



# Red Hat OpenStack Platform 12

## Storage Guide

Understanding, using, and managing persistent storage in OpenStack



# Red Hat OpenStack Platform 12 Storage Guide

---

Understanding, using, and managing persistent storage in OpenStack

OpenStack Team  
rhos-docs@redhat.com

## Legal Notice

Copyright © 2018 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, 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 Software Collections 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

<b>PREFACE</b> .....	<b>4</b>
<b>CHAPTER 1. INTRODUCTION TO PERSISTENT STORAGE IN OPENSTACK</b> .....	<b>5</b>
1.1. SCALABILITY AND BACK END	6
1.2. ACCESSIBILITY AND ADMINISTRATION	6
1.3. SECURITY	6
1.4. REDUNDANCY AND DISASTER RECOVERY	7
<b>CHAPTER 2. BLOCK STORAGE AND VOLUMES</b> .....	<b>8</b>
2.1. BACK ENDS	8
2.2. BLOCK STORAGE SERVICE ADMINISTRATION	8
2.2.1. Group Volume Settings with Volume Types	8
2.2.1.1. List a Host Driver's Capabilities	9
2.2.1.2. Create and Configure a Volume Type	10
2.2.1.3. Edit a Volume Type	10
2.2.1.4. Delete a Volume Type	11
2.2.1.5. Create and Configure Private Volume Types	11
2.2.2. Create and Configure an Internal Tenant for the Block Storage Service	12
2.2.3. Configure and Enable the Image-Volume Cache	13
2.2.4. Use Quality-of-Service Specifications	14
2.2.4.1. Create and Configure a QOS Spec	14
2.2.4.2. Set Capacity-Derived QoS Limits	15
2.2.4.3. Associate a QOS Spec with a Volume Type	16
2.2.4.4. Disassociate a QOS Spec from a Volume Type	16
2.2.5. Configure Volume Encryption	16
2.2.5.1. Configure Volume Type Encryption	16
2.2.6. Configure How Volumes are Allocated to Multiple Back Ends	17
2.2.7. Configure and Use Consistency Groups	18
2.2.7.1. Set Up Consistency Groups	19
2.2.7.2. Create and Manage Consistency Groups	20
2.2.7.3. Create and Manage Consistency Group Snapshots	21
2.2.7.4. Clone Consistency Groups	21
2.2.8. Backup Administration	22
2.2.8.1. View and Modify a Tenant's Backup Quota	22
2.2.8.2. Enable Volume Backup Management Through the Dashboard	22
2.2.8.3. Set an NFS Share as a Backup Repository	23
2.2.8.3.1. Set a Different Backup File Size	24
2.3. BASIC VOLUME USAGE AND CONFIGURATION	24
2.3.1. Create a Volume	24
2.3.2. Specify Back End for Volume Creation	25
2.3.3. Edit a Volume's Name or Description	26
2.3.4. Delete a Volume	26
2.3.5. Attach and Detach a Volume to an Instance	26
2.3.5.1. Attach a Volume to an Instance	26
2.3.5.2. Detach a Volume From an Instance	27
2.3.6. Set a Volume to Read-Only	27
2.3.7. Change a Volume's Owner	27
2.3.7.1. Transfer a Volume from the Command Line	27
2.3.7.2. Transfer a Volume Using the Dashboard	28
2.3.8. Create, Use, or Delete Volume Snapshots	29
2.3.8.1. Protected and Unprotected Snapshots in a Red Hat Ceph Back End	29
2.3.9. Upload a Volume to the Image Service	30

2.3.10. Changing a Volume's Type (Volume Re-typing)	30
2.4. ADVANCED VOLUME CONFIGURATION	30
2.4.1. Back Up and Restore a Volume	31
2.4.1.1. Create a Full Volume Backup	31
2.4.1.1.1. Create a Volume Backup as an Admin	32
2.4.1.2. Create an Incremental Volume Backup	33
2.4.1.3. Restore a Volume After a Block Storage Database Loss	33
2.4.1.4. Restore a Volume from a Backup	34
2.4.2. Migrate a Volume	34
2.4.2.1. Migrate Between Hosts	35
2.4.2.2. Migrate Between Back Ends	35
<b>CHAPTER 3. OBJECT STORAGE</b>	<b>36</b>
3.1. OBJECT STORAGE SERVICE ADMINISTRATION	36
3.1.1. Configure Erasure Coding	36
3.1.1.1. Configure an Erasure Coding Policy	36
3.1.1.2. Configure an Object Storage Ring	38
3.1.1.3. Using Erasure Coding	39
3.1.2. Configure Fast-POST	39
3.1.3. Set Object Storage as a Back End for the Image Service	40
3.1.4. Enable At-Rest Encryption	41
3.1.5. Deploying a Standalone Object Storage Cluster	42
3.1.5.1. Creating the roles_data.yaml File	42
3.1.5.2. Deploying the New Roles	44
3.2. BASIC CONTAINER MANAGEMENT	44
3.2.1. Create a Container	45
3.2.2. Create a Pseudo Folder for a Container	45
3.2.3. Delete a Container	45
3.2.4. Upload an Object	46
3.2.5. Copy an Object	46
3.2.6. Delete an Object	46
<b>CHAPTER 4. FILE SHARES</b>	<b>48</b>
4.1. BACK ENDS	48
4.2. CREATE AND MANAGE SHARE TYPES	48
4.3. CREATE A SHARE	49
4.4. LIST SHARES AND EXPORT INFORMATION	50
4.5. GRANT SHARE ACCESS	50
4.6. REVOKE ACCESS TO A SHARE	51
4.7. MOUNT A SHARE ON AN INSTANCE	51
4.8. DELETE A SHARE	52
4.9. DISABLE CREATION OF PUBLIC SHARES AND SHARE TYPES THROUGH THE DASHBOARD	52



## PREFACE

Red Hat OpenStack Platform (Red Hat OpenStack Platform) provides the foundation to build a private or public Infrastructure-as-a-Service (IaaS) cloud on top of Red Hat Enterprise Linux. It offers a massively scalable, fault-tolerant platform for the development of cloud-enabled workloads.

This guide discusses procedures for creating and managing persistent storage. Within OpenStack, 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 using either the OpenStack dashboard or the command-line clients. Most procedures can be carried out using either method; some of the more advanced procedures can only be executed 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:

```
# yum install crudini -y
```



# CHAPTER 1. INTRODUCTION TO PERSISTENT STORAGE IN OPENSTACK


OpenStack 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, on the other hand, 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. OpenStack 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.

Volumes also provide inherent redundancy and disaster recovery through backups and snapshots. In addition, you can also encrypt volumes for added security. For more information about volumes, see [Chapter 2, Block Storage and Volumes](#).



### NOTE

Instances can also be configured to use absolutely no ephemeral storage. In such cases, the Block Storage service can write images to a volume; in turn, the volume can be used as a bootable root volume for an instance.

## 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 through **containers**.

While a volume's contents can only be accessed 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. For example, the Data Processing service (**openstack-sahara**) can manage all of its binaries, data input, data output, and templates directly through the Object Storage service.

## Shares

The Shared File System 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 using Red Hat Ceph as a Block Storage back end, you can scale storage capacity and redundancy by adding more Ceph OSD (Object Storage Daemon) nodes. Both Block Storage and Object Storage services support Red Hat Ceph 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 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 available underlying block storage. You can also scale capacity by adding more storage devices to the storage node.

The Shared File System Service provisions shares backed by storage from a separate *storage pool*. This pool (which is typically managed by a third-party back end service) provides the share with storage at the file system level. The Shared File System Service can use both NetApp and CephFS, which can be configured to use a storage pool of pre-created volumes which provisioned shares can use for storage. In either deployment, scaling involves adding more volumes to the pool.

## 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 System 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 provides basic data security through *volume encryption*. With this, you can configure a volume type to be encrypted through a static key; the key will then be used for encrypting all volumes created from the configured volume type. See [Section 2.2.5, “Configure Volume Encryption”](#) for more details.

Object and container security, on the other hand, is configured at the service and node level. The Object Storage service 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’s network security in order to protect object data.

The Shared File System service can secure shares through access restriction, whether by instance IP,

user/group, or TLS certificate. In addition, some Shared File System service deployments can feature a 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.

## 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 2.2.7, “Configure and Use Consistency Groups”](#) for more details.

Finally, the Block Storage service also features *volume replication*. This allows you to configure volumes to replicate content between each other, thereby providing basic redundancy.



### NOTE

Volume replication is only available through specific third-party back ends and their respective drivers.

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 System 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.

### 2.1. 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, we advise that you deploy a more robust back end for an Enterprise environment.

When deploying Red Hat OpenStack Platform for the environment, we recommend using the 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. 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](#) (from the [Advanced Overcloud Customization](#) guide).

#### 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 the following:

- [Dell EMC PS Series](#)
- [Dell Storage Center](#)
- [NetApp](#) (for supported appliances)
- [Fujitsu Eternus](#) is also supported as a back end, but is not yet integrated into the Director.

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

### 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.

#### 2.2.1. Group Volume Settings with Volume Types

OpenStack allows you to create volume types, which allows you apply the type's associated settings. You can apply these settings during volume creation ([Section 2.3.1, "Create a Volume"](#)) or even afterwards ([Section 2.3.10, "Changing a Volume's Type \(Volume Re-typing\)"](#)). For example, you can associate:

- Whether or not a volume is encrypted ([Section 2.2.5.1, "Configure Volume Type Encryption"](#))
- Which back end a volume should use ([Section 2.3.2, "Specify Back End for Volume Creation"](#) and [Section 2.4.2.2, "Migrate Between Back Ends"](#))

- 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.



## NOTE

Available and supported Extra Specs vary per volume driver. Consult your volume driver's documentation for a list of valid Extra Specs.

### 2.2.1.1. List a Host Driver's Capabilities

Available and supported Extra Specs vary per back end driver. Consult the driver's 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.

### 2.2.1.2. 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 2.2.1.5, “Create and Configure Private Volume Types”](#).



#### NOTE

You can also associate a QOS Spec to the volume type. For details, refer to [Section 2.2.4.3, “Associate a QOS Spec with a Volume Type”](#).

### 2.2.1.3. 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. To do this, select the setting's **Edit** action.

- Delete existing settings associated with the volume type. To do this, select the extra specs' check box and click **Delete Extra Specs** in this and the next dialog screen.

#### 2.2.1.4. Delete a Volume Type

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

#### 2.2.1.5. Create and Configure Private Volume Types

By default, all volume types are visible to all tenants. You can override this during volume type creation and set it to **private**. To do so, you will need to set the type's **Is\_Public** flag to **False**.

Private volume types are useful for restricting access to certain volume settings. 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 --os-volume-api-version 2 type-create --is-public false VTYPE
```

+ Replace *VTYPE* with the name of the private volume type.

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 --os-volume-api-version 2 type-list --all
```

This command will list both public and private volume types, and will also include the name and ID of each one. You will 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 --os-volume-api-version 2 type-access-add --volume-type VTYPEID -
  -project-id TENANTID
```

Where:

- *VTYPEID* is the ID of the private volume type.
- *TENANTID* is the ID of the project/tenant you are granting access to *VTYPEID*.

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

```
# cinder --os-volume-api-version 2 type-access-list --volume-type VTYPE
```

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

```
# cinder --os-volume-api-version 2 type-access-remove --volume-type VTYPE
  --project-id TENANTID
```

**NOTE**

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

### 2.2.2. 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/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 |
+-----+-----+
```


Note that creating the project and user will display their respective IDs. Configure the Block Storage service to use both project and user as the internal project through their IDs. To do so, run the following on each Block Storage node:

```
# crudini --set /etc/cinder/cinder.conf DEFAULT
cinder_internal_tenant_project_id TENANTID
# crudini --set /etc/cinder/cinder.conf DEFAULT
cinder_internal_tenant_user_id USERID
```

Replace *TENANTID* and *USERID* with the respective IDs of the **cinder-internal** project and user, which you created earlier. For example, using the IDs supplied above:

```
# crudini --set /etc/cinder/cinder.conf DEFAULT
cinder_internal_tenant_project_id cb91e1fe446a45628bb2b139d7dccaef
# crudini --set /etc/cinder/cinder.conf DEFAULT
cinder_internal_tenant_user_id 84e9672c64f041d6bfa7a930f558d946
```



**NOTE**

This procedure involves configuring a service outside of the director. As such, you may need to repeat it the next time you re-deploy or update the overcloud.

### 2.2.3. 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 2.2.2, “Create and Configure an Internal Tenant for the Block Storage Service”](#).

To enable and configure the Image-Volume cache on a back end (*BACKEND*), run the following commands:

```
# crudini --set /etc/cinder/cinder.conf BACKEND image_volume_cache_enabled
True
```

Replace *BACKEND* with the name of the target back end (specifically, its **volume\_backend\_name** value).

By default, the Image-Volume cache size is only limited by the back end. To configure a maximum size (*MAXSIZE*, in GB):

```
# crudini --set /etc/cinder/cinder.conf BACKEND
image_volume_cache_max_size_gb MAXSIZE
```

Alternatively, you can also set a maximum number of images (*MAXNUMBER*). To do so:

```
# crudini --set /etc/cinder/cinder.conf BACKEND
image_volume_cache_max_count 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 configuring the Image-Volume cache, restart the Block Storage service. To do so, log in to any Controller node as the **heat-admin** user and run:

```
# pcs resource restart openstack-cinder-volume
```



## NOTE

This procedure involves configuring a service outside of the director. As such, you may need to repeat it the next time you re-deploy or update the overcloud.

### 2.2.4. 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.

#### 2.2.4.1. 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
<b>back-end</b>	QOS policy will be applied to the Block Storage back end.
<b>front-end</b>	QOS policy will be applied to Compute.
<b>both</b>	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.

- 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 a capacity-derived read limit of **500** IOPS, use the following Key/Value pair:

```
read_iops_sec_per_gb=500
```



#### NOTE

For more information about capacity-derived QoS limits, see [Section 2.2.4.2, “Set Capacity-Derived QoS Limits”](#).

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

### 2.2.4.2. 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 2.2.1, “Group Volume Settings with Volume Types”](#) and [Section 2.2.1.2, “Create and Configure a Volume Type”](#). For instructions on how to set QoS specs, [Section 2.2.4, “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 2.3.10, “Changing a Volume’s Type \(Volume Re-typing\)”](#) for information about volume re-typing.

### 2.2.4.3. 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.

### 2.2.4.4. 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.

## 2.2.5. Configure Volume Encryption

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



### IMPORTANT


At present, volume encryption is only supported on volumes backed by block devices. Encryption of network-attached volumes (such as RBD) or file-based volumes (such as NFS) is still unsupported.

Volume encryption is applied through volume type. See [Section 2.2.5.1, “Configure Volume Type Encryption”](#) for information on encrypted volume types.

### 2.2.5.1. Configure Volume Type Encryption

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 2.2.1.2, “Create and Configure a Volume Type”](#). Specifically, select the encrypted volume type from the Type drop-down list in the **Create Volume** window (see to [Section 2.3, “Basic Volume Usage and Configuration”](#)).

You can also re-configure the encryption settings of an encrypted volume type. To do so:

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**, use 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. 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 (when the OpenStack Data Processing service is enabled)

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.

Alternatively, you can also manually configure the filter scheduler. To do so:

1. Log in as **heat-admin** to the node hosting the Block Storage service.
2. Enable the **FilterScheduler** scheduler driver.

```
$ sudo crudini --set /etc/cinder/cinder.conf DEFAULT
scheduler_driver cinder.scheduler.filter_scheduler.FilterScheduler
```

3. Set which filters should be active:

```
$ sudo crudini --set /etc/cinder/cinder.conf DEFAULT
scheduler_default_filters
AvailabilityZoneFilter, CapacityFilter, CapabilitiesFilter
```

4. Configure how the scheduler should select a suitable back end. If you want the scheduler:

- To always choose the back end with the most available free space, run:

```
$ sudo crudini --set /etc/cinder/cinder.conf DEFAULT
scheduler_default_weighters AllocatedCapacityWeighter
$ sudo crudini --set /etc/cinder/cinder.conf DEFAULT
allocated_capacity_weight_multiplier -1.0
```

- To choose randomly among all suitable back ends, run:

```
$ crudini --set /etc/cinder/cinder.conf DEFAULT
scheduler_default_weighters ChanceWeighter
```

5. Restart the Block Storage scheduler to apply your settings. To do so, log in to any Controller node as the **heat-admin** user and run:

```
$ pcs resource restart openstack-cinder-scheduler
```



## NOTE

This procedure involves configuring a service outside of the director. As such, you may need to repeat it the next time you re-deploy or update the overcloud.

### 2.2.7. Configure and Use Consistency Groups

The Block Storage service allows you to set *consistency groups*. With this, you can group multiple volumes together as a single entity. This, in turn, allows you to perform operations on multiple volumes at once, rather than individually. Specifically, this release allows you to use consistency groups to create snapshots for multiple volumes simultaneously. By extension, this will also allow you to restore or clone those volumes simultaneously.

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

As of this release, consistency groups are only supported by the drivers of the following storage back ends:

- EMC VMAX
- EMC VNX
- EMC ScaleIO
- EMC ExtremIO
- HP 3Par StorServ
- IBM DS8000
- IBM StorwizeSVC
- IBM XIV
- NetApp Data ONTAP
- NetApp ESERIES
- NetApp SolidFire

### 2.2.7.1. Set Up Consistency Groups

By default, Block Storage security policy disables consistency groups APIs. You need to enable it here before using the feature. To do so, edit the related consistency group entries in `/etc/cinder/policy.json` of the node hosting the Block Storage API service (namely, **openstack-cinder-api**). The entries appear as follows:

```
"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",
```

For increased security, set the permissions for both consistency group API and volume type management API 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",
```

So, to enable the consistency group APIs as recommended, edit their entries as follows:

```
"consistencygroup:create" : "rule:admin_api",
"consistencygroup:delete": "rule:admin_api",
"consistencygroup:update": "rule:admin_api",
"consistencygroup:get": "rule:admin_api",
"consistencygroup:get_all": "rule:admin_api",
"consistencygroup:create_cgsnapshot" : "rule:admin_api",
"consistencygroup:delete_cgsnapshot": "rule:admin_api",
"consistencygroup:get_cgsnapshot": "rule:admin_api",
"consistencygroup:get_all_cgsnapshots": "rule:admin_api",
```

## NOTE

You can also make the consistency groups feature available to all users. To do so, 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**:

```
"consistencygroup:create" : "rule:admin_or_owner",
"consistencygroup:delete": "rule:admin_or_owner",
"consistencygroup:update": "rule:admin_or_owner",
"consistencygroup:get": "rule:admin_or_owner",
"consistencygroup:get_all": "rule:admin_or_owner",
"consistencygroup:create_cgsnapshot" : "rule:admin_or_owner",
"consistencygroup:delete_cgsnapshot": "rule:admin_or_owner",
"consistencygroup:get_cgsnapshot": "rule:admin_or_owner",
"consistencygroup:get_all_cgsnapshots": "rule:admin_or_owner",
```

After enabling the consistency group APIs, restart the Block Storage API service. To do so, log in to any Controller node as the **heat-admin** user and run:

```
# pcs resource restart openstack-cinder-api
```

## NOTE

This procedure involves configuring a service outside of the director. As such, you may need to repeat it the next time you re-deploy or update the overcloud.

### 2.2.7.2. Create and Manage Consistency Groups

After enabling the consistency groups API, you can then start creating consistency groups. To do so:

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 should appear next in the **Volume Consistency Groups** table.

You can change the name or description of a consistency group by selecting **Edit Consistency Group** from its **Action** column.

In addition, you can also add or remove volumes from a consistency group directly. To do so:

1. As an admin user in the dashboard, select **Project > Compute > Volumes > Volume Consistency Groups**.
2. Find the consistency group you want to configure. In the **Actions** column of that consistency group, select **Manage Volumes**. Doing so will launch 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.
3. Click **Edit Consistency Group**.

### 2.2.7.3. Create and Manage Consistency Group Snapshots

After adding volumes to a consistency group, you can now create snapshots from it. Before doing so, first log in as **admin** user from the command line on the node hosting the **openstack-cinder-api** and run:

```
# export OS_VOLUME_API_VERSION=2
```

Doing so will configure the client to use version **2** of **openstack-cinder-api**.

To list all available consistency groups (along with their respective IDs, which you will need later):

```
# cinder consisgroup-list
```

To create snapshots using the consistency group, run:

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

Where:

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

To display a list of all available consistency group snapshots, run:

```
# cinder cgsnapshot-list
```

### 2.2.7.4. Clone Consistency Groups

Consistency groups can also be used 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.

To clone an existing consistency group:

```
# cinder consistgroup-create-from-src --source-cg CGNAMEID --name CGNAME --  
description "DESCRIPTION"
```

Where: - *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 consistgroup-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.

## 2.2.8. Backup Administration

The following sections discuss how to customize the Block Storage service's volume backup settings.

### 2.2.8.1. View and Modify a Tenant's Backup Quota

Unlike most tenant storage quotas (number of volumes, volume storage, snapshots, etc.), backup quotas cannot be modified through the dashboard yet.

Backup quotas can only be modified through the command-line interface; namely, through the **cinder quota-update** command.

To view the storage quotas of a specific tenant (*TENANTNAME*), run:

```
# cinder quota-show TENANTNAME
```

To update the maximum number of backups (*MAXNUM*) that can be created in a specific tenant, run:

```
# cinder quota-update --backups MAXNUM TENANTNAME
```

To update the maximum total size of all backups (*MAXGB*) within a specific tenant, run:

```
# cinder quota-update --backup-gigabytes MAXGB TENANTNAME
```

To view the storage quota usage of a specific tenant, run:

```
# cinder quota-usage TENANTNAME
```

### 2.2.8.2. Enable Volume Backup Management Through the Dashboard

You can now create, view, delete, and restore volume backups through the dashboard. To perform any of these functions, go to the **Project > Compute > Volumes > Volume Backups** tab.

However, the **Volume Backups** tab is not enabled by default. To enable it, configure the dashboard accordingly:

1. Open `/etc/openstack-dashboard/local_settings`.
2. Search for the following setting:

```
OPENSTACK_CINDER_FEATURES = {
    'enable_backup': False,
}
```

Change this setting to:

```
OPENSTACK_CINDER_FEATURES = {
    'enable_backup': True,
}
```

3. Restart the dashboard by restarting the `httpd` service:

```
# systemctl restart httpd.service
```



#### NOTE

This procedure involves configuring a service outside of the director. As such, you may need to repeat it the next time you re-deploy or update the overcloud.

### 2.2.8.3. Set an NFS Share as a Backup Repository

By default, the Block Storage service uses the Object Storage service as a repository for backups. You can configure the Block Storage service to use an existing NFS share as a backup repository instead. To do so:

1. Log in to the node hosting the backup service (**openstack-cinder-backup**) as a user with administrative privileges.
2. Configure the Block Storage service to use the NFS backup driver (`cinder.backup.drivers.nfs`):

```
# crudini --set /etc/cinder/cinder.conf DEFAULT backup_driver
cinder.backup.drivers.nfs
```

3. Set the details of the NFS share that you want to use as a backup repository:

```
# crudini --set /etc/cinder/cinder.conf DEFAULT backup_share
NFSHOST:PATH
```

Where:

- `NFSHOST` is the IP address or hostname of the NFS server.
  - `PATH` is the absolute path of the NFS share on `NFSHOST`.
4. If you want to set any optional mount settings for the NFS share, run:

```
# crudini --set /etc/cinder/cinder.conf DEFAULT backup_mount_options
NFSMOUNTOPTS
```

Where *NFSMOUNTOPTS* is a comma-separated list of NFS mount options (for example, *rw,sync*). For more information on supported mount options, see the **man** pages for **nfs** and **mount**.

- Restart the Block Storage backup service to apply your changes:

```
# systemctl restart openstack-cinder-backup.service
```



## NOTE

This procedure involves configuring a service outside of the director. As such, you may need to repeat it the next time you re-deploy or update the overcloud.

### 2.2.8.3.1. Set a Different Backup File Size

The backup service limits backup files sizes to a maximum **backup file size**. If you are backing up a volume that exceeds this size, the resulting backup will be split into multiple chunks. The default backup file size is 1.8GB.

To set a different backup file size, run:

```
# crudini --set /etc/cinder/cinder.conf DEFAULT backup_file_size SIZE
```

Replace *SIZE* with the file size you want, in bytes. Restart the Block Storage backup service to apply your changes:

```
# systemctl restart openstack-cinder-backup.service
```



## NOTE

This procedure involves configuring a service outside of the director. As such, you may need to repeat it the next time you re-deploy or update the overcloud.

## 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.

### 2.3.1. Create a Volume

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

Field	Description
Volume name	Name of the volume.

Field	Description
Description	Optional, short description of the volume.
Type	Optional volume type (see <a href="#">Section 2.2.1</a> , “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 <a href="#">Section 2.3.2</a> , “Specify Back End for Volume Creation” for details.
Size (GB)	Volume size (in gigabytes).
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 on availability zones and host aggregates, see <a href="#">Manage Host Aggregates</a> .

### 3. Specify a **Volume Source**:

Source	Description
No source, empty volume	The volume will be empty, and will not contain a file system or partition table.
Snapshot	Use an existing snapshot as a volume source. If you select this option, a new Use snapshot as a source list appears; you can then choose a snapshot from the list. For more information about volume snapshots, refer to <a href="#">Section 2.3.8</a> , “Create, Use, or Delete Volume Snapshots”.
Image	Use an existing image as a volume source. If you select this option, a new Use image as a source lists appears; 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 volume as a source list appears; 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’s type later on. For details, see [Section 2.3.10](#), “Changing a Volume’s Type (Volume Re-typing)”.

## 2.3.2. Specify Back End for Volume Creation

Whenever multiple Block Storage back ends are configured, you will also need to create a volume type

for each back end. You can then use the type to specify which back end should be used for a created volume. For more information about volume types, see [Section 2.2.1, “Group Volume Settings with Volume Types”](#).

To specify a back end when creating a volume, select its corresponding volume type from the Type drop-down list (see [Section 2.3.1, “Create a Volume”](#)).

If you do not specify a back end during volume creation, the Block Storage service will automatically choose one for you. By default, the service will choose 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 2.2.6, “Configure How Volumes are Allocated to Multiple Back Ends”](#).

### 2.3.3. Edit a Volume’s 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, both Compute and Block Storage services must be configured to use the same static key. For information on how to set up the requirements for volume encryption, refer to [Section 2.2.5, “Configure Volume Encryption”](#).

### 2.3.4. 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 2.3.8, “Create, Use, or Delete Volume Snapshots”](#).

### 2.3.5. 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 on instances, see **Manage Instances** in the **Instances and Images Guide** available at [Red Hat OpenStack Platform](#).

#### 2.3.5.1. Attach a Volume to an Instance

1. In the dashboard, select **Project > Compute > Volumes**.
2. Select the volume’s **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 wish to attach the volume.
4. Click **Attach Volume**.

### 2.3.5.2. Detach 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.

### 2.3.6. Set a Volume to Read-Only

You can give multiple users shared access to a single volume without allowing them to edit its contents. To do so, set the volume to **read-only** using the following command:

```
# cinder readonly-mode-update VOLUME true
```

Replace **VOLUME** with the **ID** of the target volume.

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

```
# cinder readonly-mode-update VOLUME false
```

### 2.3.7. Change a Volume's Owner

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

#### 2.3.7.1. Transfer a Volume from the Command Line

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

```
# cinder list
```

3. Initiate the volume transfer:

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

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.

- 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
```

### 2.3.7.2. Transfer a Volume Using the Dashboard

#### Create a volume transfer from the dashboard

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

- 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.

### 2.3.8. Create, Use, or Delete Volume Snapshots

You can preserve a volume's state 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. In addition, you cannot delete a volume if it has existing snapshots. Volume backups are used to prevent data loss, whereas snapshots are used to facilitate cloning.

For this reason, snapshot back ends are typically co-located with volume back ends in order to minimize latency during cloning. By contrast, a backup repository is usually located in a different location (eg. 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, refer to [Section 2.4.1, "Back Up and Restore a Volume"](#)

To create a volume snapshot:

1. In the dashboard, select **Project > Compute > Volumes**.
2. Select the target volume's **Create Snapshot** action.
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 once it appears in the **Volume Snapshots** table. To do so, select the snapshot's **Create Volume** action. For more information about volume creation, see [Section 2.3.1, "Create a Volume"](#).

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

If your OpenStack deployment uses a Red Hat Ceph back end, see [Section 2.3.8.1, "Protected and Unprotected Snapshots in a Red Hat Ceph Back End"](#) for more information on snapshot security and troubleshooting.

#### 2.3.8.1. Protected and Unprotected Snapshots in a Red Hat Ceph Back End

When using Red Hat Ceph 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 related instructions, see [Protecting a Snapshot](#) and [Unprotecting a Snapshot](#).

### 2.3.9. 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**. The QEMU disk image utility uploads a new image of the chosen format using the name you provided.

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 **Instances and Images Guide** available at [Red Hat OpenStack Platform](#).

### 2.3.10. Changing a Volume's Type (Volume Re-typing)

*Volume re-typing* is the process of applying a volume type (and, in turn, its settings) to an already existing volume. For more information about volume types, see [Section 2.2.1, "Group Volume Settings with Volume Types"](#).

A volume can be re-typed whether or not it has an existing volume type. In either case, a re-type will only be successful if the Extra Specs of the volume type can be applied to the volume. Volume re-typing is useful for applying pre-defined settings or storage attributes to an existing volume, such as when you want to:

- Migrate the volume to a different back end ([Section 2.4.2.2, "Migrate Between Back Ends"](#)).
- Change the volume's storage class/tier.

Users with no administrative privileges can only re-type volumes they own. To perform a volume re-type:

1. In the dashboard, select **Project > Compute > Volumes**.
2. In the **Actions** column of the volume to be migrated, select **Change Volume Type**.
3. In the **Change Volume Type** dialog, select the target volume type defining the new back end from the **Type** drop-down list.



#### NOTE

If you are migrating the volume to another back end, select **On Demand** from the **Migration Policy** drop-down list. For more information, see [Section 2.4.2.2, "Migrate Between Back Ends"](#).

4. Click **Change Volume Type** to start the migration.

## 2.4. ADVANCED VOLUME CONFIGURATION

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

### 2.4.1. Back Up and Restore a Volume

A volume backup is a persistent copy of a volume's contents. Volume backups are typically created as object stores, and are managed through the Object Storage service by default. You can, however, set up a different repository for your backups; OpenStack supports Red Hat Ceph and NFS as alternative back ends for backups.

When creating a volume backup, all of the backup's metadata is stored in the Block Storage service's database. The `cinder` utility uses this metadata when restoring a volume from the backup. As such, when recovering from a catastrophic database loss, you must restore the Block Storage service's database first before restoring any volumes from backups. This also presumes that the Block Storage service database is being restored with all the original volume backup metadata intact.

If you wish to configure only a subset of volume backups to survive a catastrophic database loss, you can also export the backup's metadata. In doing so, you can then re-import the metadata to the Block Storage database later on, and restore the volume backup as normal.



#### 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. In addition, you cannot delete a volume if it has existing snapshots. Volume backups are used to prevent data loss, whereas snapshots are used to facilitate cloning.

For this reason, snapshot back ends are typically co-located with volume back ends in order to minimize latency during cloning. By contrast, a backup repository is usually located in a different location (eg. 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 snapshots, refer to [Section 2.3.8, "Create, Use, or Delete Volume Snapshots"](#).

#### 2.4.1.1. Create a Full Volume Backup

To back up a volume, use the `cinder backup-create` command. By default, this command will create a full backup of the volume. If the volume has existing backups, you can choose to create an **incremental** backup instead (see [Section 2.4.1.2, "Create an Incremental Volume Backup"](#) for details.)

You can create backups of volumes you have access to. This means that users with administrative privileges can back up any volume, regardless of owner. For more information, see [Section 2.4.1.1.1, "Create a Volume Backup as an Admin"](#).

1. View the **ID** or **Display Name** of the volume you wish to back up:


```
# cinder list
```

2. Back up the volume:

```
# cinder backup-create VOLUME
```

Replace *VOLUME* with the **ID** or **Display Name** of the volume you want to back up. For example:

Property	Value
id	e9d15fc7-eeae-4ca4-aa72-d52536dc551d
name	None
volume_id	5f75430a-abff-4cc7-b74e-f808234fa6c5

**NOTE**

The **volume\_id** of the resulting backup is identical to the **ID** of the source volume.

- Verify that the volume backup creation is complete:

```
# cinder backup-list
```

The volume backup creation is complete when the **Status** of the backup entry is **available**.

At this point, you can also export and store the volume backup's metadata. This allows you to restore the volume backup, even if the Block Storage database suffers a catastrophic loss. To do so, run:

```
# cinder --os-volume-api-version 2 backup-export BACKUPID
```

Where *BACKUPID* is the ID or name of the volume backup. For example,

Property	Value
backup_service	cinder.backup.drivers.swift
backup_url	eyJzdGF0dXMiOiAiYXZhaWxhYmxlIiwgIm9iam...
	...4NS02ZmY4MzBhZWYwNWUiLCAic2l6ZSI6IDF9

The volume backup metadata consists of the **backup\_service** and **backup\_url** values.

#### 2.4.1.1.1. Create a Volume Backup as an Admin

Users with administrative privileges (such as the default **admin** account) can back up any volume managed by OpenStack. When an admin backs up a volume owned by a non-admin user, the backup is hidden from the volume owner by default.

As an admin, you can still back up a volume **and** make the backup available to a specific tenant. To do so, run:

```
# cinder --os-auth-url KEYSTONEURL --os-tenant-name TENANTNAME --os-username USERNAME --os-password PASSWD backup-create VOLUME
```

Where:

- TENANTNAME* is the name of the tenant where you want to make the backup available.

- *USERNAME* and *PASSWD* are the username/password credentials of a user within *TENANTNAME*.
- *VOLUME* is the name or ID of the volume you want to back up.
- *KEYSTONEURL* is the URL endpoint of the Identity service (typically `http://IP:5000/v2`, where *IP* is the IP address of the Identity service host).

When performing this operation, the resulting backup's size will count against the quota of *TENANTNAME* rather than the admin's tenant.

#### 2.4.1.2. Create an Incremental Volume Backup

By default, the `cinder backup-create` command will create a full backup of a volume. However, if the volume has existing backups, you can choose to create an **incremental** backup.

An incremental backup captures any changes to the volume since the last backup (full or incremental). Performing numerous, regular, full back ups of a volume can become resource-intensive as the volume's size increases over time. In this regard, incremental backups allow you to capture periodic changes to volumes while minimizing resource usage.

To create an incremental volume backup, use the `--incremental` option. As in:

```
# cinder backup-create VOLUME --incremental
```

Replace *VOLUME* with the **ID** or **Display Name** of the volume you want to back up. Incremental backups are fully supported on NFS and Object Storage backup repositories.



#### NOTE

You cannot delete a full backup if it already has an incremental backup. In addition, if a full backup has multiple incremental backups, you can only delete the latest one.



#### WARNING

When using Red Hat Ceph Storage as a back end for both Block Storage (cinder) volumes and backups, any attempt to perform an incremental backup will result in a full backup instead, without any warning. This is a known issue ([BZ#1463058](#)).

#### 2.4.1.3. Restore a Volume After a Block Storage Database Loss

Typically, a Block Storage database loss prevents you from restoring a volume backup. This is because the Block Storage database contains metadata required by the volume backup service (openstack-cinder-backup). This metadata consists of `backup_service` and `backup_url` values, which you can export after creating the volume backup (as shown in [Section 2.4.1.1, "Create a Full Volume Backup"](#)).

If you exported and stored this metadata, then you can import it to a new Block Storage database (thereby allowing you to restore the volume backup).

1. As a user with administrative privileges, run:

```
# cinder --os-volume-api-version 2 backup-import backup_service
backup_url
```

Where *backup\_service* and *backup\_url* are from the metadata you exported. For example, using the exported metadata from [Section 2.4.1.1, “Create a Full Volume Backup”](#):

```
# cinder --os-volume-api-version 2 backup-import
cinder.backup.drivers.swift eyJzdGF0dXMi...c2l6ZSI6IDF9
+-----+-----+-----+-----+
| Property |           Value           |
+-----+-----+-----+-----+
|   id    | 77951e2f-4aff-4365-8c64-f833802eaa43 |
|  name   |           None            |
+-----+-----+-----+-----+
```

2. After the metadata is imported into the Block Storage service database, you can restore the volume as normal (see [Section 2.4.1.4, “Restore a Volume from a Backup”](#)).

#### 2.4.1.4. Restore a Volume from a Backup

1. Find the **ID** of the volume backup you wish to use:

```
# cinder backup-list
```

The **Volume ID** should match the ID of the volume you wish to restore.

2. Restore the volume backup:

```
# cinder backup-restore BACKUP_ID
```

Where *BACKUP\_ID* is the ID of the volume backup you wish to use.

3. If you no longer need the backup, delete it:

```
# cinder backup-delete BACKUP_ID
```

#### 2.4.2. Migrate a Volume

The Block Storage service allows you to migrate volumes between hosts or back ends. Volume migration has some limitations:

- The volume can not be in-use (attached to an instance) or have snapshots.
- The target of the in-use volume migration requires ISCSI block-backed devices and can not use non-block devices, such as Ceph RADOS Block Device (RBD).
- Migrations between volumes on different back ends (and thus drivers) are not optimized.

The speed of any migration depends upon your host setup and configuration. With driver-assisted migration, the data movement takes place at the storage backplane instead of inside of the OpenStack Block Storage service. Otherwise, data is copied from one host to another through the Block Storage service.

### 2.4.2.1. Migrate Between Hosts

When migrating a volume between hosts, both hosts must reside on the same back end. Use the dashboard to migrate a volume between hosts:

1. In the dashboard, select **Admin > Volumes**.
2. In the **Actions** column of the volume to be migrated, select **Migrate Volume**.
3. In the **Migrate Volume** dialog, select the target host from the **Destination Host** drop-down list.



#### NOTE

If you wish to bypass any driver optimizations for the host migration, select the **Force Host Copy** checkbox.

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

### 2.4.2.2. Migrate Between Back Ends

Migrating a volume between back ends, on the other hand, involves **volume re-typing**. This means that in order to migrate to a new back end:

1. The new back end must be specified as an **Extra Spec** in a separate *target volume type*.
2. All other Extra Specs defined in the target volume type must be compatible with the volume's original volume type.

See [Section 2.2.1, “Group Volume Settings with Volume Types”](#) and [Section 2.3.2, “Specify Back End for Volume Creation”](#) for more details.

When defining the back end as an Extra Spec, use **volume\_backend\_name** as the **Key**. Its corresponding value will be the back end's name, as defined in the Block Storage configuration file (`/etc/cinder/cinder.conf`). In this file, each back end is defined in its own section, and its corresponding name is set in the **volume\_backend\_name** parameter.

Once you have a back end defined in a target volume type, you can migrate a volume to that back end through **re-typing**. This involves re-applying the target volume type to a volume, thereby applying the new back end settings. See [Section 2.3.10, “Changing a Volume's Type \(Volume Re-typing\)”](#) for instructions.

To do so:

1. In the dashboard, select **Project > Compute > Volumes**.
2. In the **Actions** column of the volume to be migrated, select **Change Volume Type**.
3. In the **Change Volume Type** dialog, select the target volume type defining the new back end from the **Type** drop-down list.
4. Select **On Demand** from the **Migration Policy** drop-down list.
5. Click **Change Volume Type** to start the migration.

## CHAPTER 3. OBJECT STORAGE

OpenStack Object Storage (**openstack-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 SERVICE ADMINISTRATION

The following procedures explain how to further customize the Object Storage service to suit your needs.

#### 3.1.1. Configure Erasure Coding

Erasure coding (EC) is a method of data protection in which the data is broken into fragments, expanded and encoded with redundant data pieces and stored across a set of different locations or storage media. It uses a smaller volume of storage to attain the required durability than traditional replication. When compared to replication factor of 3, savings of 50% may be attained with careful deployment. However, depending on the workload, erasure coding may incur a performance penalty.

Erasure coding is supported for Object Storage service as a Storage Policy. A Storage Policy allows segmenting the cluster for various purposes through the creation of multiple object rings. Red Hat recommends you split off devices used by erasure coding and replication Storage Policies. This way behavior of the cluster is easier to analyze.

The direction you choose depends on why the erasure coding policy is being deployed. Some of the main considerations are:

- Layout of existing infrastructure.
- Cost of adding dedicated erasure coding nodes (or just dedicated erasure coding devices).
- Intended usage model(s).

##### 3.1.1.1. Configure an Erasure Coding Policy

To use erasure coding, first define an additional policy for it in `/etc/swift/swift.conf`. The [Sample Erasure Coding Policy](#) below shows a typical example:

##### Sample Erasure Coding Policy

```
[storage-policy:1]
default = no
name = ec104
alias = myec,erasure_coding
policy_type = erasure_coding
ec_type = jerasure_rs_vand
ec_num_data_fragments = 10
ec_num_parity_fragments = 4
ec_object_segment_size = 1048576
```



**NOTE**

Object Storage policy headers (for example, **[storage-policy:1]**) include an index number — in this case, **1**. The Object Storage index count begins at **0**, and therefore the [Sample Erasure Coding Policy](#) above assumes that another policy with the policy index **0** already exists. For example:

```
[storage-policy:0]
name = default
default = yes
```

After defining the erasure coding policy, you will need to create and configure its associated object storage ring. See [Section 3.1.1.2, “Configure an Object Storage Ring”](#) for instructions.

The following table explains the different parameters used in [Sample Erasure Coding Policy](#):

**Table 3.1. Storage Policy Parameters**


Term	Description
<b>default</b>	Sets whether this policy is the default one or not ( <b>yes/no</b> ). This is used if there are multiple policies defined in <b>/etc/swift/swift.conf</b> .
<b>name</b>	This is a standard storage policy parameter.
<b>alias</b>	Comma-delimited list of other names that the policy is known as.
<b>policy_type</b>	Set this to <b>erasure_coding</b> to indicate that this is an erasure coding policy.
<b>ec_type</b>	This specifies the erasure coding scheme that is to be used. See <a href="#">Table 3.2, “Supported Erasure Coding Schemes”</a> for a list of supported values.
<b>ec_num_data_fragments</b>	The total number of fragments that will be comprised of data.
<b>ec_num_parity_fragments</b>	The total number of fragments that will be comprised of parity.
<b>ec_object_segment_size</b>	The amount of data that will be buffered up before feeding a segment into the encoder/decoder. The default value is <b>1048576</b> .

**Table 3.2. Supported Erasure Coding Schemes**

Scheme	Description/Reference
<b>liberasurecode_rs_vand</b>	Vandermonde Reed-Solomon encoding, software-only backend implemented by <a href="#">liberasurecode</a>

Scheme	Description/Reference
<code>jerasure_rs_vand</code>	Vandermonde Reed-Solomon encoding, based on <a href="#">Jerasure</a>
<code>jerasure_rs_cauchy</code>	Cauchy Reed-Solomon encoding (Jerasure variant), based on <a href="#">Jerasure</a>
<code>flat_xor_hd_3</code> , <code>flat_xor_hd_4</code>	Flat-XOR based HD combination codes, <a href="#">liberasurecode</a>
<code>isa_l_rs_vand</code>	Intel Storage Acceleration Library (ISA-L) - SIMD accelerated <a href="#">Erasure Coding backends</a>

When the Object Storage service encodes an object, it breaks it into fragments. It is important during configuration to know how many of those fragments are data and how many are parity. So in the [Sample Erasure Coding Policy](#), an object will be broken into 14 different fragments, 10 of them will be made up of actual object data and 4 of them will be made of parity data (calculations depending on `ec_type`). With such a configuration, the system can sustain 4 disk failures before the data is lost. Other commonly used configurations are 4+2 (with 4 data fragments and 2 parity fragments) or 8+3 (with 8 data fragments and 3 parity fragments).



#### NOTE

This procedure involves configuring a service outside of the director. As such, you may need to repeat it the next time you re-deploy or update the overcloud.



#### IMPORTANT

Once you have deployed a policy and have created objects with that policy, these configurations options cannot be changed. In case a change in the configuration is desired, you must create a new policy and migrate the data to a new container. However, once defined, policy indices cannot be discarded. If policies are to be retired, they may be retired, but not be removed. There is essentially no performance penalty for having old policies around, but a minor administrative overhead.

### 3.1.1.2. Configure an Object Storage Ring

Object Storage uses a data structure called the **Ring** to distribute a partition space across the cluster. This partition space is core to the data durability engine in Object Storage service. It allows the Object Storage service to quickly and easily synchronize each partition across the cluster. When any component in Swift needs to interact with data, a quick lookup is done locally in the Ring to determine the possible partitions for each object.

The Object Storage service already has three rings to store different types of data. There is one for account information, another for containers (so that it's convenient to organize objects under an account) and another for the object replicas. To support erasure codes, there will be an additional ring that is created to store erasure code objects.

To create a typical replication ring, for example, you can use the following command:

```
# swift-ring-builder object-1.builder create 10 3 1
```

Where 3 is the number of replicas.

In order to create an erasure coding object ring, you need to use the number of fragments in place of the number of replicas, for example:

```
# swift-ring-builder object-1.builder create 10 14 1
```

Where 14 is for a 10+4 configuration with 10 data fragments and 4 parity fragments.

Consider the performance impacts when deciding which devices to use in the erasure coding policy's object ring. We recommend that you run some performance benchmarking in a test environment for the configuration before deployment. After you have configured your erasure coding policy in the `/etc/swift/swift.conf` and created your object ring, your application is ready to start using erasure coding by creating a container with the specified policy name and interacting as usual.

### 3.1.1.3. Using Erasure Coding

After defining a new Erasure Coding policy and configuring its object storage ring, you can use it when creating a new container for the first time. If you have multiple policies defined, each container will be created with the default storage policy assigned.

To use a non-default storage policy with a new container, you need to submit a special metadata header when creating the container. For example, to use the [Sample Erasure Coding Policy](#) policy defined in [Section 3.1.1.1, "Configure an Erasure Coding Policy"](#) on a new container (`CONTAINERNAME`), run:

```
# swift post -H "X-Storage-Policy:ec104" CONTAINERNAME
```

## 3.1.2. Configure Fast-POST

By default, the Object Storage service will copy an object whole whenever any part of its metadata changes. You can prevent this using the `fast-post` feature. This is useful if you want to save time in changing 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. To do this:

1. Log in to the nodes hosting the Object Storage proxy service.
2. Open `/etc/swift/proxy-server.conf`.
3. Under the `[app:proxy-server]` section, set `object_post_as_copy` to `false`:

```
[app:proxy-server]
use = egg:swift#proxy
set log_name = proxy-server
...
object_post_as_copy = false
...
```

4. Restart the Object Storage proxy services on the node:

```
# systemctl restart openstack-swift-proxy.service
# systemctl restart openstack-swift-object-expirer.service
```

If any storage policies other than **policy-0** are listed in `/etc/swift/swift.conf`, run the following as well:

```
# systemctl restart openstack-swift-container-reconciler.service
```



#### NOTE

This procedure involves configuring a service outside of the director. As such, you may need to repeat it the next time you re-deploy or update the overcloud.



#### NOTE

On a typical overcloud deployment, the Object Storage service is installed on the Controller nodes.

### 3.1.3. Set Object Storage as a Back End for the Image Service

The OpenStack Image service, by default, saves images and instance snapshots to the local filesystem in `/var/lib/glance/images/`. Alternatively, you can configure the Image service to save images and snapshots to the Object Storage service (when available).

To do so, perform the following procedure:

1. Log into the node running the Image service (the controller node also running Identity) as root and source your OpenStack credentials (this is typically a file named **openrc**).

```
# source ~/openrc
```

2. Verify that the Image service is part of the tenant **services** with role **admin**.

```
# openstack user role list --project services glance
```

One of the roles returned should be **admin**.

3. Open the `/etc/glance/glance.conf` file and comment out the following lines:

```
##### DEFAULT OPTIONS #####
#default_store = file
#filesystem_store_datadir = /var/lib/glance/images/
```

4. In the same file, add the following lines to the **DEFAULT OPTIONS** section.

```
default_store = swift
swift_store_auth_address = http://KEYSTONEIP:35357/v2.0/
swift_store_user = service:glance
swift_store_key = ADMINPW
swift_store_create_container_on_put = True
```

Where:

- **KEYSTONEIP** is the IP address of the Identity service, and

- **ADMINPW** is the value of admin password attribute in the `/etc/glance/glance-api.conf` file.

5. Apply the changes by restarting the Image service:

```
# systemctl restart openstack-glance-api
# systemctl restart openstack-glance-registry
```

From this point onwards, images uploaded to the Image service (whether through the Dashboard or **glance**) should now be saved to an Object Storage container named **glance**. This container exists in the service account.

To verify whether newly-created images are saved to the Image service, run:

```
# ls /var/lib/glance/images
```

Once the Dashboard or the **glance image-list** reports the image is active, you can verify whether it is in Object Storage by running the following command:

```
# swift --os-auth-url http://KEYSTONEIP:5000/v2.0 --os-tenant-name service
--os-username glance --os-password ADMINPW list glance
```



#### NOTE

This procedure involves configuring a service outside of the director. As such, you may need to repeat it the next time you re-deploy or update the overcloud.

### 3.1.4. Enable At-Rest Encryption



#### NOTE

At-rest encryption is marked as Technology Preview for this release. For more information on the support scope for features marked as technology previews, see <https://access.redhat.com/support/offerings/techpreview/>

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.

Even with encryption, it is important to use a different method to dispose of the root disk of the proxy, as access to both this and the storage node disks may allow an attacker to decrypt the data.

The following changes need to be made on the overcloud.

To enable encryption of swift object stores:

1. In `/etc/swift/proxy-server.conf`, add these lines, replacing the value of **encryption\_root\_secret** with your encryption key:

```
[pipeline:main]
pipeline = catch_errors healthcheck proxy-logging cache ratelimit
bulk tempurl formpost authtoken keystone staticweb keymaster
encryption proxy-logging proxy-server
```

```
[filter:keymaster]
use = egg:swift#keymaster
encryption_root_secret =

[filter:encryption]
use = egg:swift#encryption
# disable_encryption = False
```

- Restart the proxy service for the changes to come into effect:

```
# systemctl restart openstack-swift-proxy.service
```


To verify that objects are now being encrypted:

- Add a new object to test:

```
# swift upload container1 testobj
```

- Verify the data is encrypted by viewing it with `cat`. For example:

```
# cat
/srv/node/sdb/objects/603/424/35c4ea21cc96663792465462570b2424/14773
99821.58002.data
```



#### NOTE

Changing the encryption key and then attempting to download previously uploaded files will cause an error and may result in data loss.

### 3.1.5. 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 Custom Role File](#) section has information on roles.

#### 3.1.5.1. Creating the `roles_data.yaml` File

- Copy the `roles_data.yaml` from `/usr/share/openstack-tripleo-heat-templates`.
- Edit the new file.
- Remove unneeded controller roles, for example `Aodh*`, `Ceilometer*`, `Ceph*`, `Cinder*`, `Glance*`, `Heat*`, `Ironic*`, `Manila*`, `Mistral*`, `Nova*`, `Octavia*`, `Swift*`.
- Locate the `ObjectStorage` role within `roles_data.yaml`.
- Copy this role to a new role within that same file and name it `ObjectProxy`.
- 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

```

### 3.1.5.2. 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. 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.

### 3.2.1. Create 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
<b>Private</b>	Limits access to a user in the current project.
<b>Public</b>	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 one 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. See [Section 3.1.1.1, “Configure an Erasure Coding Policy”](#) for a sample policy that uses erasure coding.
- *CONTAINERNAME* is the name of the container.

### 3.2.2. Create 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**.

### 3.2.3. Delete 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 3.2.6, “Delete an Object”](#)).
3. Select **Delete Container** in the container’s arrow menu.

4. Click **Delete Container** to confirm the container's removal.

### 3.2.4. Upload 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**.

### 3.2.5. Copy 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**.

### 3.2.6. Delete 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. FILE SHARES

The OpenStack Shared File Systems service (**openstack-manila**) provides the means to easily provision shared file systems that can be consumed by multiple instances. In the past, OpenStack users needed to manually deploy shared file systems before mounting them on instances. The OpenStack Shared File Systems service, on the other hand, allows users to easily provision shares from a pre-configured storage pool, ready to be mounted securely. This pool, in turn, can be independently managed and scaled to meet demand.

### 4.1. BACK ENDS

When deploying Red Hat OpenStack Platform for the environment, we recommend using the director. Doing so helps ensure the proper configuration of each service, including the Shared File System service (and, by extension, its back end). The director also has several integrated back end configurations:

- [NetApp](#)
- [CephFS](#) (Technology Preview)

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

### 4.2. CREATE AND MANAGE SHARE TYPES

When creating a share on a non-default back end, you need to explicitly specify which back end to use. To make the process seamless for users, create a *share type* and associate it with the **share\_backend\_name** value of each back end. To create a share type named *SHARETYPE*, run the following as an OpenStack admin:

```
# manila type-create --is_public false SHARETYPE DHSS
```

*DHSS* (or *driver handles share servers*) specifies whether the share type will use drivers that handle share servers. This should match the value of **driver\_handles\_share\_servers** in your back end definition (**true** or **false**).

The **--is\_public false** option indicates that the share type should not be accessible to the public. If you want the share type to be public, remove this option.

For example, to create a non-public share type named **general** based on a back end configured with **driver\_handles\_share\_servers=false**:

```
# manila type-create --is_public false general false
```

After creating the share type, you can map it to a back end. To do so:

```
# manila type-key SHARETYPE set share_backend_name='SHAREBACKEND'
```

Replace *SHAREBACKEND* with the name of the back end — specifically, its **share\_backend\_name** value. With this, the Shared File System Service will use the *SHAREBACKEND* back end whenever you invoke *SHARETYPE* during share creation (as in [Section 4.3, “Create a Share”](#)). For example, to map the share type **general** to the back end **cdotSingleSVM**:

```
# manila type-key general set share_backend_name='cdotSingleSVM'
```

## TIP

As an OpenStack admin, you can prevent users from creating public share types through the dashboard. See [Section 4.9, “Disable Creation of Public Shares and Share Types Through the Dashboard”](#) for related information.

With this release, the Shared File System Service only supports NetApp storage controllers as back ends. See [Creating and Defining Manila Share Types](#) for information on *extra specs* supported by NetApp.

## 4.3. CREATE A SHARE

To create a share, log in to the Shared File System service host and run:

```
# manila create --share-type SHARETYPE --name SHARENAME PROTO GB
```

Where:

- *SHARETYPE* applies settings associated with the specified share type.
- *SHARENAME* is the name of the share.
- *PROTO* is the share protocol you want to use.
- *GB* is the size of the share, in GB.

For example, in [Section 4.2, “Create and Manage Share Types”](#) we created a share type named **general**, which we mapped to the back end **cdotSingleSVM**. To create a 1 GB NFS share named **share-01** on the **cdotSingleSVM** back end, run:

```
# manila create --share-type general --name share-01 nfs 10
+-----+-----+
| Property          | Value                                |
+-----+-----+
| status            | creating                             |
| description       | None                                  |
| availability_zone | nova                                  |
| share_network_id  | None                                  |
| export_locations  | []                                    |
| share_server_id   | None                                  |
| host              | None                                  |
| snapshot_id       | None                                  |
| is_public         | False                                 |
| id                | d760eee8-1d91-48c4-8f9a-ad07072e17a2 |
| size              | 10                                    |
| name              | share-01                              |
| share_type        | 8245657b-ab9e-4db1-8224-451c32d6b5ea |
| created_at        | 2015-09-29T16:27:54.092272          |
| export_location   | None                                  |
| share_proto       | NFS                                    |
+-----+-----+
```

```
| project_id          | a19dc7ec562c4ed48cea58d22eb0d3c7 |
| metadata           | {}                                 |
+-----+-----+
```

## 4.4. LIST SHARES AND EXPORT INFORMATION

To verify that the shares were created successfully:

```
# manila list
+-----+-----+-----+-----+
| ID          | Name      | ... | Status    ...
+-----+-----+-----+-----+
| d760eee8-1d91-48c4-8f9a-ad07072e17a2 | share-01 | ... | available ...
+-----+-----+-----+-----+
```

The **manila list** command will also display the **export location** of the share:

```
+-----+-----+
| Export location          ...
+-----+-----+
| 10.70.37.46:/manila-nfs-volume-01/share-d760eee8-1d91-...
+-----+-----+
```

This information will be used later when mounting the share ([Section 4.7, “Mount a Share on an Instance”](#)).

## 4.5. GRANT SHARE ACCESS

Before you can mount a share on an instance, grant the instance access to the share first:

```
# manila access-allow SHAREID IDENT IDENTKEY
```

Where:

- *SHAREID* is the ID of the share created in [Section 4.3, “Create a Share”](#).
- *IDENT* is the method that the Shared File System service should use to authenticate a share user or instance.
- The *IDENTKEY* varies depending on what identifying method you choose as *IDENT*:
  - **cert**: this method is used for authenticating an instance through TLS certificates.
  - **user**: use this to authenticate by user or group name.
  - **ip**: use this to authenticate an instance through its IP address.

For example, to grant read-write access to an instance (identified by the IP **10.70.36.85**), run:

```
# manila access-allow d760eee8-1d91-48c4-8f9a-ad07072e17a2 ip 10.70.36.85
+-----+-----+-----+
| Property      | Value          |
+-----+-----+-----+
| share_id     | d760eee8-1d91-48c4-8f9a-ad07072e17a2 |
```

```
| deleted      | False |
| created_at   | 2015-09-29T16:35:33.862114 |
| updated_at   | None |
| access_type  | ip |
| access_to    | 10.70.36.85 |
| access_level | rw |
| state        | new |
| deleted_at   | None |
| id           | b4e990d7-e9d1-4801-bcbe-a860fc1401d1 |
+-----+-----+
```

Note that access to the share has its own ID (*ACCESSID*), **b4e990d7-e9d1-4801-bcbe-a860fc1401d1**.

To verify that the access configuration was successful:

```
# manila access-list d760eee8-1d91-48c4-8f9a-ad07072e17a2
+-----+-----+-----+-----+
| id           | access type | access to | access level | ...
+-----+-----+-----+-----+
| b4e990d7-e9d1-4801-bcbe-... | ip         | 10.70.36.85 | rw         | ...
+-----+-----+-----+-----+
```

## 4.6. REVOKE ACCESS TO A SHARE

To revoke previously-granted access to a share, you need to delete the access to the share:

```
# manila access-deny SHAREID ACCESSID
```

For example, to revoke the access granted earlier in [Section 4.5, “Grant Share Access”](#):

```
# manila access-list d760eee8-1d91-48c4-8f9a-ad07072e17a2
+-----+-----+-----+-----+
| id           | access type | access to | access level | ...
+-----+-----+-----+-----+
| b4e990d7-e9d1-4801-bcbe-... | ip         | 10.70.36.85 | rw         | ...
+-----+-----+-----+-----+
# manila access-deny d760eee8-1d91-48c4-8f9a-ad07072e17a2 b4e990d7-e9d1-4801-bcbe-a860fc1401d1
```

At this point, the instance will no longer be able to use the mounted share.

## 4.7. MOUNT A SHARE ON AN INSTANCE

After configuring the share to authenticate an instance, you can then mount the share. For example, to mount the share from [Section 4.3, “Create a Share”](#) to /mnt on the instance from [Section 4.5, “Grant Share Access”](#), log in to the instance and mount as normal:

```
# ssh root@10.70.36.85
# mount -t nfs -o vers=3 10.70.37.46:/manila-nfs-volume-01/share-d760eee8-1d91-48c4-8f9a-ad07072e17a2 /mnt
```

See [Section 4.4, “List Shares and Export Information”](#) to learn how to view a share’s export information.

## 4.8. DELETE A SHARE

To delete a share:

```
# manila delete SHAREID
```

For example:

```
# manila delete d760eee8-1d91-48c4-8f9a-ad07072e17a2
```

## 4.9. DISABLE CREATION OF PUBLIC SHARES AND SHARE TYPES THROUGH THE DASHBOARD

The dashboard allows users to set shares or share types as *public* during creation through checkboxes. When creating a share, the **Make visible for all** checkbox sets the share as public:



# Create Share

---

Share Name \*

Description

Share Protocol \*

Size (GiB) \*

Share Type \*

Availability Zone

Metadata

Make visible for all ?

The share type creation interface features a **Public** checkbox that does the same:

# Create Share Type

Name \*

Driver handles share servers \*

Extra specs

Public 

You can disable these checkboxes to limit the number of public shares and share types created by users. To do so, log in to a Controller node and add the following to `/etc/openstack-dashboard/local_settings`:

```
OPENSTACK_MANILA_FEATURES = {
    'enable_public_shares': False
    'enable_public_share_type_creation': False, }
```

The `'enable_public_shares': False` setting disables the **Make visible for all** checkbox. Likewise, `'enable_public_share_type_creation': False` disables the **Public** checkbox in the share type creation interface.

If you deployed OpenStack with High Availability, apply the same change to all Controller nodes. Then, restart the **httpd** service from any Controller node:

```
# pcs resource restart httpd-clone
```

Otherwise, restart the **httpd** service using:

```
# systemctl restart httpd
```

To view your changes, reload your dashboard. The checkboxes should no longer appear in their respective interfaces.

**IMPORTANT**

Disabling these checkboxes does not prevent users or OpenStack admins from setting shares or share types to public *after* creating them.