end, see [Section 2.3.10.1, "Protected and unprotected snapshots in a Red Hat Ceph Storage back end"](#) for more information about snapshot security and troubleshooting.



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.

2.3.10.1. Protected and unprotected snapshots in a Red Hat Ceph Storage back end

When using Red Hat Ceph Storage as a back end for your OpenStack deployment, you can set a snapshot to *protected* in the back end. Attempting to delete protected snapshots through OpenStack (as in, through the dashboard or the **cinder snapshot-delete** command) will fail.

When this occurs, set the snapshot to *unprotected* in the Red Hat Ceph back end first. Afterwards, you should be able to delete the snapshot through OpenStack as normal.

For more information, see the following links in the *Red Hat Ceph Storage Block Device Guide* :

- [Protecting a block device snapshot](#)
- [Unprotecting a block device snapshot](#)

2.3.11. Use a snapshot to restore to the last state of a volume

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) API microversion 3.40 or later.
- You must have created at least one snapshot for the volume.

Procedure

1. Log in to the undercloud as the **stack** user.

2. Source the **overcloudrc** file:

```
[stack@undercloud ~] $ source overcloudrc
```

3. 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.

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

```
$ cinder snapshot-list
```

5. 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.

6. 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
```

7. 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.

2.3.11.1. Verifying that your revert is successful

Procedure

- 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.

2.3.12. Upload a volume to the Image Service

You can upload an existing volume as an image to the Image service directly. To do so:

1. In the dashboard, select **Project > Compute > Volumes**
2. Select the target volume's **Upload to Image** action.
3. Provide an **Image Name** for the volume and select a **Disk Format** from the list.
4. 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 [Manage images](#) in the *Creating and Managing Images* guide.

2.3.13. 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:** The default policy allows volume owners and administrators to retype a volume. The retype operation is the most common way to move volumes between back ends.
- **Migrate:** The default policy allows only administrators to migrate a volume. 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 [Migrate a volume](#).

Restrictions

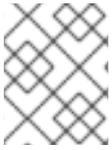
Red Hat supports moving volumes between back ends within and across availability zones (AZs), but with the following restrictions:

- 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.

2.3.14. 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.

2.3.14.1. Generic volume migration

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.

2.3.15. 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), or RBD.

2.3.16. Volume retyping

Volume retyping is the standard way to move volumes from one back end to another. The operation requires the administrator to define the appropriate volume types for the different back ends. The default policy allows volume owners and administrators to retype volumes.

When you retype a volume, you apply a volume type and its settings to an already existing volume. For more information about volume types, see [Section 2.2.2, “Group Volume Settings with Volume Types”](#).

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 pre-defined 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.

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 [Section 2.3.16.2, “Retyping a volume from the command line”](#).

Restrictions

- You cannot retype all volumes. For more information about moving volumes between back ends, see [Section 2.3.13, “Moving volumes between back ends”](#).
- You cannot retype unencrypted volumes to be encrypted volume types, but you can retype encrypted volumes to be unencrypted.
- Users with no administrative privileges can only retype volumes that they own.

2.3.16.1. Retyping a volume from the dashboard UI

Use the dashboard UI to retype a 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

- A successful undercloud installation. For more information, see [Installing director on the undercloud](#).

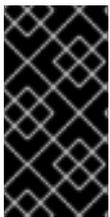
- A successful overcloud deployment. For more information, see [Creating a basic overcloud with CLI tools](#).
- Access to the Red Hat OpenStack Platform (RHOSP) Dashboard (horizon). For more information, see [Overcloud deployment output](#).

Procedure

1. In the dashboard, select **Project > Compute > Volumes**
2. In the **Actions** column of the volume you want to migrate, select **Change Volume Type**.
3. In the **Change Volume Type** dialog, select the target volume type and define the new back end from the **Type** list.
4. If you are migrating the volume to another back end, select **On Demand** from the **Migration Policy** list. For more information, see [Section 2.3.13, “Moving volumes between back ends”](#).
5. Click **Change Volume Type** to start the migration.

2.3.16.2. Retyping a volume from the command line

Similar to the dashboard UI procedure, you can retype a volume from the command line.



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

- A successful undercloud installation. For more information, see [Installing director on the undercloud](#).
- A successful overcloud deployment. For more information, see [Creating a basic overcloud with CLI tools](#).

Procedure

1. Enter the following command to retype a volume:

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

2. 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 id> <new volume type name>
```



NOTE

As the retype operation progresses, the volume status changes to **retyping**.

- Enter the following command and review the **volume_type** field to confirm that the retype operation succeeded. The **volume_type** field shows the new volume type.

```
$ cinder show <volume id>
```



NOTE

When you initiate a retype operation, the volume name is duplicated. If you enter the **cinder show** command with the volume name, the cinder client will return an error similar to **ERROR: Multiple volume matches found for '<volume name>'**. To avoid this error, use the volume ID instead.

2.3.17. 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.



NOTE

This feature is available in this release as a *Technology Preview*, and therefore is not fully supported by Red Hat. It should only be used for testing, and should not be deployed in a production environment. For more information about Technology Preview features, see [Scope of Coverage Details](#).

Prerequisites

- A successful undercloud installation. For more information, see [Installing the undercloud](#).

Procedure

1. Log in to the undercloud host as the **stack** user.

2. Source the undercloud credentials file:

```
$ source ~/stackrc
```

3. Create a new environment file, or modify an existing environment file. In this example, create a new file **lvm2-filtering.yaml**:

```
$ touch ~/lvm2-filtering.yaml
```

4. Include the following parameter in the environment file:

```
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.*
```

5. Run the **openstack overcloud deploy** command and include the environment file that contains the LVM2 filtering configuration, as well as any other environment files that are relevant to your overcloud deployment:

```
$ openstack overcloud deploy --templates \
  <environment-files> \
  -e lvm2-filtering.yaml
```

2.4. ADVANCED VOLUME CONFIGURATION

The following procedures describe how to perform advanced volume management procedures.



IMPORTANT

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

2.4.1. Migrate a Volume

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. The default policy allows only administrators to migrate volumes. Do not change the default policy.

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 [Section 2.3.13, “Moving volumes between back ends”](#).

2.4.1.1. Migrate between back ends

Use the dashboard UI to migrate a volume between back ends.

Procedure

1. In the dashboard, select **Admin > Volumes**.
2. In the **Actions** column of the volume you want to migrate, select **Migrate Volume**.
3. 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.

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

2.4.1.2. Migrating between back ends from the command line

- A successful undercloud installation. For more information, see [Installing director on the undercloud](#).

Procedure

1. 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 back end names take the form **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**.



NOTE

Use of LVM is for example only. LVM is not supported in production environments.

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

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

2.4.1.3. Verifying volume migration

When you create a volume, the **migration_status** value equals **None**. When you initiate the migration, the status changes to **migrating**. When the migration completes, the status changes to either **success** or **error**.

After the Block Storage service accepts the migration request, the cinder client responds with a message similar to **Request to migrate volume <volume id> has been accepted**. However, it takes time for the migration to complete. As an administrator, you can verify the status of the migration.

Procedure

1. Enter the following command and review the **migration_status** field:

```
$ cinder show <volume id>
```



NOTE

When you initiate a generic volume migration, the volume name is duplicated. If you enter the **cinder show** command with the volume name, the cinder client returns an error similar to **ERROR: Multiple volume matches found for '<volume name>'**. To avoid this error, use the volume ID instead.

After a successful migration, the **host** field matches the **<new host>** value set in the **cinder migrate** command.

2.4.2. Encrypting unencrypted volumes

To encrypt an unencrypted volume, you must either back up the unencrypted volume and then restore it to a new encrypted volume, or create an Image service (glance) image from the unencrypted volume and then create a new encrypted volume from the image.

Prerequisites

- An unencrypted volume that you want to encrypt.

Procedure

1. If the **cinder-backup** service is available, back up the current unencrypted volume:

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

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

2. 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.

3. 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.

If the **cinder-backup** service is unavailable, use the **upload-to-image** command to create an image of the unencrypted volume, and then create a new encrypted volume from the image.

1. Create an Image service 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.

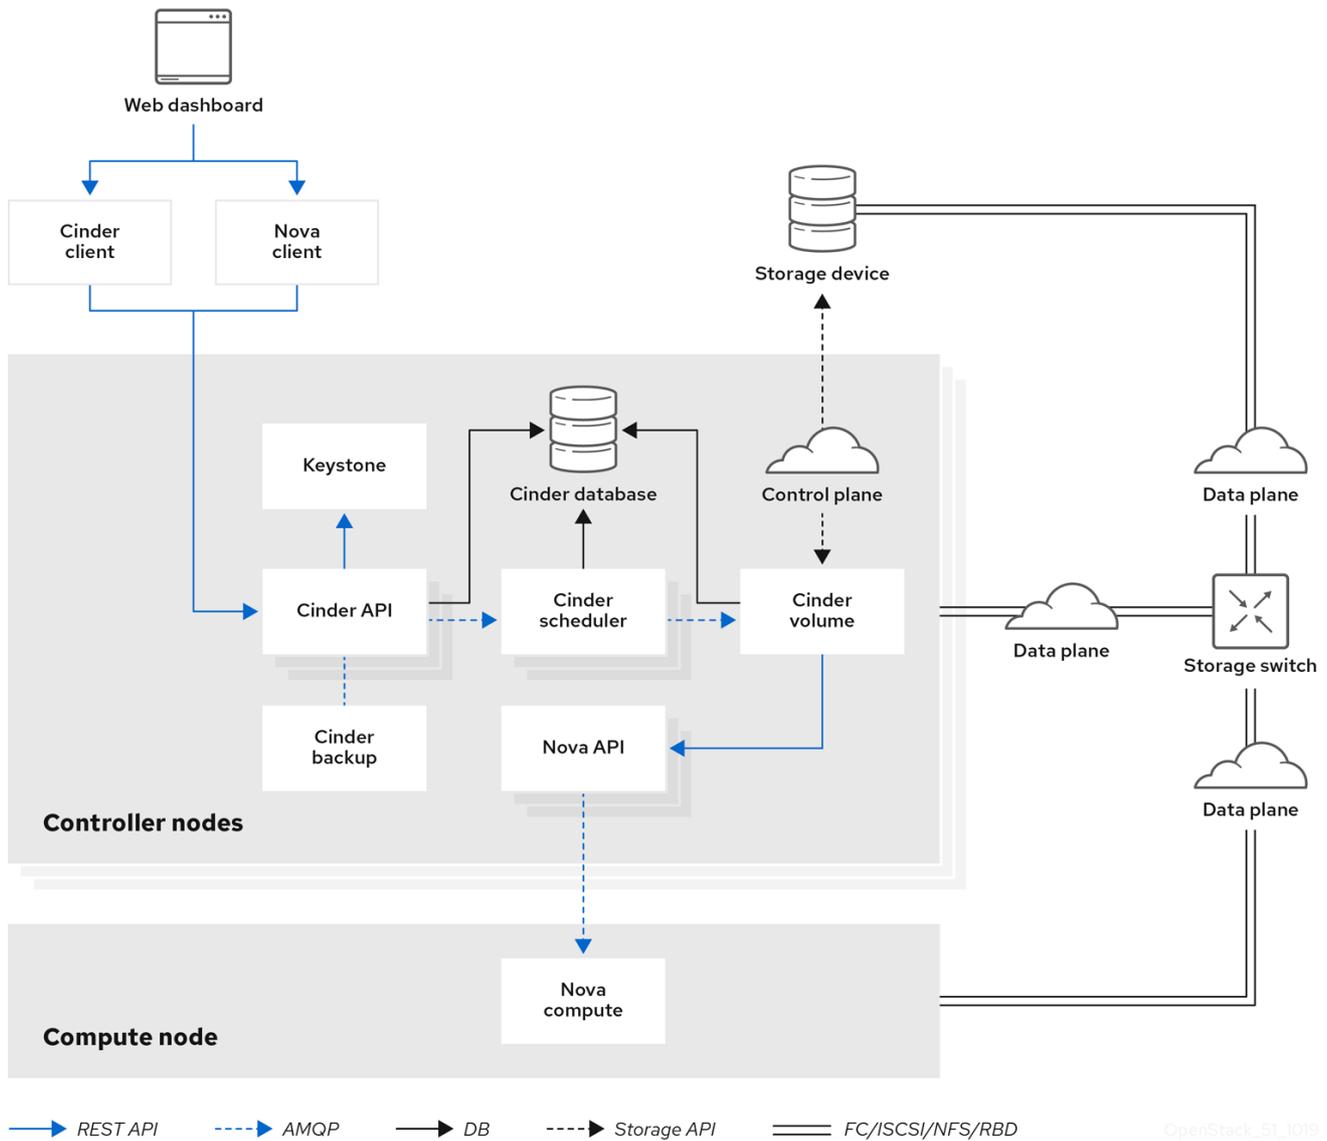
2. Create a new volume from the image that is 1GB larger than the image:

```
$ 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.

2.5. 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. You can configure multipath on new and existing overcloud deployments.



2.5.1. Configuring multipath on new deployments

Complete this procedure to configure multipath on a new overcloud deployment.

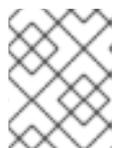
For information about how to configure multipath on existing overcloud deployments, see [Section 2.5.2, "Configuring multipath on existing deployments"](#).

Prerequisites

The overcloud Controller and Compute nodes must have access to the Red Hat Enterprise Linux server repository. For more information, see [Downloading the base cloud image](#) in the *Director Installation and Usage* guide.

Procedure

1. Configure the overcloud.



NOTE

For more information, see [Configuring a basic overcloud with CLI tools](#) in the *Director Installation and Usage* guide.

2. Update the heat template to enable multipath:

```
parameter_defaults:
  NovaLibvirtVolumeUseMultipath: true
  NovaComputeOptVolumes:
    - /etc/multipath.conf:/etc/multipath.conf:ro
    - /etc/multipath/;/etc/multipath/:rw
  CinderVolumeOptVolumes:
    - /etc/multipath.conf:/etc/multipath.conf:ro
    - /etc/multipath/;/etc/multipath/:rw
```

3. Optional: If you are using Block Storage (cinder) as an Image service (glance) back end, you must also complete the following steps:

- a. Add the following **GlanceApiOptVolumes** configuration to the heat template:

```
parameter_defaults:
  GlanceApiOptVolumes:
    - /etc/multipath.conf:/etc/multipath.conf:ro
    - /etc/multipath/;/etc/multipath/:rw
```

- b. Set the **ControllerExtraConfig** parameter in the following way:

```
parameter_defaults:
  ControllerExtraConfig:
    glance::config::api_config:
      default_backend/cinder_use_multipath:
        value: true
```

Note

Ensure that both **default_backend** and the **GlanceBackendID** heat template default value match.

4. For every configured back end, set **use_multipath_for_image_xfer** to **true**:

```
parameter_defaults:
  ExtraConfig:
    cinder::config::cinder_config:
      <backend>/use_multipath_for_image_xfer:
        value: true
```

5. Deploy the overcloud:

```
$ openstack overcloud deploy
```



NOTE

For information about creating the overcloud using overcloud parameters, see [Creating the Overcloud with the CLI Tools](#) in the *Director Installation and Usage* guide.

6. Before containers are running, install multipath on all Controller and Compute nodes:

```
$ sudo dnf install -y device-mapper-multipath
```

**NOTE**

Director provides a set of hooks to support custom configuration for specific node roles after the first boot completes and before the core configuration begins. For more information about custom overcloud configuration, see [Pre-Configuration: Customizing Specific Overcloud Roles](#) in the *Advanced Overcloud Customization* guide.

- Configure the multipath daemon on all Controller and Compute nodes:

```
$ mpathconf --enable --with_multipathd y --user_friendly_names n --find_multipaths y
```

**NOTE**

The example code creates 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 the [Configuring device mapper multipath](#) guide.

- Run the following command on all Controller and Compute nodes to prevent partition creation:

```
$ sed -i "s/^defaults {/defaults {\n\tskip_kpartx yes\n\trecheck_wwid yes/" /etc/multipath.conf
```

**NOTE**

Setting **skip_kpartx** to **yes** prevents kpartx on the Compute node from automatically creating partitions on the device, which prevents unnecessary device mapper entries. For more information about configuration attributes, see [Modifying the DM-Multipath configuration file](#) in the *Configuring device mapper multipath* guide.

- Start the multipath daemon on all Controller and Compute nodes:

```
$ systemctl enable --now multipathd
```

2.5.2. Configuring multipath on existing deployments

Configure multipath on existing deployments so that your workloads can use multipath functionality.

**NOTE**

Any new workloads that you create after you configure multipath on existing deployments are multipath-aware by default. If you have any pre-existing workloads, you must shelve and unshelve the instances to enable multipath on these instances.

For more information about how to configure multipath on new overcloud deployments, see [Section 2.5.1, “Configuring multipath on new deployments”](#).

Prerequisites

The overcloud Controller and Compute nodes must have access to the Red Hat Enterprise Linux server repository. For more information, see [Downloading the base cloud image](#) in the *Director Installation and Usage* guide.

Procedure

1. Verify that multipath is installed on all Controller and Compute nodes:

```
$ rpm -qa | grep device-mapper-multipath
device-mapper-multipath-0.4.9-127.el8.x86_64
device-mapper-multipath-libs-0.4.9-127.el8.x86_64
```

If multipath is not installed, install it on all Controller and Compute nodes:

```
$ sudo dnf install -y device-mapper-multipath
```

2. Configure the multipath daemon on all Controller and Compute nodes:

```
$ mpathconf --enable --with_multipathd y --user_friendly_names n --find_multipaths y
```



NOTE

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

3. Run the following command on all Controller and Compute nodes to prevent partition creation:

```
$ sed -i "s/^defaults {/defaults {\n\tskip_kpartx yes\n\trecheck_wwid yes/" /etc/multipath.conf
```



NOTE

Setting **skip_kpartx** to **yes** prevents kpartx on the Compute node from automatically creating partitions on the device, which prevents unnecessary device mapper entries. For more information about configuration attributes, see [Modifying the DM-Multipath configuration file](#) in the *Configuring device mapper multipath* guide.

4. Start the multipath daemon on all Controller and Compute nodes:

```
$ systemctl enable --now multipathd
```

5. Update the heat template to enable multipath:

```
parameter_defaults:
  NovaLibvirtVolumeUseMultipath: true
  NovaComputeOptVolumes:
```

```

- /etc/multipath.conf:/etc/multipath.conf:ro
- /etc/multipath/;/etc/multipath/:rw
CinderVolumeOptVolumes:
- /etc/multipath.conf:/etc/multipath.conf:ro
- /etc/multipath/;/etc/multipath/:rw

```

6. Optional: If you are using Block Storage (cinder) as an Image service (glance) back end, you must also complete the following steps:

- a. Add the following **GlanceApiOptVolumes** configuration to the heat template:

```

parameter_defaults:
  GlanceApiOptVolumes:
    - /etc/multipath.conf:/etc/multipath.conf:ro
    - /etc/multipath/;/etc/multipath/:rw

```

- b. Set the **ControllerExtraConfig** parameter in the following way:

```

parameter_defaults:
  ControllerExtraConfig:
    glance::config::api_config:
      default_backend/cinder_use_multipath:
        value: true

```

Note

Ensure that both **default_backend** and the **GlanceBackendID** heat template default value match.

7. For every configured back end, set **use_multipath_for_image_xfer** to **true**:

```

parameter_defaults:
  ExtraConfig:
    cinder::config::cinder_config:
      <backend>/use_multipath_for_image_xfer:
        value: true

```

8. Run the following command to update the overcloud:

```
$ openstack overcloud deploy
```



NOTE

When you run the **openstack overcloud deploy** command to install and configure multipath, you must pass all previous roles and environment files that you used to deploy the overcloud, such as **--templates**, **--roles-file**, **-e** for all environment files, and **--timeout**. Failure to pass all previous roles and environment files can cause problems with your overcloud deployment. For more information about using overcloud parameters, see [Creating the Overcloud with the CLI Tools](#) in the *Director Installation and Usage* guide.

2.5.3. Verifying multipath configuration

This procedure describes how to verify multipath configuration on new or existing overcloud deployments.

Procedure

1. Create a VM.
2. Attach a non-encrypted volume to the VM.
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 VM that you booted.

4. Retrieve the virsh name of the instance:

```
$ nova show INSTANCE | grep instance_name
```

Replace *INSTANCE* with the name of the VM that you booted.

5. Get the IP address of the Compute node:

```
$ . stackrc
$ nova list | grep compute_name
```

Replace *compute_name* with the name from the output of the **nova show *INSTANCE*** command.

6. SSH into the Compute node that runs the VM:

```
$ ssh heat-admin@COMPUTE_NODE_IP
```

Replace *COMPUTE_NODE_IP* with the IP address of the Compute node.

7. Log in to the container that runs virsh:

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

8. 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. OBJECT STORAGE SERVICE

OpenStack Object Storage (**swift**) stores its objects (data) in containers, which 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 might include photos, text files, or images. Stored objects are not compressed.

3.1. OBJECT STORAGE RINGS

Object Storage 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. It allows the Object Storage service to 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 different types of data: one for account information, another for containers (to facilitate organizing objects under an account), and another for object replicas.

3.1.1. Rebalancing rings

When you change the Object Storage environment by adding or removing storage capacity, nodes, or disks, you must rebalance the rings. You can run **openstack overcloud deploy** to rebalance the rings, but this method redeploys the entire overcloud. This can be cumbersome, especially if you have a large overcloud. Alternatively, run the following command on the undercloud to rebalance the rings:

```
source ~/stackrc
ansible-playbook -i /usr/bin/tripleo-ansible-inventory
/usr/share/openstack-tripleo-common/playbooks/swift_ring_rebalance.yaml
```

3.1.2. Checking cluster health

The Object Storage service 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 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.

To use **swift-recon**, log in to one of the controller nodes and run the following command:

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

```

=====-->
Starting reconnaissance on 3 hosts (object)
=====
12-14 14:55:47] Checking async pendings
[async_pending] - No hosts returned valid data.
=====
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.
=====
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
=====
12-14 14:55:47] Checking ring md5sums
3/3 hosts matched, 0 error[s] while checking hosts.
=====
12-14 14:55:47] Checking swift.conf md5sum
3/3 hosts matched, 0 error[s] while checking hosts.
=====
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
=====
12-14 14:55:47] Checking time-sync
3/3 hosts matched, 0 error[s] while checking hosts.
=====

```



NOTE

As an alternative, 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 can't update database files fast enough, some updates will occur asynchronously. These numbers should decrease over time.
- Replication metrics: Notice the replication timestamps; full replication passes should happen frequently and there should be 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 ring files are consistent on all nodes.
- Quarantined files: There should be no (or very few) quarantined files for all nodes.
- Time-sync: All nodes must be synchronized.

3.1.3. Increasing ring partition power

The ring power determines the partition to which a resource (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.

3.1.4. Creating 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 may differ from those of the original nodes. Without custom adjustments, the larger capacity of the new nodes may be underutilized. Or, if weights change in the rings, data dispersion can become uneven, which reduces safety.

Automation may 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.

3.2. OBJECT STORAGE SERVICE ADMINISTRATION

The following procedures explain how to customize the Object Storage service.

3.2.1. Configuring fast-post

By default, the Object Storage service 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 by doing the following:

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
```

3.2.2. Enabling at-rest encryption

By default, objects uploaded to Object Storage are kept 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.

You can use OpenStack Key Manager (barbican) to encrypt at-rest swift objects. For more information, see [Encrypt at-rest swift objects](#).

3.2.3. Deploying a standalone Object Storage cluster

You can use the composable role concept to deploy a standalone Object Storage (openstack-swift) cluster with the bare minimum of additional services (for example Keystone, HAProxy). The [Creating a roles_data File](#) section has information on roles.

3.2.3.1. Creating the roles_data.yaml File

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*, Mistral*, 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

3.2.3.2. Deploying the New Roles

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

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

3.2.4. Using external SAN disks

By default, when the Red Hat OpenStack Platform 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 impacts during node failure or other system issues.

In similar 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, Red Hat strongly recommends that you use additional local disks for Object Storage instead to meet performance and disk space requirements. For more information, see [Object Storage](#) in the *Deployment Recommendations for Specific Red Hat OpenStack Platform Services* guide.

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:

- The Object Storage service stores telemetry data and Image service (glance) images by default. Glance images require more disk space, but from a performance perspective, the impact of storing glance images impacts performance less than storing telemetry data. Storing and processing telemetry data requires increased performance. 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.

3.2.4.1. SAN disk deployment configuration

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

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

3.3. INSTALL AND CONFIGURE STORAGE NODES FOR RED HAT ENTERPRISE LINUX

To use Red Hat OpenStack Platform Object Storage service (swift) on external storage nodes, you must install and configure the storage nodes that operate the account, container, and object services processes. This configuration references two storage nodes, each of which contain two empty local block storage devices.

NOTE

The internal network for the Object Storage service is not authenticated. For security purposes, Red Hat recommends that you keep the storage nodes on a dedicated network or VLAN.

NOTE

The instructions use `/dev/sdb` and `/dev/sdc` as device names, but you can substitute the values for specific nodes in your environment.

3.3.1. Preparing storage devices

Before you install and configure the Object Storage service on the storage nodes, you must prepare the storage devices.



NOTE

Perform all of these steps on each storage node.

Procedure

1. Install the supporting utility packages:

```
# yum install xfsprogs rsync
```

2. Format the **/dev/sdb** and **/dev/sdc** devices as XFS:

```
# mkfs.xfs /dev/sdb  
# mkfs.xfs /dev/sdc
```

3. Create the mount point directory structure:

```
# mkdir -p /srv/node/sdb  
# mkdir -p /srv/node/sdc
```

4. Edit the **/etc/fstab** file and add the following data to it:

```
/dev/sdb /srv/node/sdb xfs defaults 0 2  
/dev/sdc /srv/node/sdc xfs defaults 0 2
```

5. Mount the devices:

```
# mount /srv/node/sdb  
# mount /srv/node/sdc
```

6. Create or edit the **/etc/rsyncd.conf** file to contain the following data:

```
uid = swift  
gid = swift  
log file = /var/log/rsyncd.log  
pid file = /var/run/rsyncd.pid  
address = MANAGEMENT_INTERFACE_IP_ADDRESS
```

```
[account]  
max connections = 2  
path = /srv/node/  
read only = False  
lock file = /var/lock/account.lock
```

```
[container]  
max connections = 2  
path = /srv/node/  
read only = False
```

```
lock file = /var/lock/container.lock
```

```
[object]
max connections = 2
path = /srv/node/
read only = False
lock file = /var/lock/object.lock
```

Replace *MANAGEMENT_INTERFACE_IP_ADDRESS* with the IP address of the management network on the storage node.

7. Start the **rsyncd** service and configure it to start when the system boots:

```
# systemctl enable rsyncd.service
# systemctl start rsyncd.service
```

3.3.2. Configuring components

Configure the account, container, and object storage servers.

Procedure

1. Install the packages:

```
# yum install openstack-swift-account openstack-swift-container \
openstack-swift-object
```

2. Edit the **/etc/swift/account-server.conf** file and complete the following actions:

- a. In the **[DEFAULT]** section, configure the bind IP address, bind port, user, configuration directory, and mount point directory:

```
[DEFAULT]
...
bind_ip = MANAGEMENT_INTERFACE_IP_ADDRESS
bind_port = 6202
user = swift
swift_dir = /etc/swift
devices = /srv/node
mount_check = True
```

Replace *MANAGEMENT_INTERFACE_IP_ADDRESS* with the IP address of the management network on the storage node.

- b. In the **[pipeline:main]** section, enable the **healthcheck** and **recon** modules:

```
[pipeline:main]
pipeline = healthcheck recon account-server
```

- c. In the **[filter:recon]** section, configure the recon cache directory:

```
[filter:recon]
use = egg:swift#recon
...
```

```
recon_cache_path = /var/cache/swift
```

3. Open the default firewall port for the account service:

```
# firewall-cmd --permanent --add-port=6202/tcp
```

4. Edit the **/etc/swift/container-server.conf** file and complete the following actions:

- a. In the **[DEFAULT]** section, configure the bind IP address, bind port, user, configuration directory, and mount point directory:

```
[DEFAULT]
...
bind_ip = MANAGEMENT_INTERFACE_IP_ADDRESS
bind_port = 6201
user = swift
swift_dir = /etc/swift
devices = /srv/node
mount_check = True
```

Replace *MANAGEMENT_INTERFACE_IP_ADDRESS* with the IP address of the management network on the storage node.

- b. In the **[pipeline:main]** section, enable the **healthcheck** and **recon** modules:

```
[pipeline:main]
pipeline = healthcheck recon container-server
```

- c. In the **[filter:recon]** section, configure the recon cache directory:

```
[filter:recon]
use = egg:swift#recon
...
recon_cache_path = /var/cache/swift
```

5. Open the default firewall port for the container service:

```
# firewall-cmd --permanent --add-port=6201/tcp
```

6. Edit the **/etc/swift/object-server.conf** file and complete the following actions:

- a. In the **[DEFAULT]** section, configure the bind IP address, bind port, user, configuration directory, and mount point directory:

```
[DEFAULT]
...
bind_ip = MANAGEMENT_INTERFACE_IP_ADDRESS
bind_port = 6200
user = swift
swift_dir = /etc/swift
devices = /srv/node
mount_check = True
```

Replace `MANAGEMENT_INTERFACE_IP_ADDRESS` with the IP address of the management network on the storage node.

- b. In the `[pipeline:main]` section, enable the `healthcheck` and `recon` modules:

```
[pipeline:main]
pipeline = healthcheck recon object-server
```

- c. In the `[filter:recon]` section, configure the `recon_cache_path` and the `recon_lock_path` directories:

```
[filter:recon]
use = egg:swift#recon
...
recon_cache_path = /var/cache/swift
recon_lock_path = /var/lock
```

7. Open the default firewall port for the object service:

```
# firewall-cmd --permanent --add-port=6200/tcp
```

8. Ensure that the ownership of the mount point directory structure is correct:

```
# chown -R swift:swift /srv/node
```

9. Create the `recon` directory and ensure proper ownership of it:

```
# mkdir -p /var/cache/swift
# chown -R root:swift /var/cache/swift
# chmod -R 775 /var/cache/swift
```

3.4. BASIC CONTAINER MANAGEMENT

To help with organization, pseudo-folders are logical devices that can contain objects (and can 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.

3.4.1. Creating a container

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.

Type	Description
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**.

New containers use the default storage policy. If you have multiple storage policies defined (for example, a default and another that enables erasure coding), you can configure a container to use a non-default storage policy through the command line. To do so, run:

```
# swift post -H "X-Storage-Policy:POLICY" CONTAINERNAME
```

Where:

- *POLICY* is the name or alias of the policy you want the container to use.
- *CONTAINERNAME* is the name of the container.

3.4.2. Creating a pseudo folder for a container

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**.

3.4.3. Deleting a container

1. In the dashboard, select **Project > Object Store > Containers**
2. Browse for the container in the **Containers** section, and ensure all objects have been deleted (see [Section 3.4.6, "Deleting an object"](#)).
3. Select **Delete Container** in the container's arrow menu.
4. Click **Delete Container** to confirm the container's removal.

3.4.4. Uploading an object

If you do not upload an actual file, the object is still created (as placeholder) and can later be used to upload the file.

1. In the dashboard, select **Project > Object Store > Containers**
2. Click the name of the container in which the uploaded object will be placed; 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:
 - Pseudo-folders can be specified in the name 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 object's contents.
5. Click **Upload Object**.

3.4.5. Copying an object

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.
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**.

3.4.6. Deleting an object

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's removal.

CHAPTER 4. SHARED FILE SYSTEMS SERVICE

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.

Use the following concepts and procedures to understand and use the Shared File Systems service:

- To choose a back end, see [Section 4.1, “Shared File Systems service \(manila\) back ends”](#) .
- To understand your networking choices, see [Section 4.1.1, “Networking for shared file systems”](#) .
- To create and manage shares, see:
 - [Section 4.1.2, “Creating a share type”](#)
 - [Section 4.1.4, “Discovering share types”](#)
 - [Section 4.1.5, “Creating a share”](#)
 - [Section 4.1.6, “Listing shares and exporting information”](#)
- To manage network connectivity, see:
 - [Section 4.2, “Ensuring network connectivity to the share”](#)
 - [Section 4.2.1, “Connecting to a shared network to access shares”](#)
- To manage share access, see:
 - [Section 4.3.1, “Granting access to a share”](#)
 - [Section 4.3.2, “Revoking access to a share”](#)
- To manage shares, see:
 - [Section 4.4.1, “Listing shares export locations”](#)
 - [Section 4.4.2, “Mounting the share”](#)
 - [Section 4.5, “Deleting a share”](#)

- To manage quotas, see:
 - [Section 4.6.1, “Listing quotas”](#)
- To troubleshoot issues, see:
 - [Section 4.7.1, “Fixing create share or create share group failures”](#)
 - [Section 4.7.2, “Debugging share mounting failures”](#)

4.1. SHARED FILE SYSTEMS SERVICE (MANILA) BACK ENDS

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

- CephFS through NFS. For more information, see [Deploying the Shared File Systems service with CephFS through NFS](#).
- Native CephFS. For more information, see [CephFS Back End Guide for the Shared File System Service](#).
- 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](#).

4.1.1. 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 4.2, “Ensuring network connectivity to the share”](#)

4.1.2. Creating a share type

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, but you can create private share types for use within specific projects. For more information about how to make private share types, or to set additional share-type options, see the [Security and Hardening Guide](#).

4.1.3. 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 4.1. 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.

Capability	Values	Description
availability_zones	a list of one or more availability zones	Grants permission to create shares only on the availability zones listed.

4.1.4. Discovering share types

As a cloud user, you must specify a share type when you create a share.

Procedure

1. Discover the available share types:

```
$ manila type-list
```

The command output lists the name and ID of the share type.

4.1.5. Creating a share

Create a share to read and write data.

To create a 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 CephFS native, **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 4.1.2, “Creating a share type”](#), the cloud administrator created a **default** share type that selects a CephFS back end and another share type named **netapp** that selects a NetApp back end.

Procedure

1. Use the example share types to create a 10 GB NFS share named **share-01** on the CephFS NFS back end. This example uses CephFS with NFS:

```
(user) [stack@undercloud-0 ~]$ manila create --name share-01 nfs 10
```

- Optional: Create a 20 GB NFS share named **share-02** on the NetApp back end:

```
(user) [stack@undercloud-0 ~]$ manila create --name share-02 --share-type netapp --share-network mynet nfs 20
```

4.1.6. Listing shares and exporting information

To verify that you successfully created the shares, complete the following steps.

Procedure

- List the shares:

```
(user) [stack@undercloud-0 ~]$ manila list
```

Name	...	Status	...	ID
8c3bedd8-bc82-4100-a65d-53ec51b5fe81	share-01	...	available ...	

- View the export locations of the share:

```
(user) [stack@undercloud-0 ~]$ manila share-export-location-list share-01
```

Path
172.17.5.13:/volumes/_nogroup/e840b4ae-6a04-49ee-9d6e-67d4999fbc01

- View the parameters for the share:

```
manila share-export-location-show <id>
```



NOTE

This information is used to mount the share in [Section 4.4.2, "Mounting the share"](#).

4.2. 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 4.2.1, "Connecting to a shared network to access shares"](#) .

4.2.1. 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) [stack@undercloud-0 ~]$ 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) [stack@undercloud-0 ~]$ 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) [stack@undercloud-0 ~]$ openstack server add port instance0 nfs-port0
(user) [stack@undercloud-0 ~]$ 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.

4.2.2. 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 issues with DHCPv6 on the secondary interface. Use this procedure to configure an IPv6 interface manually on the instance. For more information, see [BZ#1859695](#).

Prerequisites

- Connect to a shared network to access shares.

Procedure

1. Log in to the instance.
2. Configure the IPv6 interface address:

```
$ ip address add fd00:fd00:fd00:7000::c/64 dev eth1
```

3. Activate the interface:

```
$ 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:

```
$ dnf install -y telnet
$ telnet fd00:fd00:fd00:7000::21 2049
```

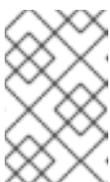
4.3. GRANT SHARE ACCESS

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:

```
# 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 4.1.5, "Creating a share"](#).
- **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 user name.



NOTE

The type of access depends on the protocol of the share. For NFS shares, only **ip** access type is allowed. For CIFS, **user** access type is appropriate. For native CephFS shares, you must use **cephx**.

- **accesslevel** - optional, default is **rw**
 - **rw**: read-write access to shares.
 - **ro**: read-only access to shares.
- **clientidentifier** - varies depending on **acesstype**.
 - Use an IP address for **ip accesstype**.
 - Use CIFS user or group for **user accesstype**.
 - Use a user name string for **cephx accesstype**.

4.3.1. Granting access to a share

You must grant end user clients access to the share so that users can read data from and write data to the share.

Use this procedure to 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 desired.



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) [stack@undercloud-0 ~]$ 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) [stack@undercloud-0 ~]$ 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) [stack@undercloud-0 ~]$ manila access-list share-01

+-----+-----+-----+-----+ ...
| id      | access_type | access_to | access_level | state | ...
+-----+-----+-----+-----+
| 875c6251-... | ip      | 172.17.5.160 | rw      | active | ...
+-----+-----+-----+-----+ ...

```

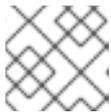
4.3.2. Revoking access to a share

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

1. Revoke access to a share:

```
# manila access-deny <share> <accessid>
```



NOTE

Replace **<share>** with either the share name or the share ID.

For example:

```

(user) [stack@undercloud-0 ~]$ manila access-list share-01
+-----+-----+-----+-----+ ...
| id      | access_type | access_to | access_level | state | ...
+-----+-----+-----+-----+
| 875c6251-... | ip      | 172.17.5.160 | rw      | active | ...
+-----+-----+-----+-----+

(user) [stack@undercloud-0 ~]$ manila access-deny share-01 875c6251-c17e-4c45-8516-
fe0928004fff

(user) [stack@undercloud-0 ~]$ 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.

4.4. MOUNT SHARE ON COMPUTE INSTANCES

After you grant access to clients, shares can be mounted and used by them. 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.

4.4.1. Listing shares export locations

Retrieve the export locations of shares so that you can mount a share.

Procedure

1. Retrieve the export location of a share:

```
(user) [stack@undercloud-0 ~]$ 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.

4.4.2. Mounting the share

Mount a share on the client to enable access to data.

For information about creating and granting share access, see the following procedures:

- [Section 4.1.5, "Creating a share"](#)
- [Section 4.3, "Grant share access"](#)

Procedure

1. Log in to the instance and run the following command:

```
(user) [stack@undercloud-0 ~]$ openstack server ssh demo-instance0 --login root
# hostname
demo-instance0
```

2. Mount the share on an IPv4 network by using the export location:

```
# mount -t nfs -v 172.17.5.13:/volumes/_nogroup/e840b4ae-6a04-49ee-9d6e-67d4999fbc01 /mnt
```

4.5. DELETING A SHARE

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.

Procedure

1. Delete a share:

```
# manila delete <share>
```



NOTE

In the example command, <share> can be either the share name or the share ID.

For example:

```
# manila delete share-01
```

4.6. CHANGE 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.

4.6.1. Listing quotas

As a cloud administrator, you can list the quotas for a project or user by using the **manila quota-show** command. If you include the optional **--user** parameter, you can view the quota for this user in the specified project. If you omit this parameter, you get the quotas for the specified project.

You can update and delete quotas. You can update the shares, snapshots, gigabytes, snapshot-gigabytes, share-networks, share_groups, share_group_snapshots, and share-type quotas.

Procedure

1. To see the usage statements, run the following commands:

```
# manila help quota-show  
# manila help quota-update  
# manila help quota-delete
```

4.7. TROUBLESHOOTING FAILURES


```
| 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.

5. 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.


```

True          |      |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+

```

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           | {}                                |
+-----+-----+

```

11. 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 |
+-----+-----+-----+-----+-----+-----+

```

4.7.2. Debugging share mounting failures

If you experience trouble when you mount shares, use these verification steps to identify the root cause of the issue.

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 logs and contact Red Hat Support.