



Red Hat OpenStack Platform 16.0

Undercloud and Control Plane Back Up and Restore

Procedures for backing up and restoring the undercloud and the overcloud control plane during updates and upgrades

Red Hat OpenStack Platform 16.0 Undercloud and Control Plane Back Up and Restore

Procedures for backing up and restoring the undercloud and the overcloud control plane during updates and upgrades

Legal Notice

Copyright © 2020 Red Hat, Inc.

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

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

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

This guide explains how to install and configure Relax-and-Recover (ReaR) on the undercloud and overcloud control plane nodes; how to back up the undercloud and Control Plane nodes before updates and upgrades; and, how to restore the undercloud and Control Plane nodes if an error occurs while performing updates or upgrades.

Table of Contents

CHAPTER 1. INTRODUCTION TO UNDERCLOUD AND CONTROL PLANE BACK UP AND RESTORE	3
1.1. BACKGROUND	3
1.2. BACKUP MANAGEMENT OPTIONS	3
CHAPTER 2. PREPARING THE BACKUP NODE	5
2.1. PREPARING THE NFS SERVER	5
2.2. CREATING AND EXPORTING THE BACKUP DIRECTORY	5
CHAPTER 3. INSTALLING AND CONFIGURING REAR	7
3.1. INSTALLING THE REQUIRED PACKAGES	7
3.2. CREATING THE CONFIGURATION FILES	7
CHAPTER 4. EXECUTING THE BACKUP PROCEDURE	9
4.1. PERFORMING PREREQUISITE TASKS BEFORE BACKING UP THE UNDERCLOUD	9
4.2. BACKING UP THE UNDERCLOUD	9
4.3. BACKING UP THE CONTROL PLANE	10
CHAPTER 5. EXECUTING THE RESTORE PROCEDURE	13
5.1. RESTORING THE UNDERCLOUD	13
5.2. RESTORING THE CONTROL PLANE	14
CHAPTER 6. BACKING UP AND RESTORING THE UNDERCLOUD AND CONTROL PLANE NODES WITH COLLOCATED CEPH MONITORS	15

CHAPTER 1. INTRODUCTION TO UNDERCLOUD AND CONTROL PLANE BACK UP AND RESTORE

The Undercloud and Control Plane Back Up and Restore procedure provides steps for backing up the state of the Red Hat OpenStack Platform 16.0 undercloud and overcloud Controller nodes, hereinafter referred to as control plane nodes, before updates and upgrades. Use the procedure to restore the undercloud and the overcloud control plane nodes to their previous state if an error occurs during an update or upgrade.

1.1. BACKGROUND

The Undercloud and Control Plane Back Up and Restore procedure uses the open source Relax and Recover (ReaR) disaster recovery solution, written in Bash. ReaR creates a bootable image consisting of the latest state of an undercloud or a Control Plane node. ReaR also enables a system administrator to select files for backup.

ReaR supports numerous boot media formats, including:

- ISO
- USB
- eSATA
- PXE

The examples in this document were tested using the **ISO** boot format.

ReaR can transport the boot images using multiple protocols, including:

- HTTP/HTTPS
- SSH/SCP
- FTP/SFTP
- NFS
- CIFS (SMB)

For the purposes of backing up and restoring the Red Hat OpenStack Platform 16.0 undercloud and overcloud Control Plane nodes, the examples in this document were tested using NFS.

1.2. BACKUP MANAGEMENT OPTIONS

ReaR can use both internal and external backup management options.

Internal backup management

Internal backup options include:

- **tar**
- **rsync**

External backup management

External backup management options include both open source and proprietary solutions. Open source solutions include:

- Bacula
- Bareos

Proprietary solutions include:

- EMC NetWorker (Legato)
- HP DataProtector
- IBM Tivoli Storage Manager (TSM)
- Symantec NetBackup

CHAPTER 2. PREPARING THE BACKUP NODE

Before you back up the undercloud or control plane nodes, prepare the backup node to accept the backup images.

2.1. PREPARING THE NFS SERVER

ReaR can use multiple transport methods. Red Hat supports back up and restore with ReaR using NFS.

1. Install the NFS server on the backup node.

```
[root@backup ~]# dnf install -y nfs-utils
```

2. Add the NFS service to the firewall to ensure ports **111** and **2049** are open. For example:

```
[root@backup ~]# firewall-cmd --add-service=nfs
[root@backup ~]# firewall-cmd --add-service=nfs --permanent
```

3. Enable the NFS server and start it.

```
[root@backup ~]# systemctl enable nfs-server
[root@backup ~]# systemctl restart nfs-server
```

2.2. CREATING AND EXPORTING THE BACKUP DIRECTORY

To copy backup ISO images from the undercloud or control plane nodes to the backup node, you must create a backup directory.

Prerequisites

- You installed and enabled the NFS server. For more information, see [Preparing the NFS server](#).

Procedure

1. Create the backup directory:

```
[root@backup ~]# mkdir /ctl_plane_backups
```

2. Export the directory. Replace **<ip-addr>/24** with the IP address and subnet mask of the network:

```
[root@backup ~]# cat >> /etc/exports << EOF
/ctl_plane_backups <ip-addr>/24(rw,sync,no_root_squash,no_subtree_check) /
EOF
```

The entries in the **/etc/exports** file are in a space-delimited list. If the undercloud and the overcloud control plane nodes use different networks or subnets, repeat this step for each network or subnet, as shown in this example:

```
cat >> /etc/exports << EOF
/ctl_plane_backups 192.168.24.0/24(rw,sync,no_root_squash,no_subtree_check) /
ctl_plane_backups 10.0.0.0/24(rw,sync,no_root_squash,no_subtree_check) /
```

```
ctl_plane_backups 172.16.0.0/24(rw,sync,no_root_squash,no_subtree_check)
EOF
```

3. Restart the NFS server:

```
[root@backup ~]# systemctl restart nfs-server
```

4. Verify that the entries are correctly configured in the NFS server:

```
[root@backup ~]# showmount -e `hostname`
```

CHAPTER 3. INSTALLING AND CONFIGURING REAR

Before you back up the undercloud and the overcloud control plane nodes, you must first install and configure Relax and Recover (ReaR) on the undercloud and on each control plane node.

3.1. INSTALLING THE REQUIRED PACKAGES

You must install the Relax and Recover (ReaR) packages and packages for generating ISO images on the undercloud node and on each control plane node.

Procedure

1. Install the required packages on the undercloud and on each control plane node. For example:

```
[root@controller-x ~]# dnf install rear genisoimage nfs-utils -y
```

2. Create a backup directory on the undercloud and on each control plane node. For example:

```
[root@controller-x ~]# mkdir -p /ctl_plane_backups
```

3. Mount the **ctl_plane_backups** NFS directory from the backup node running NFS on the undercloud and on each control plane node. For example:

```
[root@controller-x ~]# mount -t nfs <ip-addr>:/ctl_plane_backups /ctl_plane_backups
```

Replace **<ip-addr>** with the IP address of the backup node running the NFS server.


3.2. CREATING THE CONFIGURATION FILES

As the **root** user on the undercloud and on each control plane node, perform the following steps:

1. Create the ReaR configuration file:

```
[root@controller-x ~]# mkdir -p /etc/rear
[root@controller-x ~]# tee -a "/etc/rear/local.conf" > /dev/null <<'EOF'
OUTPUT=ISO
OUTPUT_URL=nfs://<ip-addr>/ctl_plane_backups
ISO_PREFIX=<SERVER_NAME-X>
BACKUP=NETFS
BACKUP_PROG_COMPRESS_OPTIONS=( --gzip )
BACKUP_PROG_COMPRESS_SUFFIX=".gz"
BACKUP_PROG_EXCLUDE=( '/tmp/*' '/data/*' )
BACKUP_URL=nfs://<ip-addr>/ctl_plane_backups
BACKUP_PROG_EXCLUDE=( "${BACKUP_PROG_EXCLUDE[@]}" '/media' '/var/tmp'
'/var/crash' )
BACKUP_PROG_OPTIONS+=( --anchored --xattrs-include='*.*' --xattrs )
EOF
```

Replace **<SERVER_NAME-X>** with the hostname of the node. For example, if the node hostname is **controller-0**, replace **<SERVER_NAME-X>** with **controller-0**. Replace **<ip-addr>** with the IP address of the backup node running the NFS server configured in [Chapter 2, Preparing the backup node](#).

**IMPORTANT**

If the undercloud or control plane nodes use UEFI as their boot mode, you must add **USING_UEFI_BOOTLOADER=1** to the configuration file too.

2. Create the **rescue.conf** file:

```
[root@controller-x ~]# tee -a "/etc/rear/rescue.conf" > /dev/null <<'EOF'  
BACKUP_PROG_OPTIONS+=( --anchored --xattrs-include='*.*' --xattrs )  
EOF
```

CHAPTER 4. EXECUTING THE BACKUP PROCEDURE

Before you perform a fast forward upgrade, back up the undercloud and the overcloud control plane nodes so that you can restore them to their previous state if an error occurs.



NOTE

Before you back up the undercloud and overcloud, ensure that you do not perform any operations on the overcloud from the undercloud.

4.1. PERFORMING PREREQUISITE TASKS BEFORE BACKING UP THE UNDERCLOUD

Do not perform an undercloud backup when you deploy the undercloud or when you make changes to an existing undercloud. To prevent data corruptions, confirm that there are no stack failures, ongoing tasks, and that all OpenStack services except for **mariadb** are stopped before you back up the undercloud node.

Procedure

1. List failures for all available stacks:

```
(undercloud) [stack@undercloud-0 ~]$ source stackrc && for i in `openstack stack list -c 'Stack Name' -f value`; do openstack stack failures list $i; done
```

2. Verify that there are no ongoing tasks in the cloud:

```
(undercloud) [stack@undercloud-0 ~]$ openstack stack list --nested | grep -v "_COMPLETE"
```

If the command returns no results, there are no ongoing tasks.

3. Stop all OpenStack services in the cloud:

```
# systemctl stop tripleo_*
```

4. Start the **tripleo_mysql** service:

```
# systemctl start tripleo_mysql
```

5. Verify that the **tripleo_mysql** service is running:

```
# systemctl status tripleo_mysql
```

4.2. BACKING UP THE UNDERCLOUD

To back up the undercloud node, you must log in as the root user on the undercloud node. As a precaution, you must back up the database to ensure that you can restore it.

Prerequisites

- You have created and exported the backup directory. For more information, see [Creating and exporting the backup directory](#).
- You have performed prerequisite tasks before backing up the undercloud. For more information, see [Performing prerequisite tasks before backing up the undercloud](#).
- You have installed and configured ReaR on the undercloud node. For more information, see [Install and Configure ReaR](#).

Procedure

1. Locate the database password.

```
[root@undercloud stack]# PASSWORD=$(/bin/hiera -c /etc/puppet/hiera.yaml
mysql::server::root_password)
```

2. Back up the databases:

```
# podman exec mysql bash -c "mysql -uroot -p$PASSWORD -s -N -e \"SELECT
CONCAT(\"\\\"SHOW GRANTS FOR \",user,\"@\",host,\";\\\") FROM mysql.user where
(length(user) > 0 and user NOT LIKE 'root')\" | xargs -n1 mysql -uroot -p$PASSWORD -s -N -
e | sed 's/$/;/' " > openstack-backup-mysql-grants.sql
```

```
# podman exec mysql bash -c "mysql -uroot -p$PASSWORD -s -N -e \"select distinct
table_schema from information_schema.tables where engine='innodb' and table_schema !=
'mysql';\" | xargs mysqldump -uroot -p$PASSWORD --single-transaction --databases" >
openstack-backup-mysql.sql
```

3. Stop the **mariadb** database service:

```
[root@undercloud stack]# systemctl stop tripleo_mysql
```

4. Create the backup:

```
[root@undercloud stack]# rear -d -v mkbackup
```

You can find the backup ISO file that you create with ReaR on the backup node in the **/ctl_plane_backups** directory.

4.3. BACKING UP THE CONTROL PLANE

To back up the control plane, you must first stop the pacemaker cluster and all containers operating on the control plane nodes. Do not operate the stack to ensure state consistency. After you complete the backup procedure, start the pacemaker cluster and the containers.

As a precaution, you must back up the database to ensure that you can restore the database after you restart the pacemaker cluster and containers.

Back up the control plane nodes simultaneously.

Prerequisites

- You have created and exported the backup directory. For more information, see [Creating and exporting the backup directory](#).
- You have installed and configured ReaR on the undercloud node. For more information, see [Install and Configure ReaR](#).

Procedure

1. Locate the database password:


```
[heat-admin@overcloud-controller-x ~]# PASSWORD=$(/bin/hiera -c /etc/puppet/hiera.yaml
mysql::server::root_password)
```

2. Back up the databases:

```
[heat-admin@overcloud-controller-x ~]# podman exec galera-bundle-podman-X bash -c
"mysql -uroot -p$PASSWORD -s -N -e \"SELECT CONCAT('\\\\\\\"SHOW GRANTS FOR
\",user,\"@\",host,\";\\\\\\\"') FROM mysql.user where (length(user) > 0 and user NOT LIKE
'root')\" | xargs -n1 mysql -uroot -p$PASSWORD -s -N -e | sed 's/$/;/' " > openstack-backup-
mysql-grants.sql
```

```
[heat-admin@overcloud-controller-x ~]# podman exec galera-bundle-podman-X bash -c
"mysql -uroot -p$PASSWORD -s -N -e \"select distinct table_schema from
information_schema.tables where engine='innodb' and table_schema != 'mysql';\" | xargs
mysqldump -uroot -p$PASSWORD --single-transaction --databases" > openstack-backup-
mysql.sql
```

3. On one of control plane nodes, stop the pacemaker cluster:



IMPORTANT

Do not operate the stack. When you stop the pacemaker cluster and the containers, this results in the temporary interruption of control plane services to Compute nodes. There is also disruption to network connectivity, Ceph, and the NFS data plane service. You cannot create instances, migrate instances, authenticate requests, or monitor the health of the cluster until the pacemaker cluster and the containers return to service following the final step of this procedure.

```
[heat-admin@overcloud-controller-x ~]# pcs cluster stop --all
```

4. On each control plane node, stop the containers.

- a. Stop the containers:

```
[heat-admin@overcloud-controller-x ~]# systemctl stop tripleo_*
```

- b. Stop the **ceph-mon@controller.service** container:

```
[heat-admin@overcloud-controller-x ~]# sudo systemctl stop ceph-mon@$(hostname -s)
```

- c. Stop the **ceph-mgr@controller.service** container:

```
[heat-admin@overcloud-controller-x ~]# sudo systemctl stop ceph-mgr@$(hostname -s)
```

5. To back up the control plane, run the following command as **root** in the command line interface of each control plane node:

```
[heat-admin@overcloud-controller-x ~]# rear -d -v mkbackup
```

You can find the backup ISO file that you create with ReaR on the backup node in the **/ctl_plane_backups** directory.



NOTE

When you execute the backup command, you might see warning messages regarding the **tar** command and sockets that are ignored during the tar process, similar to the following:

```
WARNING: tar ended with return code 1 and below output:
```

```
---snip---
```

```
tar: /var/spool/postfix/public/qmgr: socket ignored
```

```
...
```

```
...
```

This message indicates that files have been modified during the archiving process and the backup might be inconsistent. Relax-and-Recover continues to operate, however, it is important that you verify the backup to ensure that you can use this backup to recover your system.

6. When the backup procedure generates ISO images for each of the control plane nodes, restart the pacemaker cluster and the containers:

On one of the control plane nodes, enter the following command:

```
[heat-admin@overcloud-controller-x ~]# pcs cluster start --all
```

7. On each control plane node, start the containers.

- a. Start the **ceph-mon@controller.service** container:

```
[heat-admin@overcloud-controller-x ~]# systemctl start ceph-mon@$(hostname -s)
```

- b. Start the **ceph-mgr@controller.service** container:

```
[heat-admin@overcloud-controller-x ~]# systemctl start ceph-mgr@$(hostname -s)
```


CHAPTER 5. EXECUTING THE RESTORE PROCEDURE

If an error occurs during an update or upgrade, you can restore either the undercloud or overcloud control plane nodes or both so that they assume their previous state.

Use the following general steps:

1. Burn the bootable ISO image to a DVD or load it through ILO remote access.
2. Boot the node that requires restoration from the recovery medium.
3. Select **Recover <hostname>**, where **<hostname>** is the name of the node to restore.
4. Log in as user **root**.
5. Recover the backup.

5.1. RESTORING THE UNDERCLOUD

If an error occurs during a fast-forward upgrade, you can restore the undercloud node to its previously saved state using the ISO image created using the [Section 4.2, "Backing up the undercloud"](#) procedure. The backup procedure stores the ISO images on the backup node in the folders created during the [Section 2.2, "Creating and exporting the backup directory"](#) step.

Procedure

1. Shutdown the undercloud node. Ensure that the undercloud node is shutdown completely before you proceed.
2. Restore the undercloud node by booting it with the ISO image created during the backup process. The ISO image is located under the **/ctl_plane_backups** directory of the Backup node.
3. When the **Relax-and-Recover** boot menu appears, select **Recover <Undercloud Node>** where **<Undercloud Node>** is the name of the undercloud node.
4. Log in as user **root**.
The following message displays:

```
Welcome to Relax-and-Recover. Run "rear recover" to restore your system!
RESCUE <Undercloud Node>:~ # rear recover
```

The image restore progresses quickly. When it is complete, the console echoes the following message:

```
Finished recovering your system
Exiting rear recover
Running exit tasks
```

5. When the command line interface is available, the image is restored. Switch the node off.

```
RESCUE <Undercloud Node>:~ # poweroff
```

On boot up, the node resumes with its previous state.

5.2. RESTORING THE CONTROL PLANE

If an error occurs during a fast-forward upgrade, you can use the ISO images created using the [Section 4.3, "Backing up the control plane"](#) procedure to restore the control plane nodes to their previously saved state. To restore the control plane, you must restore all control plane nodes to the previous state to ensure state consistency.



NOTE

Red Hat supports backups of Red Hat OpenStack Platform with native SDNs, such as Open vSwitch (OVS) and the default Open Virtual Network (OVN). For information about third-party SDNs, refer to the third-party SDN documentation.

1. Shutdown each control plane node. Ensure that the control plane nodes are shutdown completely before you proceed.
2. Restore the control plane nodes by booting them with the ISO image that you created during the backup process. The ISO images are located under the `/ctl_plane_backups` directory of the backup node.
3. When the **Relax-and-Recover** boot menu appears, select **Recover <Control Plane Node>** where `<Control Plane Node>` is the name of the control plane node. The following message displays:

```
Welcome to Relax-and-Recover. Run "rear recover" to restore your system!
RESCUE <Control Plane Node>:~ # rear recover
```

The image restore progresses quickly. When the restore completes, the console echoes the following message:

```
Finished recovering your system
Exiting rear recover
Running exit tasks
```

When the command line interface is available, the image is restored. Switch the node off.

```
RESCUE <Control Plane Node>:~ # poweroff
```

Set the boot sequence to the normal boot device. On boot up, the node resumes with its previous state.

4. To ensure that the services are running correctly, check the status of pacemaker. Log in to a controller as **root** user and run the following command:


```
# pcs status
```
5. To view the status of the overcloud, use Tempest. For more information about Tempest, see Chapter 4 of the [OpenStack Integration Test Suite Guide](#).

CHAPTER 6. BACKING UP AND RESTORING THE UNDERCLOUD AND CONTROL PLANE NODES WITH COLLOCATED CEPH MONITORS

If an error occurs during an update or upgrade, you can use ReaR backups to restore either the undercloud or overcloud control plane nodes, or both, to their previous state.

Prerequisites

- Install and configure ReaR. For more information, see [Install and configure ReaR](#).
- Prepare the backup node. For more information, see [Prepare the backup node](#).
- Execute the backup procedure. For more information, see [Execute the backup procedure](#).

Procedure

1. On the backup node, export the NFS directory to host the Ceph backups. Replace **<IP_ADDRESS/24>** with the IP address and subnet mask of the network:

```
[root@backup ~]# cat >> /etc/exports << EOF
/ceph_backups <IP_ADDRESS/24>(rw,sync,no_root_squash,no_subtree_check)
EOF
```

2. On the undercloud node, source the undercloud credentials and run the following script:

```
# source stackrc
```

```
#!/bin/bash
for i in `openstack server list -c Name -c Networks -f value | grep controller | awk -F=' ' '{print $2}' | awk -F' ' '{print $1}'`; do ssh -q heat-admin@$i 'sudo systemctl stop ceph-mon@$(hostname -s) ceph-mgr@$(hostname -s)'; done
```

To verify that the **ceph-mgr@controller.service** container has stopped, enter the following command:

```
[heat-admin@overcloud-controller-x ~]# sudo podman ps | grep ceph
```

3. On the undercloud node, source the undercloud credentials and run the following script. Replace **<BACKUP_NODE_IP_ADDRESS>** with the IP address of the backup node:

```
# source stackrc
```

```
#!/bin/bash
for i in `openstack server list -c Name -c Networks -f value | grep controller | awk -F=' ' '{print $2}' | awk -F' ' '{print $1}'`; do ssh -q heat-admin@$i 'sudo mkdir /ceph_backups'; done
```

```
#!/bin/bash
for i in `openstack server list -c Name -c Networks -f value | grep controller | awk -F=' ' '{print $2}' | awk -F' ' '{print $1}'`; do ssh -q heat-admin@$i 'sudo mount -t nfs <BACKUP_NODE_IP_ADDRESS>:/ceph_backups /ceph_backups'; done
```

```
#!/bin/bash
for i in `openstack server list -c Name -c Networks -f value | grep controller | awk -F=' ' '{print $2}' | awk -F' ' '{print $1}'`; do ssh -q heat-admin@$i 'sudo mkdir /ceph_backups/${hostname -s}'; done

#!/bin/bash
for i in `openstack server list -c Name -c Networks -f value | grep controller | awk -F=' ' '{print $2}' | awk -F' ' '{print $1}'`; do ssh -q heat-admin@$i 'sudo tar -zcv --xattrs-include=*. * --xattrs --xattrs-include=security.capability --xattrs-include=security.selinux --acls -f /ceph_backups/${hostname -s}/${hostname -s}.tar.gz /var/lib/ceph'; done
```

4. On the node that you want to restore, complete the following tasks:
 - a. Power off the node before you proceed.
 - b. Restore the node with the ReaR backup file that you have created during the backup process. The file is located in the **/ceph_backups** directory of the backup node.
 - c. From the **Relax-and-Recover** boot menu, select **Recover <CONTROL_PLANE_NODE>**, where **<CONTROL_PLANE_NODE>** is the name of the control plane node.
 - d. At the prompt, enter the following command:

```
RESCUE <CONTROL_PLANE_NODE> :~ # rear recover
```

When the image restoration process completes, the console displays the following message:

```
Finished recovering your system
Exiting rear recover
Running exit tasks
```

5. For the node that you want to restore, copy the Ceph backup from the **/ceph_backups** directory into the **/var/lib/ceph** directory:
 - a. Identify the system mount points:

```
RESCUE <CONTROL_PLANE_NODE>:~# df -h
Filesystem      Size  Used Avail Use% Mounted on
devtmpfs        16G   0  16G   0% /dev
tmpfs           16G   0  16G   0% /dev/shm
tmpfs           16G  8.4M  16G   1% /run
tmpfs           16G   0  16G   0% /sys/fs/cgroup
/dev/vda2       30G  13G  18G  41% /mnt/local
```

The **/dev/vda2** file system is mounted on **/mnt/local**.

- b. Create a temporary directory:

```
RESCUE <CONTROL_PLANE_NODE>:~ # mkdir /tmp/restore
RESCUE <CONTROL_PLANE_NODE>:~ # mount -v -t nfs -o rw,noatime
<BACKUP_NODE_IP_ADDRESS>:/ceph_backups /tmp/restore/
```

- c. On the control plane node, remove the existing **/var/lib/ceph** directory:

```
RESCUE <CONTROL_PLANE_NODE>:~ # rm -rf /mnt/local/var/lib/ceph/*
```

- d. Restore the previous Ceph maps. Replace **<CONTROL_PLANE_NODE>** with the name of your control plane node:

```
RESCUE <CONTROL_PLANE_NODE>:~ # tar -xvC /mnt/local/ -f  
/tmp/restore/<CONTROL_PLANE_NODE>/<CONTROL_PLANE_NODE>.tar.gz --xattrs -  
-xattrs-include='*.*' var/lib/ceph
```

- e. Verify that the files are restored:

```
RESCUE <CONTROL_PLANE_NODE>:~ # ls -l  
total 0  
drwxr-xr-x 2 root 107 26 Jun 18 18:52 bootstrap-mds  
drwxr-xr-x 2 root 107 26 Jun 18 18:52 bootstrap-osd  
drwxr-xr-x 2 root 107 26 Jun 18 18:52 bootstrap-rbd  
drwxr-xr-x 2 root 107 26 Jun 18 18:52 bootstrap-rgw  
drwxr-xr-x 3 root 107 31 Jun 18 18:52 mds  
drwxr-xr-x 3 root 107 31 Jun 18 18:52 mgr  
drwxr-xr-x 3 root 107 31 Jun 18 18:52 mon  
drwxr-xr-x 2 root 107 6 Jun 18 18:52 osd  
drwxr-xr-x 3 root 107 35 Jun 18 18:52 radosgw  
drwxr-xr-x 2 root 107 6 Jun 18 18:52 tmp
```

6. Power off the node:

```
RESCUE <CONTROL_PLANE_NODE> :~ # poweroff
```

7. Power on the node. The node resumes its previous state.