Red Hat OpenStack Platform 11

Red Hat Ceph Storage for the Overcloud

Configuring an Overcloud to Use Red Hat Ceph Storage

Red Hat OpenStack Platform 11 Red Hat Ceph Storage for the Overcloud

Configuring an Overcloud to Use Red Hat Ceph Storage

OpenStack Team
rhos-docs@redhat.com

Legal Notice

Abstract

This guide provides information on using the Red Hat OpenStack Platform director to create an Overcloud that uses Red Hat Ceph Storage. This includes recommendations for your Red Hat Ceph Storage environment and instructions on how to implement an Overcloud with Ceph Storage nodes.

Table of Contents

CHAPTER 1. INTRODUCTION

Red Hat OpenStack Platform director creates a cloud environment called the **Overcloud**. The director provides the ability to configure extra features for an Overcloud. One of these extra features includes integration with Red Hat Ceph Storage. This includes both Ceph Storage clusters created with the director or existing Ceph Storage clusters. This guide provides information for integrating Ceph Storage into your Overcloud through the director and configuration examples.

1.1. DEFINING CEPH STORAGE

Red Hat Ceph Storage is a distributed data object store designed to provide excellent performance, reliability, and scalability. Distributed object stores are the future of storage, because they accommodate unstructured data, and because clients can use modern object interfaces and legacy interfaces simultaneously. At the heart of every Ceph deployment is the **Ceph Storage Cluster**, which consists of two types of daemons:

Ceph OSD (Object Storage Daemon)

Ceph OSDs store data on behalf of Ceph clients. Additionally, Ceph OSDs utilize the CPU and memory of Ceph nodes to perform data replication, rebalancing, recovery, monitoring and reporting functions.

Ceph Monitor

A Ceph monitor maintains a master copy of the Ceph storage cluster map with the current state of the storage cluster.

For more information about Red Hat Ceph Storage, see the Red Hat Ceph Storage Architecture Guide.

> **IMPORTANT**
>
> This guide only integrates Ceph Block storage and the Ceph Object Gateway (RGW). It does not include Ceph File (CephFS) storage.

1.2. USING CEPH STORAGE IN RED HAT OPENSTACK PLATFORM

Red Hat OpenStack Platform director provides two main methods for integrating Red Hat Ceph Storage into an Overcloud.

Creating an Overcloud with its own Ceph Storage Cluster

The director has the ability to create a Ceph Storage Cluster during the creation on the Overcloud. The director creates a set of Ceph Storage nodes that use the Ceph OSD to store the data. In addition, the director installs the Ceph Monitor service on the Overcloud's Controller nodes. This means if an organization creates an Overcloud with three highly available controller nodes, the Ceph Monitor also becomes a highly available service.

Integrating a Existing Ceph Storage into an Overcloud

If you already have an existing Ceph Storage Cluster, you can integrate this during an Overcloud deployment. This means you manage and scale the cluster outside of the Overcloud configuration.

1.3. SETTING REQUIREMENTS

This guide acts as supplementary information for the Director Installation and Usage guide. This means the Requirements section also applies to this guide. Implement these requirements as necessary.

If using the Red Hat OpenStack Platform director to create Ceph Storage nodes, note the following requirements for these nodes:

Processor

64-bit x86 processor with support for the Intel 64 or AMD64 CPU extensions.

Memory

Memory requirements depend on the amount of storage space. Ideally, use at minimum 1 GB of memory per 1 TB of hard disk space.

Disk Space

Storage requirements depends on the amount of memory. Ideally, use at minimum 1 GB of memory per 1 TB of hard disk space.

Disk Layout

The recommended Red Hat Ceph Storage node configuration requires at least three or more disks in a layout similar to the following:

- **/dev/sda** - The root disk. The director copies the main Overcloud image to the disk.

- **/dev/sdb** - The journal disk. This disk divides into partitions for Ceph OSD journals. For example, /dev/sdb1, /dev/sdb2, /dev/sdb3, and onward. The journal disk is usually a solid state drive (SSD) to aid with system performance.

- **/dev/sdc** and onward - The OSD disks. Use as many disks as necessary for your storage requirements.

IMPORTANT

Erase all existing partitions on the disks targeted for journaling and OSDs before deploying the Overcloud. In addition, the Ceph Storage OSDs and journal disks require GPT disk labels, which you can configure as a part of the deployment. See Section 2.8, "Cleaning Ceph Storage Node Disks" for more information.

Network Interface Cards

A minimum of one 1 Gbps Network Interface Cards, although it is recommended to use at least two NICs in a production environment. Use additional network interface cards for bonded interfaces or to delegate tagged VLAN traffic. It is recommended to use a 10 Gbps interface for storage node, especially if creating an OpenStack Platform environment that serves a high volume of traffic.

Intelligent Platform Management Interface (IPMI)

Each Ceph node requires IPMI functionality on the server's motherboard.

This guide also requires the following:

- An Undercloud host with the Red Hat OpenStack Platform director installed. See Installing the Undercloud.

- Any additional hardware recommendation for Red Hat Ceph Storage. See the Red Hat Ceph Storage Hardware Guide for these recommendations.

IMPORTANT

The Ceph Monitor service is installed on the Overcloud's Controller nodes. This means you must provide adequate resources to alleviate performance issues. Ensure the Controller nodes in your environment use at least 16 GB of RAM for memory and solid-state drive (SSD) storage for the Ceph monitor data.

1.4. DEFINING THE SCENARIOS

This guide uses two scenarios:

- The first scenario creates an Overcloud with a Ceph Storage Cluster. This means the director deploys the Ceph Storage Cluster.

- The second scenario integrates an existing Ceph Storage Cluster with an Overcloud. This means you manage the Ceph Storage Cluster separate from Overcloud management.

CHAPTER 2. CREATING AN OVERCLOUD WITH CEPH STORAGE NODES

This chapter describes how to use the director to create an Overcloud that includes its own Ceph Storage Cluster. For instructions on how to create an Overcloud and integrate it with an existing Ceph Storage Cluster, see Chapter 3, *Integrating an Existing Ceph Storage Cluster with an Overcloud* instead.

The scenario described in this chapter consists of nine nodes in the Overcloud:

- Three Controller nodes with high availability. This includes the Ceph Monitor service on each node.

- Three Red Hat Ceph Storage nodes in a cluster. These nodes contain the Ceph OSD service and act as the actual storage.

- Three Compute nodes.

All machines in this scenario are bare metal systems using IPMI for power management. These nodes do not require an operating system because the director copies a Red Hat Enterprise Linux 7 image to each node.

The director communicates to each node through the Provisioning network during the introspection and provisioning processes. All nodes connect to this network through the native VLAN. For this example, we use 192.0.2.0/24 as the Provisioning subnet with the following IP address assignments:

Node Name	IP Address	MAC Address	IPMI IP Address
Director	192.0.2.1	aa:aa:aa:aa:aa:aa	
Controller 1	DHCP defined	b1:b1:b1:b1:b1:b1	192.0.2.205
Controller 2	DHCP defined	b2:b2:b2:b2:b2:b2	192.0.2.206
Controller 3	DHCP defined	b3:b3:b3:b3:b3:b3	192.0.2.207
Compute 1	DHCP defined	c1:c1:c1:c1:c1:c1	192.0.2.208
Compute 2	DHCP defined	c2:c2:c2:c2:c2:c2	192.0.2.209
Compute 3	DHCP defined	c3:c3:c3:c3:c3:c3	192.0.2.210
Ceph 1	DHCP defined	d1:d1:d1:d1:d1:d1	192.0.2.211
Ceph 2	DHCP defined	d2:d2:d2:d2:d2:d2	192.0.2.212
Ceph 3	DHCP defined	d3:d3:d3:d3:d3:d3	192.0.2.213

2.1. INITIALIZING THE STACK USER

Log into the director host as the **stack** user and run the following command to initialize your director configuration:

```
$ source ~/stackrc
```

This sets up environment variables containing authentication details to access the director's CLI tools.

2.2. REGISTERING NODES

A node definition template (**instackenv.json**) is a JSON format file and contains the hardware and power management details for registering nodes. For example:

```
{
    "nodes":[
        {
            "mac":[
                "b1:b1:b1:b1:b1:b1"
            ],
            "cpu":"4",
            "memory":"6144",
            "disk":"40",
            "arch":"x86_64",
            "pm_type":"pxe_ipmitool",
            "pm_user":"admin",
            "pm_password":"p@55w0rd!",
            "pm_addr":"192.0.2.205"
        },
        {
            "mac":[
                "b2:b2:b2:b2:b2:b2"
            ],
            "cpu":"4",
            "memory":"6144",
            "disk":"40",
            "arch":"x86_64",
            "pm_type":"pxe_ipmitool",
            "pm_user":"admin",
            "pm_password":"p@55w0rd!",
            "pm_addr":"192.0.2.206"
        },
        {
            "mac":[
                "b3:b3:b3:b3:b3:b3"
            ],
            "cpu":"4",
            "memory":"6144",
            "disk":"40",
            "arch":"x86_64",
            "pm_type":"pxe_ipmitool",
            "pm_user":"admin",
            "pm_password":"p@55w0rd!",
            "pm_addr":"192.0.2.207"
        },
        {
            "mac":[
```

```
                "c1:c1:c1:c1:c1:c1"
            ],
            "cpu":"4",
            "memory":"6144",
            "disk":"40",
            "arch":"x86_64",
            "pm_type":"pxe_ipmitool",
            "pm_user":"admin",
            "pm_password":"p@55w0rd!",
            "pm_addr":"192.0.2.208"
        },
        {
            "mac":[
                "c2:c2:c2:c2:c2:c2"
            ],
            "cpu":"4",
            "memory":"6144",
            "disk":"40",
            "arch":"x86_64",
            "pm_type":"pxe_ipmitool",
            "pm_user":"admin",
            "pm_password":"p@55w0rd!",
            "pm_addr":"192.0.2.209"
        },
        {
            "mac":[
                "c3:c3:c3:c3:c3:c3"
            ],
            "cpu":"4",
            "memory":"6144",
            "disk":"40",
            "arch":"x86_64",
            "pm_type":"pxe_ipmitool",
            "pm_user":"admin",
            "pm_password":"p@55w0rd!",
            "pm_addr":"192.0.2.210"
        },
        {
            "mac":[
                "d1:d1:d1:d1:d1:d1"
            ],
            "cpu":"4",
            "memory":"6144",
            "disk":"40",
            "arch":"x86_64",
            "pm_type":"pxe_ipmitool",
            "pm_user":"admin",
            "pm_password":"p@55w0rd!",
            "pm_addr":"192.0.2.211"
        },
        {
            "mac":[
                "d2:d2:d2:d2:d2:d2"
            ],
            "cpu":"4",
            "memory":"6144",
```

```
                "disk":"40",
                "arch":"x86_64",
                "pm_type":"pxe_ipmitool",
                "pm_user":"admin",
                "pm_password":"p@55w0rd!",
                "pm_addr":"192.0.2.212"
        },
        {
                "mac":[
                    "d3:d3:d3:d3:d3:d3"
                ],
                "cpu":"4",
                "memory":"6144",
                "disk":"40",
                "arch":"x86_64",
                "pm_type":"pxe_ipmitool",
                "pm_user":"admin",
                "pm_password":"p@55w0rd!",
                "pm_addr":"192.0.2.213"
        }
    ]
}
```

After creating the template, save the file to the stack user's home directory
(**/home/stack/instackenv.json**), then import it into the director. Use the following command to
accomplish this:

```
$ openstack baremetal import --json ~/instackenv.json
```

This imports the template and registers each node from the template into the director.

Assign the kernel and ramdisk images to all nodes:

```
$ openstack baremetal configure boot
```

The nodes are now registered and configured in the director.

2.3. MANUALLY TAGGING THE NODES

After registering each node, you will need to inspect the hardware and tag the node into a specific profile.
Profile tags match your nodes to flavors, and in turn the flavors are assigned to a deployment role.

To inspect and tag new nodes, follow these steps:

1. Trigger hardware introspection to retrieve the hardware attributes of each node:

   ```
   $ openstack baremetal introspection bulk start
   ```

   IMPORTANT

   Make sure this process runs to completion. This process usually takes 15 minutes
   for bare metal nodes.

2. Retrieve a list of your nodes to identify their UUIDs:

```
$ ironic node-list
```

3. Add a **profile** option to the **properties/capabilities** parameter for each node to manually tag a node to a specific profile.
   For example, a typical deployment will use three profiles: **control**, **compute**, and **ceph-storage**. The following commands tag three nodes for each profile:

```
$ ironic node-update 1a4e30da-b6dc-499d-ba87-0bd8a3819bc0 add
properties/capabilities='profile:control,boot_option:local'
$ ironic node-update 6faba1a9-e2d8-4b7c-95a2-c7fbdc12129a add
properties/capabilities='profile:control,boot_option:local'
$ ironic node-update 5e3b2f50-fcd9-4404-b0a2-59d79924b38e add
properties/capabilities='profile:control,boot_option:local'
$ ironic node-update 484587b2-b3b3-40d5-925b-a26a2fa3036f add
properties/capabilities='profile:compute,boot_option:local'
$ ironic node-update d010460b-38f2-4800-9cc4-d69f0d067efe add
properties/capabilities='profile:compute,boot_option:local'
$ ironic node-update d930e613-3e14-44b9-8240-4f3559801ea6 add
properties/capabilities='profile:compute,boot_option:local'
$ ironic node-update da0cc61b-4882-45e0-9f43-fab65cf4e52b add
properties/capabilities='profile:ceph-storage,boot_option:local'
$ ironic node-update b9f70722-e124-4650-a9b1-aade8121b5ed add
properties/capabilities='profile:ceph-storage,boot_option:local'
$ ironic node-update 68bf8f29-7731-4148-ba16-efb31ab8d34f add
properties/capabilities='profile:ceph-storage,boot_option:local'
```

TIP

You can also configure a new custom profile that will allow you to tag a node for the Ceph MON and Ceph MDS services. See Section 2.4, "Deploying Other Ceph Services on Dedicated Nodes" for details.

The addition of the **profile** option tags the nodes into each respective profiles.

NOTE

As an alternative to manual tagging, use the Automated Health Check (AHC) Tools to automatically tag larger numbers of nodes based on benchmarking data.

2.4. DEPLOYING OTHER CEPH SERVICES ON DEDICATED NODES

By default, the director deploys the Ceph MON and Ceph MDS services on the Controller nodes. This is suitable for small deployments. However, with larger deployments we advise that you deploy the Ceph MON and Ceph MDS services on dedicated nodes to improve the performance of your Ceph cluster. You can do this by creating a *custom role* for either one.

NOTE

For detailed information about custom roles, see Creating a New Role from the Advanced Overcloud Customization guide.

The director uses the following file as a default reference for all overcloud roles:

/usr/share/openstack-tripleo-heat-templates/roles_data.yaml

Copy this file to **/home/stack/templates/** so you can add custom roles to it:

```
$ cp /usr/share/openstack-tripleo-heat-templates/roles_data.yaml
/home/stack/templates/roles_data_custom.yaml
```

You invoke the **/home/stack/templates/roles_data_custom.yaml** file later during overcloud creation (Section 2.14, "Creating the Overcloud"). The following sub-sections describe how to configure custom roles for either Ceph MON and Ceph MDS services.

2.4.1. Creating a Custom Role and Flavor for the Ceph MON Service

This section describes how to create a custom role (named **CephMon**) and flavor (named **ceph-mon**) for the Ceph MON role. You should already have a copy of the default roles data file as described in Section 2.4, "Deploying Other Ceph Services on Dedicated Nodes".

1. Open the **/home/stack/templates/roles_data_custom.yaml** file.

2. Remove the service entry for the Ceph MON service (namely, **OS::TripleO::Services::CephMon**) from under the Controller role:

   ```
   [...]
   - name: Controller # the 'primary' role goes first
     CountDefault: 1
     ServicesDefault:
       - OS::TripleO::Services::CACerts
       - OS::TripleO::Services::CephMds
   # - OS::TripleO::Services::CephMon  //  ❶
       - OS::TripleO::Services::CephExternal
       - OS::TripleO::Services::CephRbdMirror
       - OS::TripleO::Services::CephRgw
       - OS::TripleO::Services::CinderApi
   [...]
   ```

   ❶ Comment out this line. This same service will be added to a custom role in the next step.

3. At the end of **roles_data_custom.yaml**, add a custom **CephMon** role containing the Ceph MON service and all the other required node services. For example:

   ```
   - name: CephMon
     ServicesDefault:
       - OS::TripleO::Services::CACerts
       - OS::TripleO::Services::CephMon
       - OS::TripleO::Services::Kernel
       - OS::TripleO::Services::Ntp
       - OS::TripleO::Services::Timezone
       - OS::TripleO::Services::Snmp
       - OS::TripleO::Services::Sshd
       - OS::TripleO::Services::TripleoPackages
       - OS::TripleO::Services::TripleoFirewall
       - OS::TripleO::Services::SensuClient
   ```

```
      - OS::TripleO::Services::FluentdClient
      - OS::TripleO::Services::AuditD
      - OS::TripleO::Services::Collectd
```

4. Using the **openstack flavor create** command, define a new flavor named **ceph-mon** for this role:

```
$ openstack flavor create --id auto --ram 6144 --disk 40 --vcpus 4
ceph-mon
```

> **NOTE**
>
> For more details about this command, run **openstack flavor create --help**.

5. Map this flavor to a new profile, also named **ceph-mon**:

```
$ openstack flavor set --property "cpu_arch"="x86_64" --property
"capabilities:boot_option"="local" --property
"capabilities:profile"="ceph-mon" ceph-mon
```

> **NOTE**
>
> For more details about this command, run **openstack flavor set --help**.

Tag nodes into the new **ceph-mon** profile:

```
$ ironic node-update UUID add properties/capabilities='profile:ceph-
mon,boot_option:local'
```

See Section 2.3, "Manually Tagging the Nodes" for more details about tagging nodes. See also Tagging Nodes Into Profiles for related information on custom role profiles.

> **IMPORTANT**
>
> After defining the custom role, you need to configure its port assignments. See Section 2.4.3, "Configuring Port Assignments for Custom Roles" for instructions.

2.4.2. Creating a Custom Role and Flavor for the Ceph MDS Service

This section describes how to create a custom role (named **CephMDS**) and flavor (named **ceph-mds**) for the Ceph MDS role. You should already have a copy of the default roles data file as described in Section 2.4, "Deploying Other Ceph Services on Dedicated Nodes".

1. Open the **/home/stack/templates/roles_data_custom.yaml** file.

2. Remove the service entry for the Ceph MDS service (namely, **OS::TripleO::Services::CephMds**) from under the Controller role:

```
[...]
- name: Controller # the 'primary' role goes first
  CountDefault: 1
```

```
    ServicesDefault:
      - OS::TripleO::Services::CACerts
  # - OS::TripleO::Services::CephMds  //  1
      - OS::TripleO::Services::CephMon
      - OS::TripleO::Services::CephExternal
      - OS::TripleO::Services::CephRbdMirror
      - OS::TripleO::Services::CephRgw
      - OS::TripleO::Services::CinderApi
  [...]
```

1   Comment out this line. This same service will be added to a custom role in the next step.

3. At the end of **roles_data_custom.yaml**, add a custom **CephMDS** role containing the Ceph MDS service and all the other required node services. For example:

```
  - name: CephMDS
    ServicesDefault:
      - OS::TripleO::Services::CACerts
      - OS::TripleO::Services::CephMds
      - OS::TripleO::Services::CephClient  //  1
      - OS::TripleO::Services::Kernel
      - OS::TripleO::Services::Ntp
      - OS::TripleO::Services::Timezone
      - OS::TripleO::Services::Snmp
      - OS::TripleO::Services::Sshd
      - OS::TripleO::Services::TripleoPackages
      - OS::TripleO::Services::TripleoFirewall
      - OS::TripleO::Services::SensuClient
      - OS::TripleO::Services::FluentdClient
      - OS::TripleO::Services::AuditD
      - OS::TripleO::Services::Collectd
```

1   The Ceph MDS service requires the admin keyring, which can be set by either Ceph MON or Ceph Client service. As we are deploying Ceph MDS on a dedicated node (without the Ceph MON service), include the Ceph Client service on the role as well.

4. Using the **openstack flavor create** command, define a new flavor named **ceph-mds** for this role:

```
$ openstack flavor create --id auto --ram 6144 --disk 40 --vcpus 4
ceph-mds
```

NOTE

For more details about this command, run **openstack flavor create -- help**.

5. Map this flavor to a new profile, also named **ceph-mds**:

```
$ openstack flavor set --property "cpu_arch"="x86_64" --property
"capabilities:boot_option"="local" --property
"capabilities:profile"="ceph-mds" ceph-mds
```

NOTE

For more details about this command, run **openstack flavor set --help**.

Tag nodes into the new **ceph-mds** profile:

```
$ ironic node-update UUID add properties/capabilities='profile:ceph-
mds,boot_option:local'
```

See Section 2.3, "Manually Tagging the Nodes" for more details about tagging nodes. See also Tagging Nodes Into Profiles for related information on custom role profiles.

IMPORTANT

After defining the custom role, you need to configure its port assignments. See Section 2.4.3, "Configuring Port Assignments for Custom Roles" for instructions.

2.4.3. Configuring Port Assignments for Custom Roles

The default Heat templates in **/usr/share/openstack-tripleo-heat-templates/** provide the necessary network settings for the default roles. These settings include how IP addresses and ports should be assigned for each service on each node.

Custom roles (like **CephMds** in Section 2.4.2, "Creating a Custom Role and Flavor for the Ceph MDS Service" and **CephMon** in Section 2.4.1, "Creating a Custom Role and Flavor for the Ceph MON Service") do not have the required port assignment Heat templates, so you need to define these yourself. To do so, create a new Heat template named **ports.yaml** in **/home/stack/templates** containing the following snippet for each custom role:

```
resource_registry:
    OS::TripleO::ROLE::Ports::ExternalPort: /usr/share/openstack-tripleo-
heat-templates/network/ports/external.yaml
    OS::TripleO::ROLE::Ports::InternalApiPort: /usr/share/openstack-
tripleo-heat-templates/network/ports/internal_api.yaml
    OS::TripleO::ROLE::Ports::StoragePort: /usr/share/openstack-tripleo-
heat-templates/network/ports/storage.yaml
    OS::TripleO::ROLE::Ports::StorageMgmtPort: /usr/share/openstack-
tripleo-heat-templates/network/ports/storage_mgmt.yaml
    OS::TripleO::ROLE::Ports::TenantPort: /usr/share/openstack-tripleo-
heat-templates/network/ports/tenant.yaml
    OS::TripleO::ROLE::Ports::ManagementPort: /usr/share/openstack-tripleo-
heat-templates/network/ports/management.yaml
```

Replace *ROLE* with the custom role name. For example, if you create the **CephMds** and **CephMon** roles (as described in Section 2.4.2, "Creating a Custom Role and Flavor for the Ceph MDS Service" and Section 2.4.1, "Creating a Custom Role and Flavor for the Ceph MON Service"), your **/home/stack/templates/ports.yaml** should contain:

```
resource_registry:
    OS::TripleO::CephMds::Ports::ExternalPort: /usr/share/openstack-
tripleo-heat-templates/network/ports/external.yaml
    OS::TripleO::CephMds::Ports::InternalApiPort: /usr/share/openstack-
tripleo-heat-templates/network/ports/internal_api.yaml
    OS::TripleO::CephMds::Ports::StoragePort: /usr/share/openstack-tripleo-
```

```
heat-templates/network/ports/storage.yaml
    OS::TripleO::CephMds::Ports::StorageMgmtPort: /usr/share/openstack-
tripleo-heat-templates/network/ports/storage_mgmt.yaml
    OS::TripleO::CephMds::Ports::TenantPort: /usr/share/openstack-tripleo-
heat-templates/network/ports/tenant.yaml
    OS::TripleO::CephMds::Ports::ManagementPort: /usr/share/openstack-
tripleo-heat-templates/network/ports/management.yaml

    OS::TripleO::CephMon::Ports::ExternalPort: /usr/share/openstack-
tripleo-heat-templates/network/ports/external.yaml
    OS::TripleO::CephMon::Ports::InternalApiPort: /usr/share/openstack-
tripleo-heat-templates/network/ports/internal_api.yaml
    OS::TripleO::CephMon::Ports::StoragePort: /usr/share/openstack-tripleo-
heat-templates/network/ports/storage.yaml
    OS::TripleO::CephMon::Ports::StorageMgmtPort: /usr/share/openstack-
tripleo-heat-templates/network/ports/storage_mgmt.yaml
    OS::TripleO::CephMon::Ports::TenantPort: /usr/share/openstack-tripleo-
heat-templates/network/ports/tenant.yaml
    OS::TripleO::CephMon::Ports::ManagementPort: /usr/share/openstack-
tripleo-heat-templates/network/ports/management.yaml
```

When you deploy the overcloud (Section 2.14, "Creating the Overcloud"), include the **/home/stack/templates/ports.yaml** environment file. See Selecting Networks to Deploy (from the Advanced Overcloud Customization guide) for more details on configuring port assignments for custom roles.

2.5. DEFINING THE ROOT DISK FOR CEPH STORAGE NODES

Most Ceph Storage nodes use multiple disks. This means the director needs to identify the disk to use for the root disk when provisioning a Ceph Storage node. There are several properties you can use to help identify the root disk:

- **model** (String): Device identifier.

- **vendor** (String): Device vendor.

- **serial** (String): Disk serial number.

- **wwn** (String): Unique storage identifier.

- **size** (Integer): Size of the device in GB.

In this example, we specify the drive to deploy the Overcloud image using the serial number of the disk to determine the root device.

First, collect a copy of each node's hardware information that the director obtained from the introspection. This information is stored in the OpenStack Object Storage server (swift). Download this information to a new directory:

```
$ mkdir swift-data
$ cd swift-data
$ export SWIFT_PASSWORD=`sudo crudini --get /etc/ironic-
inspector/inspector.conf swift password`
```

```
$ for node in $(ironic node-list | grep -v UUID| awk '{print $2}'); do
swift -U service:ironic -K $SWIFT_PASSWORD download ironic-inspector
inspector_data-$node; done
```

NOTE

This example uses the **crudini** command, which is available in the **crudini** package.

This downloads the data from each **inspector_data** object from introspection. All objects use the node UUID as part of the object name:

```
$ ls -1
inspector_data-15fc0edc-eb8d-4c7f-8dc0-a2a25d5e09e3
inspector_data-46b90a4d-769b-4b26-bb93-50eaefcdb3f4
inspector_data-662376ed-faa8-409c-b8ef-212f9754c9c7
inspector_data-6fc70fe4-92ea-457b-9713-eed499eda206
inspector_data-9238a73a-ec8b-4976-9409-3fcff9a8dca3
inspector_data-9cbfe693-8d55-47c2-a9d5-10e059a14e07
inspector_data-ad31b32d-e607-4495-815c-2b55ee04cdb1
inspector_data-d376f613-bc3e-4c4b-ad21-847c4ec850f8
```

Check the disk information for each node. The following command displays each node ID and the disk information:

```
$ for node in $(ironic node-list | grep -v UUID| awk '{print $2}'); do
echo "NODE: $node" ; cat inspector_data-$node | jq '.inventory.disks' ;
echo "-----" ; done
```

For example, the data for one node might show three disk:

```
NODE: 15fc0edc-eb8d-4c7f-8dc0-a2a25d5e09e3
[
  {
    "size": 299439751168,
    "rotational": true,
    "vendor": "DELL",
    "name": "/dev/sda",
    "wwn_vendor_extension": "0x1ea4dcc412a9632b",
    "wwn_with_extension": "0x61866da04f3807001ea4dcc412a9632b",
    "model": "PERC H330 Mini",
    "wwn": "0x61866da04f380700",
    "serial": "61866da04f3807001ea4dcc412a9632b"
  }
  {
    "size": 299439751168,
    "rotational": true,
    "vendor": "DELL",
    "name": "/dev/sdb",
    "wwn_vendor_extension": "0x1ea4e13c12e36ad6",
    "wwn_with_extension": "0x61866da04f380d001ea4e13c12e36ad6",
    "model": "PERC H330 Mini",
    "wwn": "0x61866da04f380d00",
    "serial": "61866da04f380d001ea4e13c12e36ad6"
  }
```

```
  {
    "size": 299439751168,
    "rotational": true,
    "vendor": "DELL",
    "name": "/dev/sdc",
    "wwn_vendor_extension": "0x1ea4e31e121cfb45",
    "wwn_with_extension": "0x61866da04f37fc001ea4e31e121cfb45",
    "model": "PERC H330 Mini",
    "wwn": "0x61866da04f37fc00",
    "serial": "61866da04f37fc001ea4e31e121cfb45"
  }
]
-----
```

For this example, set the root device to disk 2, which has **61866da04f37fc001ea4e31e121cfb45** as the serial number. This requires a change to the **root_device** parameter for the node definition:

```
$ ironic node-update 15fc0edc-eb8d-4c7f-8dc0-a2a25d5e09e3 add
properties/root_device='{"serial": "61866da04f37fc001ea4e31e121cfb45"}'
```

This helps the director identify the specific disk to use as the root disk. When we initiate our Overcloud creation, the director provisions this node and writes the Overcloud image to this disk. The other disks are used for mapping our Ceph Storage nodes.

> **IMPORTANT**
>
> Do not use **name** to set the root disk as this value can change when the node boots.

2.6. ENABLING CEPH STORAGE IN THE OVERCLOUD

The Overcloud image already contains the Ceph services and the necessary Puppet modules to automatically configure both the Ceph OSD nodes and the Ceph Monitor on Controller clusters. The Overcloud's Heat template collection also contains the necessary procedures to enable your Ceph Storage configuration.

However, the director requires some details to enable Ceph Storage and pass on the intended configuration. You need an *environment file* to do this:

1. Create the file **storage-environment.yaml** in your stack user's templates directory (**~/templates/**).

   > **NOTE**
   >
   > For the purposes of this document, **~/templates/storage-environment.yaml** will contain all the custom settings for your environment. It will override all the default settings applied by the director to your overcloud.

2. If needed, set the following options in **~/templates/storage-environment.yaml** as you see fit:

Option	Description	Default value
CinderEnableIscsiBackend	Enables the iSCSI backend	**false**

Option	Description	Default value
CinderEnableRbdBackend	Enables the Ceph Storage back end	**true**
CinderBackupBackend	Sets **ceph** or **swift** as the back end for volume backups; see Section 2.10, "Configuring the Backup Service to Use Ceph" for related details	**ceph**
NovaEnableRbdBackend	Enables Ceph Storage for Nova ephemeral storage	**true**.
GlanceBackend	Defines which back end the Image service should use: **rbd** (Ceph), **swift**, or **file**	**rbd**
GnocchiBackend	Defines which back end the Telemetry service should use: **rbd** (Ceph), **swift**, or **file**	**rbd**

NOTE

You can omit an option from **~/templates/storage-environment.yaml** if you intend to use the default setting.

3. Add a **parameter_defaults** section to **~/templates/storage-environment.yaml**. This section will contain custom settings for your overcloud. For example, to set **vxlan** as the network type of the networking service (**neutron**):

```
parameter_defaults:
  NeutronNetworkType: vxlan
```

The contents of your environment file will change depending on the settings you apply in the sections that follow. See Appendix A, *Sample Environment File: Creating a Ceph Cluster* for a finished example.

2.7. MAPPING THE CEPH STORAGE NODE DISK LAYOUT

The default mapping uses the root disk for Ceph Storage. However, most production environments use multiple separate disks for storage and partitions for journaling. In this situation, you define a storage map as part of your environment file (namely, **~/templates/storage-environment.yaml** from Section 2.6, "Enabling Ceph Storage in the Overcloud").

Edit the **storage-environment.yaml** file and the following snippet to the **parameter_defaults** section:

```
ExtraConfig:
  ceph::profile::params::osds:
```

This adds extra Hiera data to the Overcloud, which Puppet uses as custom parameters during configuration. Use the **ceph::profile::params::osds** parameter to map the relevant disks and journal partitions. For example, a Ceph node with four disks might have the following assignments:

- **/dev/sda** - The root disk containing the Overcloud image

- **/dev/sdb** - The disk containing the journal partitions. This is usually a solid state disk (SSD) to aid with system performance.

- **/dev/sdc** and **/dev/sdd** - The OSD disks

For this example, the mapping might contain the following:

```
ceph::profile::params::osds:
    '/dev/sdc':
      journal: '/dev/sdb'
    '/dev/sdd':
      journal: '/dev/sdb'
```

If you do not want a separate disk for journals, use co-located journals on the OSD disks. Pass a blank value to the journal parameters:

```
ceph::profile::params::osds:
    '/dev/sdb': {}
    '/dev/sdc': {}
    '/dev/sdd': {}
```

NOTE

In some nodes, disk paths (for example, **/dev/sdb**, **/dev/sdc**) may not point to the exact same block device during reboots. If this is the case with your CephStorage nodes, specify each disk through its **/dev/disk/by-path/** symlink. For example:

```
ceph::profile::params::osds:
   '/dev/disk/by-path/pci-0000:00:17.0-ata-2-part1':
     journal: '/dev/nvme0n1'
   '/dev/disk/by-path/pci-0000:00:17.0-ata-2-part2':
     journal: '/dev/nvme0n1'
```

This will ensure that the block device mapping is consistent throughout deployments.

For more information about naming conventions for storage devices, see Persistent Naming.

You can also deploy Ceph nodes with different types of disks (for example, SSD and SATA disks on the same physical host). In a typical Ceph deployment, this is configured through CRUSH maps, as described in Placing Different Pools on Different OSDS. If you are mapping such a deployment, add the following line to the **ExtraConfig** section of the **storage-environment.yaml**:

```
ceph::osd_crush_update_on_start: false
```

Afterwards, save the **~/templates/storage-environment.yaml** file so that when we deploy the Overcloud, the Ceph Storage nodes use our disk mapping. We include this file in our deployment to initiate our storage requirements.

2.8. CLEANING CEPH STORAGE NODE DISKS

The Ceph Storage OSDs and journal partitions require GPT disk labels. This means the additional disks on Ceph Storage require conversion to GPT before installing the Ceph OSD services. For this to happen, all metadata must be deleted from the disks; this will allow the director to set GPT labels on them.

You can set the director to delete all disk metadata by default by adding the following setting to **/usr/share/instack-undercloud/undercloud.conf**:

```
clean_nodes=true
```

With this option, the Bare Metal Provisioning service will run an additional step to boot the nodes and *clean* the disks each time the node is set to **available**. This adds an additional power cycle after the first introspection and before each deployment. The Bare Metal Provisioning service uses **wipefs --force --all** to perform the clean.

> **WARNING**
>
> The **wipefs --force --all** will delete all data and metadata on the disk, but does not perform a *secure erase*. A secure erase takes much longer.

2.9. DEPLOY THE CEPH OBJECT GATEWAY

The *Ceph Object Gateway* provides applications with an interface to object storage capabilities within a Ceph storage cluster. Upon deploying the Ceph Object Gateway, you can then replace the default Object Storage service (**swift**) with Ceph. For more information, see Object Gateway Guide for Red Hat Enterprise Linux.

To enable a Ceph Object Gateway in your deployment, invoke the following environment file when creating your overcloud:

/usr/share/openstack-tripleo-heat-templates/environments/ceph-radosgw.yaml

See Section 2.14, "Creating the Overcloud" for more details.

The Ceph Object Gateway acts as a drop-in replacement for the default Object Storage service. As such, all other services that normally use **swift** can seamlessly start using the Ceph Object Gateway instead without further configuration. For example, when configuring the Block Storage Backup service (**cinder-backup**) to use the Ceph Object Gateway, set **ceph** as the target back end (see Section 2.10, "Configuring the Backup Service to Use Ceph").

2.10. CONFIGURING THE BACKUP SERVICE TO USE CEPH

The Block Storage Backup service (**cinder-backup**) is disabled by default. To enable it, invoke the following environment file when creating your overcloud:

/usr/share/openstack-tripleo-heat-templates/environments/cinder-backup.yaml

See Section 2.14, "Creating the Overcloud" for more details.

When you enable **cinder-backup** (as in Section 2.9, "Deploy the Ceph Object Gateway"), you can

configure it to store backups in Ceph. This involves adding the following line to the **parameter_defaults** of your environment file (namely, **~/templates/storage-environment.yaml**):

```
CinderBackupBackend: ceph
```

2.11. CONFIGURING MULTIPLE BONDED INTERFACES PER CEPH NODE

Using a *bonded interface* allows you to combine multiple NICs to add redundancy to a network connection. If you have enough NICs on your Ceph nodes, you can take this a step further by creating *multiple bonded interfaces* per node.

With this, you can then use a bonded interface for *each* network connection required by the node. This provides both redundancy and a dedicated connection for each network.

The simplest implementation of this involves the use of two bonds, one for each storage network used by the Ceph nodes. These networks are the following:

Front-end storage network (`StorageNet`)

> The Ceph client uses this network to interact with its Ceph cluster.

Back-end storage network (`StorageMgmtNet`)

> The Ceph cluster uses this network to balance data in accordance with the *placement group* policy of the cluster. For more information, see Placement Groups (PG) (from the Red Hat Ceph Architecture Guide).

Configuring this involves customizing a network interface template, as the director does not provide any sample templates that deploy multiple bonded NICs. However, the director does provide a template that deploys a single bonded interface — namely, **/usr/share/openstack-tripleo-heat-templates/network/config/bond-with-vlans/ceph-storage.yaml**. You can add a bonded interface for your additional NICs by defining it there.

> **NOTE**
>
> For more detailed instructions on how to do this, see Creating Custom Interface Templates (from the Advanced Overcloud Customization guide). That section also explains the different components of a bridge and bonding definition.

The following snippet contains the default definition for the single bonded interface defined by **/usr/share/openstack-tripleo-heat-templates/network/config/bond-with-vlans/ceph-storage.yaml**:

```
type: ovs_bridge //        1
name: br-bond
members:
  -
    type: ovs_bond //       2
    name: bond1 //          3
    ovs_options: {get_param: BondInterfaceOvsOptions}    4
    members: //             5
      -
        type: interface
```

```
        name: nic2
        primary: true
      -
        type: interface
        name: nic3
  -
    type: vlan //  6
    device: bond1 //  7
    vlan_id: {get_param: StorageNetworkVlanID}
    addresses:
      -
        ip_netmask: {get_param: StorageIpSubnet}
  -
    type: vlan
    device: bond1
    vlan_id: {get_param: StorageMgmtNetworkVlanID}
    addresses:
      -
        ip_netmask: {get_param: StorageMgmtIpSubnet}
```

1 A single bridge named **br-bond** holds the bond defined by this template. This line defines the bridge type, namely OVS.

2 The first member of the **br-bond** bridge is the bonded interface itself, named **bond1**. This line defines the bond type of **bond1**, which is also OVS.

3 The default bond is named **bond1**, as defined in this line.

4 The **ovs_options** entry instructs director to use a specific set of *bonding module directives*. Those directives are passed through the **BondInterfaceOvsOptions**, which you can also configure in this same file. For instructions on how to configure this, see Section 2.11.1, "Configuring Bonding Module Directives".

5 The **members** section of the bond defines which network interfaces are bonded by **bond1**. In this case, the bonded interface uses **nic2** (set as the primary interface) and **nic3**.

6 The **br-bond** bridge has two other members: namely, a VLAN for both front-end (**StorageNetwork**) and back-end (**StorageMgmtNetwork**) storage networks.

7 The **device** parameter defines what device a VLAN should use. In this case, both VLANs will use the bonded interface **bond1**.

With at least two more NICs, you can define an additional bridge and bonded interface. Then, you can move one of the VLANs to the new bonded interface. This results in added throughput and reliability for both storage network connections.

When customizing **/usr/share/openstack-tripleo-heat-templates/network/config/bond-with-vlans/ceph-storage.yaml** for this purpose, it is advisable to also use Linux bonds (**type: linux_bond**) instead of the default OVS (**type: ovs_bond**). This bond type is more suitable for enterprise production deployments.

The following edited snippet defines an additional OVS bridge (**br-bond2**) which houses a new Linux bond named **bond2**. The **bond2** interface uses two additional NICs (namely, **nic4** and **nic5**) and will be used solely for back-end storage network traffic:

```
    type: ovs_bridge
    name: br-bond
    members:
      -
        type: linux_bond
        name: bond1
        bonding_options: {get_param: BondInterfaceOvsOptions} // 1
        members:
          -
            type: interface
            name: nic2
            primary: true
          -
            type: interface
            name: nic3
      -
        type: vlan
        device: bond1
        vlan_id: {get_param: StorageNetworkVlanID}
        addresses:
          -
            ip_netmask: {get_param: StorageIpSubnet}
  -
    type: ovs_bridge
    name: br-bond2
    members:
      -
        type: linux_bond
        name: bond2
        bonding_options: {get_param: BondInterfaceOvsOptions}
        members:
          -
            type: interface
            name: nic4
            primary: true
          -
            type: interface
            name: nic5
      -
        type: vlan
        device: bond1
        vlan_id: {get_param: StorageMgmtNetworkVlanID}
        addresses:
          -
            ip_netmask: {get_param: StorageMgmtIpSubnet}
```

1 As **bond1** and **bond2** are both Linux bonds (instead of OVS), they use **bonding_options** instead of **ovs_options** to set bonding directives. For related information, see Section 2.11.1, "Configuring Bonding Module Directives".

For the full contents of this customized template, see Appendix B, *Sample Custom Interface Template: Multiple Bonded Interfaces*.

2.11.1. Configuring Bonding Module Directives

After adding and configuring the bonded interfaces, use the **BondInterfaceOvsOptions** parameter to set what directives each should use. You can find this in the **parameters:** section of **/usr/share/openstack-tripleo-heat-templates/network/config/bond-with-vlans/ceph-storage.yaml**. The following snippet shows the default definition of this parameter (namely, empty):

```
BondInterfaceOvsOptions:
    default: ''
    description: The ovs_options string for the bond interface. Set
                 things like lacp=active and/or bond_mode=balance-slb
                 using this option.
    type: string
```

Define the options you need in the **default:** line. For example, to use 802.3ad (mode 4) and a LACP rate of 1 (fast), use **'mode=4 lacp_rate=1'**, as in:

```
BondInterfaceOvsOptions:
    default: 'mode=4 lacp_rate=1'
    description: The bonding_options string for the bond interface. Set
                 things like lacp=active and/or bond_mode=balance-slb
                 using this option.
    type: string
```

See Appendix C. Open vSwitch Bonding Options (from the Advanced Overcloud Optimization guide) for other supported bonding options. For the full contents of the customized **/usr/share/openstack-tripleo-heat-templates/network/config/bond-with-vlans/ceph-storage.yaml** template, see Appendix B, *Sample Custom Interface Template: Multiple Bonded Interfaces*.

2.12. CUSTOMIZING THE CEPH STORAGE CLUSTER

It is possible to override the default configuration parameters for Ceph Storage nodes using the **ExtraConfig** hook to define data to pass to the Puppet configuration. There are two methods to pass this data:

Method 1: Modifying Puppet Defaults

You customize parameters provided to the **ceph** Puppet module during the overcloud configuration. These parameters are a part of the **ceph::profile::params** Puppet class defined in **/etc/puppet/modules/ceph/manifests/profile/params.conf**. For example, following environment file snippet customizes the default **osd_journal_size** parameter from the **ceph::profile::params** class and overrides any default:

```
parameter_defaults:
  ExtraConfig:
    ceph::profile::params::osd_journal_size: 2048
```

Add this content to an environment file (for example, **ceph-settings.yaml**) and include it when you run the **openstack overcloud deploy** command in Section 2.14, "Creating the Overcloud". For example:

```
$ openstack overcloud deploy --templates -e /home/stack/templates/storage-
environment.yaml -e /home/stack/templates/ceph-settings.yaml
```

Method 2: Arbitrary Configuration Defaults

Method 2. Arbitrary Configuration Defaults

If Method 1 does not include a specific parameter you need to configure, it is possible to provide arbitrary Ceph Storage parameters using the **ceph::conf::args** Puppet class. This class accepts parameter names using a **stanza/key** format and **value** to define the parameter's value. These settings configure the **ceph.conf** file on each node. For example, to change the **max_open_files** parameter in the **global** section of the **ceph.conf** file, use the following structure in an environment file:

```
parameter_defaults:
  ExtraConfig:
    ceph::conf::args:
      global/max_open_files:
        value: 131072
```

Add this content to an environment file (for example, **ceph-settings.yaml**) and include it when you run the **openstack overcloud deploy** command in Section 2.14, "Creating the Overcloud". For example:

```
$ openstack overcloud deploy --templates -e /home/stack/templates/storage-
environment.yaml -e /home/stack/templates/ceph-settings.yaml
```

The resulting **ceph.conf** file should be populated with the following:

```
[global]
max_open_files = 131072
```

2.12.1. Assigning Custom Attributes to Different Ceph Pools

By default, Ceph pools created through the director have the same placement group (**pg_num** and **pgp_num**) and sizes. You can use either method in Section 2.12, "Customizing the Ceph Storage Cluster" to override these settings globally; that is, doing so will apply the same values for all pools.

You can also apply different attributes to each Ceph pool. To do so, use the **CephPools** resource, as in:

```
parameter_defaults:
  CephPools:
    POOL:
      size: 5
      pg_num: 128
      pgp_num: 128
```

Replace *POOL* with the name of the pool you want to configure with the **size**, **pg_num**, and **pgp_num** settings that follow.

TIP

For typical pool configurations of common Ceph use cases, see the Ceph Placement Groups (PGs) per Pool Calculator. This calculator is normally used to generate the commands for manually configuring your Ceph pools. In this deployment, the director will configure the pools based on your specifications.

2.13. ASSIGNING NODES AND FLAVORS TO ROLES

Planning an overcloud deployment involves specifying how many nodes and which flavors to assign to

each role. Like all Heat template parameters, these role specifications are declared in the **parameter_defaults** section of your environment file (in this case, **~/templates/storage-environment.yaml**).

For this purpose, use the following parameters:

Table 2.1. Roles and Flavors for Overcloud Nodes

Heat Template Parameter	Description
ControllerCount	The number of Controller nodes to scale out
OvercloudControlFlavor	The flavor to use for Controller nodes (**control**)
ComputeCount	The number of Compute nodes to scale out
OvercloudComputeFlavor	The flavor to use for Compute nodes (**compute**)
CephStorageCount	The number of Ceph storage (OSD) nodes to scale out
OvercloudCephStorageFlavor	The flavor to use for Ceph Storage (OSD) nodes (**ceph-storage**)
CephMonCount	The number of dedicated Ceph MON nodes to scale out
CephMdsCount	The number of dedicated Ceph MDS nodes to scale out
OvercloudCephMonFlavor	The flavor to use for dedicated Ceph MON nodes (**ceph-mon**)
OvercloudCephMdsFlavor	The flavor to use for dedicated Ceph MDS nodes (**ceph-mds**)

> **IMPORTANT**
>
> The **CephMonCount**, **CephMdsCount**, **OvercloudCephMonFlavor**, and **OvercloudCephMdsFlavor** parameters (along with the **ceph-mon** and **ceph-mds** flavors) will only be valid if you created a custom **CephMON** and **CephMds** role, as described in Section 2.4, "Deploying Other Ceph Services on Dedicated Nodes".

For example, to configure the overcloud to deploy three nodes for each role (Controller, Compute, Ceph-Storage, and CephMon), add the following to your **parameter_defaults**:

```
parameter_defaults:
  ControllerCount: 3
  OvercloudControlFlavor: control
  ComputeCount: 3
```

```
OvercloudComputeFlavor: compute
CephStorageCount: 3
OvercloudCephStorageFlavor: ceph-storage
CephMonCount: 3
OvercloudCephMonFlavor: ceph-mon
CephMdsCount: 3
OvercloudCephMdsFlavor: ceph-mds
```

> **NOTE**
>
> See Creating the Overcloud with the CLI Tools from the Director Installation and Usage guide for a more complete list of Heat template parameters.

2.14. CREATING THE OVERCLOUD

The creation of the Overcloud requires additional arguments for the **openstack overcloud deploy** command. For example:

```
$ openstack overcloud deploy --templates -r
/home/stack/templates/roles_data_custom.yaml \
  -e /usr/share/openstack-tripleo-heat-templates/environments/ceph-
radosgw.yaml \
  -e /usr/share/openstack-tripleo-heat-templates/environments/cinder-
backup.yaml \
  -e /home/stack/templates/storage-environment.yaml \
  -e /home/stack/templates/ports.yaml \
  --ntp-server pool.ntp.org
```

The above command uses the following options:

- **--templates** - Creates the Overcloud from the default Heat template collection (namely, **/usr/share/openstack-tripleo-heat-templates/**).

- **-r /home/stack/templates/roles_data_custom.yaml** - Specifies the customized roles definition file from Section 2.4, "Deploying Other Ceph Services on Dedicated Nodes", which adds custom roles for either Ceph MON or Ceph MDS services. These roles allow either service to be installed on dedicated nodes.

- **-e /usr/share/openstack-tripleo-heat-templates/environments/ceph-radosgw.yaml** - Enables the Ceph Object Gateway, as described in Section 2.9, "Deploy the Ceph Object Gateway".

- **-e /usr/share/openstack-tripleo-heat-templates/environments/cinder-backup.yaml** - Enables the Block Storage Backup service (**cinder-backup**), as described in Section 2.10, "Configuring the Backup Service to Use Ceph".

- **-e /home/stack/templates/storage-environment.yaml** - Adds the environment file containing our Ceph Storage configuration.

- **-e /home/stack/templates/ports.yaml** - Adds the Heat template that configures the port assignments for custom roles (see Section 2.4.3, "Configuring Port Assignments for Custom Roles" for details).

- **--ntp-server pool.ntp.org** - Sets our NTP server.

See Appendix A, *Sample Environment File: Creating a Ceph Cluster* for an overview of all the settings used in **/home/stack/templates/storage-environment.yaml**.

TIP

You can also use an *answers file* to invoke all your templates and environment files. For example, you can use the following command to deploy an identical overcloud:

```
$ openstack overcloud deploy -r
/home/stack/templates/roles_data_custom.yaml \
  --answers-file /home/stack/templates/answers.yaml --ntp-server
pool.ntp.org
```

In this case, the answers file **/home/stack/templates/answers.yaml** contains:

```
templates: /usr/share/openstack-tripleo-heat-templates/
environments:
  - /usr/share/openstack-tripleo-heat-templates/environments/ceph-
radosgw.yaml
  - /usr/share/openstack-tripleo-heat-templates/environments/cinder-
backup.yaml
  - /home/stack/templates/storage-environment.yaml
  - /home/stack/templates/ports.yaml
```

See Including Environment Files in Overcloud Creation for more details

For a full list of options, run:

```
$ openstack help overcloud deploy
```

For more information, see Creating the Overcloud with the CLI Tools in the Director Installation and Usage guide.

The Overcloud creation process begins and the director provisions your nodes. This process takes some time to complete. To view the status of the Overcloud creation, open a separate terminal as the **stack** user and run:

```
$ source ~/stackrc
$ heat stack-list --show-nested
```

2.15. ACCESSING THE OVERCLOUD

The director generates a script to configure and help authenticate interactions with your Overcloud from the director host. The director saves this file (**overcloudrc**) in your **stack** user's home directory. Run the following command to use this file:

```
$ source ~/overcloudrc
```

This loads the necessary environment variables to interact with your Overcloud from the director host's CLI. To return to interacting with the director's host, run the following command:

```
$ source ~/stackrc
```

2.16. MONITORING CEPH STORAGE NODES

After completing the Overcloud creation, we recommend that you check the status of the Ceph Storage Cluster to ensure it is working properly. To do this, log into a Controller node as the **heat-admin** user:

```
$ nova list
$ ssh heat-admin@192.168.0.25
```

Check the health of the cluster:

```
$ sudo ceph health
```

If the cluster has no issues, the command reports back **HEALTH_OK**. This means the cluster is safe to use.

Check the status of the Ceph Monitor quorum:

```
$ sudo ceph quorum_status
```

This shows the monitors participating in the quorum and which one is the leader.

Check if all Ceph OSDs are running:

```
$ ceph osd stat
```

For more information on monitoring Ceph Storage clusters, see Monitoring in the *Red Hat Ceph Storage Administration Guide*.

2.17. REBOOTING THE ENVIRONMENT

A situation might occur where you need to reboot the environment. For example, when you might need to modify the physical servers, or you might need to recover from a power outage. In this situation, it is important to make sure your Ceph Storage nodes boot correctly.

Make sure to boot the nodes in the following order:

- **Boot all Ceph Monitor nodes first** - This ensures the Ceph Monitor service is active in your high availability cluster. By default, the Ceph Monitor service is installed on the Controller node. If the Ceph Monitor is separate from the Controller in a custom role, make sure this custom Ceph Monitor role is active.

- **Boot all Ceph Storage nodes** - This ensures the Ceph OSD cluster can connect to the active Ceph Monitor cluster on the Controller nodes.

Use the following process to reboot the Ceph Storage nodes:

1. Log into a Ceph MON or Controller node and disable Ceph Storage cluster rebalancing temporarily:

   ```
   $ sudo ceph osd set noout
   $ sudo ceph osd set norebalance
   ```

2. Select the first Ceph Storage node to reboot and log into it.

3. Reboot the node:

```
$ sudo reboot
```

4. Wait until the node boots.

5. Log into the node and check the cluster status:

```
$ sudo ceph -s
```

Check that the **pgmap** reports all **pgs** as normal (**active+clean**).

6. Log out of the node, reboot the next node, and check its status. Repeat this process until you have rebooted all Ceph storage nodes.

7. When complete, log into a Ceph MON or Controller node and enable cluster rebalancing again:

```
$ sudo ceph osd unset noout
$ sudo ceph osd unset norebalance
```

8. Perform a final status check to verify the cluster reports **HEALTH_OK**:

```
$ sudo ceph status
```

If a situation occurs where all Overcloud nodes boot at the same time, the Ceph OSD services might not start correctly on the Ceph Storage nodes. In this situation, reboot the Ceph Storage OSDs so they can connect to the Ceph Monitor service. Run the following command on each Ceph Storage node:

```
$ sudo systemctl restart 'ceph*'
```

Verify a **HEALTH_OK** status of the Ceph Storage node cluster with the following command:

```
$ sudo ceph status
```

2.18. SCALING UP THE CEPH CLUSTER

You can scale up the number of Ceph Storage nodes in your overcloud by re-running the deployment with the number of Ceph Storage nodes you need.

Before doing so, ensure that you have enough nodes for the updated deployment. These nodes must be registered with the director and tagged accordingly.

Registering New Ceph Storage Nodes

To register new Ceph storage nodes with the director, follow these steps:

1. Log into the director host as the **stack** user and initialize your director configuration:

```
$ source ~/stackrc
```

2. Define the hardware and power management details for the new nodes in a new node definition template; for example, **instackenv-scale.json**.

3. Import this file to the OpenStack director:

```
$ openstack baremetal import --json ~/instackenv-scale.json
```

Importing the node definition template registers each node defined there to the director.

4. Assign the kernel and ramdisk images to all nodes:

```
$ openstack baremetal configure boot
```

NOTE

For more information about registering new nodes, see Section 2.2, "Registering Nodes".

Manually Tagging New Nodes

After registering each node, you will need to inspect the hardware and tag the node into a specific profile. Profile tags match your nodes to flavors, and in turn the flavors are assigned to a deployment role.

To inspect and tag new nodes, follow these steps:

1. Trigger hardware introspection to retrieve the hardware attributes of each node:

```
$ openstack baremetal introspection bulk start
```

IMPORTANT

Make sure this process runs to completion. This process usually takes 15 minutes for bare metal nodes.

2. Retrieve a list of your nodes to identify their UUIDs:

```
$ ironic node-list
```

3. Add a **profile** option to the **properties/capabilities** parameter for each node to manually tag a node to a specific profile.
   For example, the following commands tag three additional nodes with the **ceph-storage** profile:

```
$ ironic node-update 551d81f5-4df2-4e0f-93da-6c5de0b868f7 add
properties/capabilities='profile:ceph-storage,boot_option:local'
$ ironic node-update 5e735154-bd6b-42dd-9cc2-b6195c4196d7 add
properties/capabilities='profile:ceph-storage,boot_option:local'
$ ironic node-update 1a2b090c-299d-4c20-a25d-57dd21a7085b add
properties/capabilities='profile:ceph-storage,boot_option:local'
```

TIP

If the nodes you just tagged and registered use multiple disks, you can set the director to use a specific root disk on each node. See Section 2.5, "Defining the Root Disk for Ceph Storage Nodes" for instructions on how to do so.

Re-deploying the Overcloud with Additional Ceph Storage Nodes

After registering and tagging the new nodes, you can now scale up the number of Ceph Storage nodes by re-deploying the overcloud. When you do, set the **CephStorageCount** parameter in the **parameter_defaults** of your environment file (in this case,**~/templates/storage-environment.yaml**). In Section 2.13, "Assigning Nodes and Flavors to Roles", the overcloud is configured to deploy with 3 Ceph Storage nodes. To scale it up to 6 nodes instead, use:

```
parameter_defaults:
  ControllerCount: 3
  OvercloudControlFlavor: control
  ComputeCount: 3
  OvercloudComputeFlavor: compute
  CephStorageCount: 6
  OvercloudCephStorageFlavor: ceph-storage
  CephMonCount: 3
  OvercloudCephMonFlavor: ceph-mon
```

Upon re-deployment with this setting, the overcloud should now have 6 Ceph Storage nodes instead of 3.

2.19. SCALING DOWN AND REPLACING CEPH STORAGE NODES

In some cases, you may need to scale down your Ceph cluster, or even replace a Ceph Storage node (for example, if a Ceph Storage node is faulty). In either situation, you need to disable and rebalance any Ceph Storage node you are removing from the Overcloud to ensure no data loss. This procedure explains the process for replacing a Ceph Storage node.

> **NOTE**
>
> This procedure uses steps from the *Red Hat Ceph Storage Administration Guide* to manually remove Ceph Storage nodes. For more in-depth information about manual removal of Ceph Storage nodes, see Adding and Removing OSD Nodes from the *Red Hat Ceph Storage Administration Guide*.

Log into either a Controller node or a Ceph Storage node as the **heat-admin** user. The director's **stack** user has an SSH key to access the **heat-admin** user.

List the OSD tree and find the OSDs for your node. For example, your node to remove might contain the following OSDs:

```
-2 0.09998      host overcloud-cephstorage-0
0 0.04999        osd.0                           up  1.00000
1.00000
1 0.04999        osd.1                           up  1.00000
1.00000
```

Disable the OSDs on the Ceph Storage node. In this case, the OSD IDs are 0 and 1.

```
[heat-admin@overcloud-controller-0 ~]$ sudo ceph osd out 0
[heat-admin@overcloud-controller-0 ~]$ sudo ceph osd out 1
```

The Ceph Storage cluster begins rebalancing. Wait for this process to complete. You can follow the status using the following command:

```
[heat-admin@overcloud-controller-0 ~]$ sudo ceph -w
```

Once the Ceph cluster completes rebalancing, log into the Ceph Storage node you are removing (in this case, **overcloud-cephstorage-0**) as the **heat-admin** user and stop the node.

```
[heat-admin@overcloud-cephstorage-0 ~]$ sudo systemctl disable ceph-osd@0
[heat-admin@overcloud-cephstorage-0 ~]$ sudo systemctl disable ceph-osd@1
```

While logged into **overcloud-cephstorage-0**, remove it from the CRUSH map so that it no longer receives data.

```
[heat-admin@overcloud-cephstorage-0 ~]$ sudo ceph osd crush remove osd.0
[heat-admin@overcloud-cephstorage-0 ~]$ sudo ceph osd crush remove osd.1
```

Remove the OSD authentication key.

```
[heat-admin@overcloud-cephstorage-0 ~]$ sudo ceph auth del osd.0
[heat-admin@overcloud-cephstorage-0 ~]$ sudo ceph auth del osd.1
```

Remove the OSD from the cluster.

```
[heat-admin@overcloud-cephstorage-0 ~]$ sudo ceph osd rm 0
[heat-admin@overcloud-cephstorage-0 ~]$ sudo ceph osd rm 1
```

Leave the node and return to the director host as the **stack** user.

```
[heat-admin@overcloud-cephstorage-0 ~]$ exit
[stack@director ~]$
```

Disable the Ceph Storage node so the director does not reprovision it.

```
[stack@director ~]$ ironic node-list
[stack@director ~]$ ironic node-set-maintenance UUID true
```

Removing a Ceph Storage node requires an update to the **overcloud** stack in the director using the local template files. First identify the UUID of the Overcloud stack:

```
$ heat stack-list
```

Identify the UUIDs of the Ceph Storage node to delete:

```
$ nova list
```

Run the following command to delete the node from the stack and update the plan accordingly:

```
$ openstack overcloud node delete --stack STACK_UUID --templates -e
ENVIRONMENT_FILE NODE_UUID
```

IMPORTANT

If you passed any extra environment files when you created the overcloud, pass them again here using the **-e** option to avoid making undesired changes to the overcloud. For more information, see Modifying the Overcloud Environment (from Director Installation and Usage).

Wait until the stack completes its update. Monitor the stack update using the **heat stack-list --show-nested** command.

Add new nodes to the director's node pool and deploy them as Ceph Storage nodes. Use the **CephStorageCount** parameter in the **parameter_defaults** of your environment file (in this case, **~/templates/storage-environment.yaml**) to define the total number of Ceph Storage nodes in the Overcloud. For example:

```
parameter_defaults:
  ControllerCount: 3
  OvercloudControlFlavor: control
  ComputeCount: 3
  OvercloudComputeFlavor: compute
  CephStorageCount: 3
  OvercloudCephStorageFlavor: ceph-storage
  CephMonCount: 3
  OvercloudCephMonFlavor: ceph-mon
```

See Section 2.13, "Assigning Nodes and Flavors to Roles" for details on how to define the number of nodes per role.

Upon updating your environment file, re-deploy the overcloud as normal:

```
$ openstack overcloud deploy --templates -e ENVIRONMENT_FILES
```

The director provisions the new node and updates the entire stack with the new node's details.

Log into a Controller node as the **heat-admin** user and check the status of the Ceph Storage node. For example:

```
[heat-admin@overcloud-controller-0 ~]$ sudo ceph status
```

Confirm that the value in the **osdmap** section matches the number of desired nodes in your cluster. The Ceph Storage node you removed has now been replaced with a new node.

2.20. ADDING AND REMOVING OSD DISKS FROM CEPH STORAGE NODES

In situations when an OSD disk fails and requires a replacement, use the standard instructions from the *Red Hat Ceph Storage Administration Guide*:

- "Adding an OSD"

- "Removing an OSD"

CHAPTER 3. INTEGRATING AN EXISTING CEPH STORAGE CLUSTER WITH AN OVERCLOUD

This chapter describes how to create an Overcloud and integrate it with an existing Ceph Storage Cluster. For instructions on how to create both Overcloud and Ceph Storage Cluster, see Chapter 2, *Creating an Overcloud with Ceph Storage Nodes* instead.

The scenario described in this chapter consists of six nodes in the Overcloud:

- Three Controller nodes with high availability.

- Three Compute nodes.

The director will integrate a separate Ceph Storage Cluster with its own nodes into the Overcloud. You manage this cluster independently from the Overcloud. For example, you scale the Ceph Storage cluster using the Ceph management tools and not through the OpenStack Platform director. Consult the Red Hat Ceph documentation for more information.

All OpenStack machines are bare metal systems using IPMI for power management. These nodes do not require an operating system because the director copies a Red Hat Enterprise Linux 7 image to each node.

The director communicates to the Controller and Compute nodes through the Provisioning network during the introspection and provisioning processes. All nodes connect to this network through the native VLAN. For this example, we use 192.0.2.0/24 as the Provisioning subnet with the following IP address assignments:

Node Name	IP Address	MAC Address	IPMI IP Address	
Director	192.0.2.1	aa:aa:aa:aa:aa:aa		
Controller 1	DHCP defined	b1:b1:b1:b1:b1:b1	192.0.2.205	
Controller 2	DHCP defined	b2:b2:b2:b2:b2:b2	192.0.2.206	
Controller 3	DHCP defined	b3:b3:b3:b3:b3:b3	192.0.2.207	
Compute 1	DHCP defined	c1:c1:c1:c1:c1:c1	192.0.2.208	
Compute 2	DHCP defined	c2:c2:c2:c2:c2:c2	192.0.2.209	
Compute 3	DHCP defined	c3:c3:c3:c3:c3:c3	192.0.2.210	

3.1. CONFIGURING THE EXISTING CEPH STORAGE CLUSTER

1. Create the following pools in your Ceph cluster relevant to your environment:

   - `volumes`: Storage for OpenStack Block Storage (cinder)

   - `images`: Storage for OpenStack Image Storage (glance)

- **vms**: Storage for instances

- **backups**: Storage for OpenStack Block Storage Backup (cinder-backup)

- **metrics**: Storage for OpenStack Telemetry Metrics (gnocchi)
  Use the following commands as a guide:

```
[root@ceph ~]# ceph osd pool create volumes PGNUM
[root@ceph ~]# ceph osd pool create images PGNUM
[root@ceph ~]# ceph osd pool create vms PGNUM
[root@ceph ~]# ceph osd pool create backups PGNUM
[root@ceph ~]# ceph osd pool create metrics PGNUM
```

Replace *PGNUM* with the number of *placement groups*. We recommend approximately 100 per OSD. For example, the total number of OSDs multiplied by 100 divided by the number of replicas (**osd pool default size**). You can also use the Ceph Placement Groups (PGs) per Pool Calculator to determine a suitable value.

2. Create a **client.openstack** user in your Ceph cluster with the following capabilities:

- **cap_mon: allow r**

- **cap_osd: allow class-read object_prefix rbd_children, allow rwx pool=volumes, allow rwx pool=vms, allow rwx pool=images, allow rwx pool=backups, allow rwx pool=metrics**
  Use the following command as a guide:

```
[root@ceph ~]# ceph auth add client.openstack mon 'allow r' osd
'allow class-read object_prefix rbd_children, allow rwx
pool=volumes, allow rwx pool=vms, allow rwx pool=images, allow
rwx pool=backups, allow rwx pool=metrics'
```

3. Next, note the *Ceph client key* created for the **client.openstack** user:

```
[root@ceph ~]# ceph auth list
...
client.openstack
 key: AQDLOh1VgEp6FRAAFzT7Zw+Y9V6JJExQAsRnRQ==
 caps: [mon] allow r
 caps: [osd] allow class-read object_prefix rbd_children, allow rwx
pool=volumes, allow rwx pool=vms, allow rwx pool=images, allow rwx
pool=backups, allow rwx pool=metrics
...
```

The **key** value here (**AQDLOh1VgEp6FRAAFzT7Zw+Y9V6JJExQAsRnRQ==**) is your Ceph client key.

4. Finally, note the *file system ID* of your Ceph Storage cluster. This value is specified with the **fsid** setting in the configuration file of your cluster (under the **[global]** section):

```
[global]
fsid = 4b5c8c0a-ff60-454b-a1b4-9747aa737d19
...
```

NOTE

For more information about the Ceph Storage cluster configuration file, see Configuration Reference (from the Red Hat Ceph Storage Configuration Guide).

The Ceph client key and file system ID will both be used later in Section 3.5, "Integrating with the Existing Ceph Storage Cluster".

3.2. INITIALIZING THE STACK USER

Log into the director host as the **stack** user and run the following command to initialize your director configuration:

```
$ source ~/stackrc
```

This sets up environment variables containing authentication details to access the director's CLI tools.

3.3. REGISTERING NODES

A node definition template (**instackenv.json**) is a JSON format file and contains the hardware and power management details for registering nodes. For example:

```
{
    "nodes":[
        {
            "mac":[
                "bb:bb:bb:bb:bb:bb"
            ],
            "cpu":"4",
            "memory":"6144",
            "disk":"40",
            "arch":"x86_64",
            "pm_type":"pxe_ipmitool",
            "pm_user":"admin",
            "pm_password":"p@55w0rd!",
            "pm_addr":"192.0.2.205"
        },
        {
            "mac":[
                "cc:cc:cc:cc:cc:cc"
            ],
            "cpu":"4",
            "memory":"6144",
            "disk":"40",
            "arch":"x86_64",
            "pm_type":"pxe_ipmitool",
            "pm_user":"admin",
            "pm_password":"p@55w0rd!",
            "pm_addr":"192.0.2.206"
        },
        {
            "mac":[
                "dd:dd:dd:dd:dd:dd"
            ],
```

```
            "cpu":"4",
            "memory":"6144",
            "disk":"40",
            "arch":"x86_64",
            "pm_type":"pxe_ipmitool",
            "pm_user":"admin",
            "pm_password":"p@55w0rd!",
            "pm_addr":"192.0.2.207"
        },
        {
            "mac":[
                "ee:ee:ee:ee:ee:ee"
            ],
            "cpu":"4",
            "memory":"6144",
            "disk":"40",
            "arch":"x86_64",
            "pm_type":"pxe_ipmitool",
            "pm_user":"admin",
            "pm_password":"p@55w0rd!",
            "pm_addr":"192.0.2.208"
        }
        {
            "mac":[
                "ff:ff:ff:ff:ff:ff"
            ],
            "cpu":"4",
            "memory":"6144",
            "disk":"40",
            "arch":"x86_64",
            "pm_type":"pxe_ipmitool",
            "pm_user":"admin",
            "pm_password":"p@55w0rd!",
            "pm_addr":"192.0.2.209"
        }
        {
            "mac":[
                "gg:gg:gg:gg:gg:gg"
            ],
            "cpu":"4",
            "memory":"6144",
            "disk":"40",
            "arch":"x86_64",
            "pm_type":"pxe_ipmitool",
            "pm_user":"admin",
            "pm_password":"p@55w0rd!",
            "pm_addr":"192.0.2.210"
        }
    ]
}
```

After creating the template, save the file to the stack user's home directory
(**/home/stack/instackenv.json**), then import it into the director. Use the following command to
accomplish this:

```
$ openstack baremetal import --json ~/instackenv.json
```

–

This imports the template and registers each node from the template into the director.

Assign the kernel and ramdisk images to all nodes:

```
$ openstack baremetal configure boot
```

The nodes are now registered and configured in the director.

3.4. MANUALLY TAGGING THE NODES

After registering each node, you will need to inspect the hardware and tag the node into a specific profile. Profile tags match your nodes to flavors, and in turn the flavors are assigned to a deployment role.

To inspect and tag new nodes, follow these steps:

1. Trigger hardware introspection to retrieve the hardware attributes of each node:

   ```
   $ openstack baremetal introspection bulk start
   ```

   > **IMPORTANT**
   >
   > Make sure this process runs to completion. This process usually takes 15 minutes for bare metal nodes.

2. Retrieve a list of your nodes to identify their UUIDs:

   ```
   $ ironic node-list
   ```

3. Add a **profile** option to the **properties/capabilities** parameter for each node to manually tag a node to a specific profile.
   For example, to tag three nodes to use the **control** profile and another three nodes to use the **compute** profile, run:

   ```
   $ ironic node-update 1a4e30da-b6dc-499d-ba87-0bd8a3819bc0 add
   properties/capabilities='profile:control,boot_option:local'
   $ ironic node-update 6faba1a9-e2d8-4b7c-95a2-c7fbdc12129a add
   properties/capabilities='profile:control,boot_option:local'
   $ ironic node-update 5e3b2f50-fcd9-4404-b0a2-59d79924b38e add
   properties/capabilities='profile:control,boot_option:local'
   $ ironic node-update 484587b2-b3b3-40d5-925b-a26a2fa3036f add
   properties/capabilities='profile:compute,boot_option:local'
   $ ironic node-update d010460b-38f2-4800-9cc4-d69f0d067efe add
   properties/capabilities='profile:compute,boot_option:local'
   $ ironic node-update d930e613-3e14-44b9-8240-4f3559801ea6 add
   properties/capabilities='profile:compute,boot_option:local'
   ```

The addition of the **profile** option tags the nodes into each respective profiles.

> **NOTE**
>
> As an alternative to manual tagging, use the Automated Health Check (AHC) Tools to automatically tag larger numbers of nodes based on benchmarking data.

3.5. INTEGRATING WITH THE EXISTING CEPH STORAGE CLUSTER

Create a copy of **/usr/share/openstack-tripleo-heat-templates/environments/puppet-ceph-external.yaml** to a directory in your **stack** user's home directory:

```
[stack@director ~]# mkdir templates
[stack@director ~]# cp /usr/share/openstack-tripleo-heat-
templates/environments/puppet-ceph-external.yaml ~/templates/.
```

Edit the file and set the following parameters:

- Set the **CephExternal** resource definition to an absolute path:

```
OS::TripleO::Services::CephExternal: /usr/share/openstack-tripleo-heat-
templates/puppet/services/ceph-external.yaml
```

- Set the following three parameters using your Ceph Storage environment details:

  - **CephClientKey**: the Ceph client key of your Ceph Storage cluster. This is the value of**key** you retrieved earlier in Section 3.1, "Configuring the Existing Ceph Storage Cluster"(for example, **AQDLOh1VgEp6FRAAFzT7Zw+Y9V6JJExQAsRnRQ==**).

  - **CephExternalMonHost**: a comma-delimited list of the IPs of all MON hosts in your Ceph Storage cluster.

  - **CephClusterFSID**: the file system ID of your Ceph Storage cluster. This is the value of **fsid** in your Ceph Storage cluster configuration file, which you retrieved earlier in Section 3.1, "Configuring the Existing Ceph Storage Cluster"(for example, **4b5c8c0a-ff60-454b-a1b4-9747aa737d19**).

- Add the following parameter to set **vxlan** as the **neutron** network type:

  - **NeutronNetworkType: vxlan**

- If necessary, also set the name of the OpenStack pools and the client user using the following parameters and values:

  - **CephClientUserName: openstack**

  - **NovaRbdPoolName: vms**

  - **CinderRbdPoolName: volumes**

  - **GlanceRbdPoolName: images**

  - **CinderBackupRbdPoolName: backups**

  - **GnocchiRbdPoolName: metrics**

3.6. BACKWARDS COMPATIBILITY WITH OLDER VERSIONS OF RED HAT CEPH STORAGE

If you are integrating Red Hat OpenStack Platform with an external Ceph Storage Cluster from an earlier version (that is, Red Hat Ceph Storage 1.3), you need to enable backwards compatibility. To do so, first

create an environment file (for example, **/home/stack/templates/ceph-backwards-compatibility.yaml**) containing the following:

```
parameter_defaults:
   RbdDefaultFeatures: 1
```

Include this file in your overcloud deployment, described in Section 3.8, "Creating the Overcloud".

Alternatively, you can also uncomment the following line from **/home/stack/templates/puppet-ceph-external.yaml** (which you copied earlier in Section 3.5, "Integrating with the Existing Ceph Storage Cluster"):

```
   # RbdDefaultFeatures: 1
```

3.7. ASSIGNING NODES AND FLAVORS TO ROLES

Planning an overcloud deployment involves specifying how many nodes and which flavors to assign to each role. Like all Heat template parameters, these role specifications are declared in the **parameter_defaults** section of your environment file (in this case, **~/templates/puppet-ceph.yaml**.).

For this purpose, use the following parameters:

Table 3.1. Roles and Flavors for Overcloud Nodes

Heat Template Parameter	Description
ControllerCount	The number of Controller nodes to scale out
OvercloudControlFlavor	The flavor to use for Controller nodes (**control**)
ComputeCount	The number of Compute nodes to scale out
OvercloudComputeFlavor	The flavor to use for Compute nodes (**compute**)

For example, to configure the overcloud to deploy three nodes for each role (Controller and Compute), add the following to your **parameter_defaults**:

```
parameter_defaults:
   ControllerCount: 3
   ComputeCount: 3
   OvercloudControlFlavor: control
   OvercloudComputeFlavor: compute
```

> **NOTE**
>
> See Creating the Overcloud with the CLI Tools from the Director Installation and Usage guide for a more complete list of Heat template parameters.

3.8. CREATING THE OVERCLOUD

The creation of the Overcloud requires additional arguments for the **openstack overcloud deploy** command. For example:

```
$ openstack overcloud deploy --templates -e -e
/home/stack/templates/puppet-ceph-external.yaml --ntp-server pool.ntp.org
```

The above command uses the following options:

- **--templates** - Creates the Overcloud from the default Heat template collection (namely, **/usr/share/openstack-tripleo-heat-templates/**).

- **-e /home/stack/templates/puppet-ceph-external.yaml** - Adds an additional environment file to the Overcloud deployment. In this case, it is the storage environment file containing the configuration for the existing Ceph Storage Cluster.

- **--ntp-server pool.ntp.org** - Sets our NTP server.

For a full list of options, run:

```
$ openstack help overcloud deploy
```

For more information, see Creating the Overcloud with the CLI Tools in the Director Installation and Usage guide.

The Overcloud creation process begins and the director provisions your nodes. This process takes some time to complete. To view the status of the Overcloud creation, open a separate terminal as the **stack** user and run:

```
$ source ~/stackrc
$ heat stack-list --show-nested
```

This configures the Overcloud to use your external Ceph Storage cluster. Note that you manage this cluster independently from the Overcloud. For example, you scale the Ceph Storage cluster using the Ceph management tools and not through the OpenStack Platform director.

3.9. ACCESSING THE OVERCLOUD

The director generates a script to configure and help authenticate interactions with your Overcloud from the director host. The director saves this file (**overcloudrc**) in your **stack** user's home directory. Run the following command to use this file:

```
$ source ~/overcloudrc
```

This loads the necessary environment variables to interact with your Overcloud from the director host's CLI. To return to interacting with the director's host, run the following command:

```
$ source ~/stackrc
```

CHAPTER 4. CONCLUSION

This concludes the creation and configuration of Overcloud with Red Hat Ceph Storage. For general Overcloud post-creation functions, see Performing Tasks after Overcloud Creation in the Director Installation and Usage guide.

APPENDIX A. SAMPLE ENVIRONMENT FILE: CREATING A CEPH CLUSTER

The following custom environment file uses many of the options described throughout Chapter 2, *Creating an Overcloud with Ceph Storage Nodes*. This sample does not include any commented-out options. For an overview on environment files, see Environment Files (from the Advanced Overcloud Customization guide).

/home/stack/templates/storage-environment.yaml

```
parameter_defaults: //  ①
  CinderBackupBackend: ceph //  ②
  ExtraConfig:
    ceph::profile::params::osds: //  ③
      '/dev/sdc':
        journal: '/dev/sdb'
      '/dev/sdd':
        journal: '/dev/sdb'
    CephPools: //  ④
      volumes:
        size: 5
        pg_num: 128
        pgp_num: 128
  ControllerCount: 3 //  ⑤
  OvercloudControlFlavor: control
  ComputeCount: 3
  OvercloudComputeFlavor: compute
  CephStorageCount: 3
  OvercloudCephStorageFlavor: ceph-storage
  CephMonCount: 3
  OvercloudCephMonFlavor: ceph-mon
  CephMdsCount: 3
  OvercloudCephMdsFlavor: ceph-mds
  NeutronNetworkType: vxlan //  ⑥
```

① The **parameter_defaults** section modifies the default values for parameters in all templates. Most of the entries listed here are described in Section 2.6, "Enabling Ceph Storage in the Overcloud".

② If you are deploying the Ceph Object Gateway, you can use Ceph Object Storage (**ceph-rgw**) as a backup target. To configure this, set **CinderBackupBackend** to **swift**. See Section 2.9, "Deploy the Ceph Object Gateway" for details.

③ The **ceph::profile::params::osds::** section defines a custom disk layout, as described in Section 2.7, "Mapping the Ceph Storage Node Disk Layout".

④ The **CephPools** section sets custom attributes for any Ceph pool. This example sets custom **size**, **pg_num**, and **pgp_num** attributes for the **volumes** pool. See Section 2.12.1, "Assigning Custom Attributes to Different Ceph Pools" for more details.

⑤ For each role, the ***Count** parameters assign a number of nodes while the **Overcloud*Flavor** parameters assign a flavor. For example, **ControllerCount: 3** assigns 3 nodes to the Controller role, and **OvercloudControlFlavor: control** sets each of those roles to use the

control flavor. See Section 2.13, "Assigning Nodes and Flavors to Roles" for details.

> **NOTE**
>
> The **CephMonCount**, **CephMdsCount**, **OvercloudCephMonFlavor**, and **OvercloudCephMdsFlavor** parameters (along with the **ceph-mon** and **ceph-mds** flavors) will only be valid if you created a custom **CephMON** and **CephMds** role, as described in Section 2.4, "Deploying Other Ceph Services on Dedicated Nodes".

6 **NeutronNetworkType:** sets the network type that the **neutron** service should use (in this case, **vxlan**).

APPENDIX B. SAMPLE CUSTOM INTERFACE TEMPLATE: MULTIPLE BONDED INTERFACES

The following template is a customized version of **/usr/share/openstack-tripleo-heat-templates/network/config/bond-with-vlans/ceph-storage.yaml**. It features multiple bonded interfaces to isolate back-end and front-end storage network traffic, along with redundancy for both connections (as described in Section 2.11, "Configuring Multiple Bonded Interfaces Per Ceph Node"). It also uses custom bonding options (namely, **'mode=4 lacp_rate=1'**, as described in Section 2.11.1, "Configuring Bonding Module Directives").

/usr/share/openstack-tripleo-heat-templates/network/config/bond-with-vlans/ceph-storage.yaml (custom)

```
heat_template_version: 2015-04-30

description: >
  Software Config to drive os-net-config with 2 bonded nics on a bridge
  with VLANs attached for the ceph storage role.

parameters:
  ControlPlaneIp:
    default: ''
    description: IP address/subnet on the ctlplane network
    type: string
  ExternalIpSubnet:
    default: ''
    description: IP address/subnet on the external network
    type: string
  InternalApiIpSubnet:
    default: ''
    description: IP address/subnet on the internal API network
    type: string
  StorageIpSubnet:
    default: ''
    description: IP address/subnet on the storage network
    type: string
  StorageMgmtIpSubnet:
    default: ''
    description: IP address/subnet on the storage mgmt network
    type: string
  TenantIpSubnet:
    default: ''
    description: IP address/subnet on the tenant network
    type: string
  ManagementIpSubnet: # Only populated when including
environments/network-management.yaml
    default: ''
    description: IP address/subnet on the management network
    type: string
  BondInterfaceOvsOptions:
    default: 'mode=4 lacp_rate=1'
    description: The bonding_options string for the bond interface. Set
                 things like lacp=active and/or bond_mode=balance-slb
                 using this option.
    type: string
```

```
    constraints:
      - allowed_pattern: "^((?!balance.tcp).)*$"
        description: |
          The balance-tcp bond mode is known to cause packet loss and
          should not be used in BondInterfaceOvsOptions.
  ExternalNetworkVlanID:
    default: 10
    description: Vlan ID for the external network traffic.
    type: number
  InternalApiNetworkVlanID:
    default: 20
    description: Vlan ID for the internal_api network traffic.
    type: number
  StorageNetworkVlanID:
    default: 30
    description: Vlan ID for the storage network traffic.
    type: number
  StorageMgmtNetworkVlanID:
    default: 40
    description: Vlan ID for the storage mgmt network traffic.
    type: number
  TenantNetworkVlanID:
    default: 50
    description: Vlan ID for the tenant network traffic.
    type: number
  ManagementNetworkVlanID:
    default: 60
    description: Vlan ID for the management network traffic.
    type: number
  ControlPlaneSubnetCidr: # Override this via parameter_defaults
    default: '24'
    description: The subnet CIDR of the control plane network.
    type: string
  ControlPlaneDefaultRoute: # Override this via parameter_defaults
    description: The default route of the control plane network.
    type: string
  ExternalInterfaceDefaultRoute: # Not used by default in this template
    default: '10.0.0.1'
    description: The default route of the external network.
    type: string
  ManagementInterfaceDefaultRoute: # Commented out by default in this
template
    default: unset
    description: The default route of the management network.
    type: string
  DnsServers: # Override this via parameter_defaults
    default: []
    description: A list of DNS servers (2 max for some implementations)
that will be added to resolv.conf.
    type: comma_delimited_list
  EC2MetadataIp: # Override this via parameter_defaults
    description: The IP address of the EC2 metadata server.
    type: string

resources:
  OsNetConfigImpl:
```

```
type: OS::Heat::StructuredConfig
properties:
  group: os-apply-config
  config:
    os_net_config:
      network_config:
        -
          type: interface
          name: nic1
          use_dhcp: false
          dns_servers: {get_param: DnsServers}
          addresses:
            -
              ip_netmask:
                list_join:
                  - '/'
                  - - {get_param: ControlPlaneIp}
                    - {get_param: ControlPlaneSubnetCidr}
          routes:
            -
              ip_netmask: 169.254.169.254/32
              next_hop: {get_param: EC2MetadataIp}
            -
              default: true
              next_hop: {get_param: ControlPlaneDefaultRoute}
        -
          type: ovs_bridge
          name: br-bond
          members:
            -
              type: linux_bond
              name: bond1
              bonding_options: {get_param: BondInterfaceOvsOptions}
              members:
                -
                  type: interface
                  name: nic2
                  primary: true
                -
                  type: interface
                  name: nic3
            -
              type: vlan
              device: bond1
              vlan_id: {get_param: StorageNetworkVlanID}
              addresses:
                -
                  ip_netmask: {get_param: StorageIpSubnet}
        -
          type: ovs_bridge
          name: br-bond2
          members:
            -
              type: linux_bond
              name: bond2
              bonding_options: {get_param: BondInterfaceOvsOptions}
```

```
              members:
                -
                  type: interface
                  name: nic4
                  primary: true
                -
                  type: interface
                  name: nic5
            -
              type: vlan
              device: bond1
              vlan_id: {get_param: StorageMgmtNetworkVlanID}
              addresses:
                -
                  ip_netmask: {get_param: StorageMgmtIpSubnet}
outputs:
  OS::stack_id:
    description: The OsNetConfigImpl resource.
    value: {get_resource: OsNetConfigImpl}
```