

Dell EMC Ready Bundle for Red Hat OpenStack Platform

**Adding and Removing Compute and Storage Nodes
Version 10.0**



Dell EMC Converged Platforms and Solutions

Contents

- Trademarks.....3
- Notes, Cautions, and Warnings.....4

- Chapter 1: Executive Summary.....5
 - Intended Audience..... 6

- Chapter 2: Adding Nodes..... 7
 - Prerequisites.....8
 - Adding a Node..... 8

- Chapter 3: Removing Nodes..... 12
 - Removing a Compute Node..... 13
 - Removing a Storage Node..... 15

- Appendix A: Getting Help..... 21
 - Contacting Dell EMC.....22
 - References..... 22
 - To Learn More..... 22

Trademarks

Copyright © 2014-2017 Dell Inc. or its subsidiaries. All rights reserved.

Microsoft® and Windows® are registered trademarks of Microsoft Corporation in the United States and/or other countries.

Red Hat®, Red Hat Enterprise Linux®, and Ceph are trademarks or registered trademarks of Red Hat, Inc., registered in the U.S. and other countries. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries. Oracle® and Java® are registered trademarks of Oracle Corporation and/or its affiliates.

DISCLAIMER: 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.

The Midokura® name and logo, as well as the MidoNet® name and logo, are registered trademarks of Midokura SARL.

Notes, Cautions, and Warnings



A **Note** indicates important information that helps you make better use of your system.



A **Caution** indicates potential damage to hardware or loss of data if instructions are not followed.



A **Warning** indicates a potential for property damage, personal injury, or death.

This document is for informational purposes only and may contain typographical errors and technical inaccuracies. The content is provided as is, without express or implied warranties of any kind.

Chapter 1

Executive Summary

Topics:

- *Intended Audience*

One of the common needs of a cloud platform is to scale resources as compute and storage needs expand or contract. Over time there is often a need to add additional Compute resources, as the required VM (Virtual Machine) growth has exceeded planned capacity. Likewise, Storage resources must expand to keep up with demand. At other times, there may be a need to remove a Compute or Storage resource for upgrades, repairs, etc.

This technical guide explains the process of adding or removing the following nodes to or from an existing Dell EMC Ready Bundle for Red Hat OpenStack Platform cluster:

- Compute nodes:
 - Dell EMC PowerEdge R430
 - Dell EMC PowerEdge R630
 - Dell EMC PowerEdge R730
- Storage node:
 - Dell EMC PowerEdge R730xd

Compute nodes are used for hosting VMs in the Dell EMC Ready Bundle for Red Hat OpenStack Platform version 10.0.

Storage nodes are used for hosting Red Hat Ceph Storage version 2.0 which manages block, image, and ephemeral storage in the Dell EMC Ready Bundle for Red Hat OpenStack Platform version 10.0.

The following documentation from Red Hat should be referenced when using the steps in this document:

<https://access.redhat.com/documentation/en/red-hat-openstack-platform/9/paged/director-installation-and-usage/chapter-9-scaling-the-overcloud>

<https://access.redhat.com/documentation/en/red-hat-ceph-storage/1.3/paged/administration-guide/chapter-6-managing-cluster-size#removing-osds-manual>

Intended Audience

This guide is written for OpenStack administrators or deployment engineers who are responsible for installation and ongoing operation of OpenStack clusters. It assumes that the reader is familiar with:

- OpenStack
- Red Hat Enterprise Linux (RHEL)
- Red Hat OpenStack Platform (RHOSP) documentation
- Networking and system administration

Chapter 2

Adding Nodes

Topics:

- [Prerequisites](#)
- [Adding a Node](#)

This section describes prerequisites and procedures to add a Compute or Storage node to a Dell EMC Ready Bundle for Red Hat OpenStack Platform cluster.

Prerequisites

The following prerequisites must be met:

- Compute and Storage nodes' RAID and BIOS settings have been configured using the Dell EMC Open Source Hardware Configuration Toolkit utility
- Dell EMC Ready Bundle for Red Hat OpenStack Platform version 10.0 installed
- [Dell EMC Ready Bundle for Red Hat OpenStack Platform Deployment Guide](#) available for reference



Note: All nodes in the same roles must be of the same server models, with identical HDD, RAM, and NIC configurations. So, all Controller nodes must be identical to each other; all Compute nodes must be identical to each other; and so on. See the [Dell EMC Ready Bundle for Red Hat OpenStack Platform Reference Architecture](#) for configuration options for each node role.

Adding a Node

To add a Compute or Storage node to a Dell EMC Ready Bundle for Red Hat OpenStack Platform cluster:



Note: If you add a new Compute node, **and** your existing installation uses Instance HA technology, you will need to do a manual install of Instance HA for the new Compute node after it is added to the cluster. See the instructions in the *Scripted HA Installation* section of the [Dell EMC Ready Bundle for Red Hat OpenStack Platform Deployment Guide](#).

1. Log onto the Director Node as the *admin_user* user (or user as configured in *Director Node Configuration Parameters*, in the [Dell EMC Ready Bundle for Red Hat OpenStack Platform Deployment Guide](#)).
2. Navigate to the directory in which the Dell EMC Ready Bundle for Red Hat OpenStack Platform was deployed (e.g., the */home/admin_user/pilot* directory.)
3. Run the following command specifying the iDRAC credentials (replace example below with actual values) and the IP address of the node being added:

```
./discover_nodes/discover_nodes.py -u root \
-p 'root_password' 192.168.110.122 \
> ~/newnodes.json
```

4. If more than one node is being added, then you can specify an IP range using two IP addresses separated by a dash:

```
./discover_nodes/discover_nodes.py -u root \
-p '<root_password>' \
192.168.110.122-192.168.110.126 \
> ~/newnodes.json
```

5. When discovery is complete, examine *~/newnodes.json* to verify that it contains an entry for every cluster node to be provisioned, and no other nodes:

```
{
  "nodes": [
    {
      "pm_addr": "192.168.110.50",
      "pm_password": "calvin",
      "pm_type": "pxe_ipmitool",
      "pm_user": "root"
    }
  ]
}
```



```
}
```

The `pm_user`, `pm_password`, and `pm_addr` attribute values are examples - the actual values should conform to your node specifications. The other attributes will be adjusted during the node introspection phase.

6. Perform initial iDRAC configuration for each new node.

- a. This includes setting the appropriate NIC port to PXE boot for provisioning, as well as performing other basic iDRAC configuration.

```
$ cd ~/pilot
$ ./config_idrac.py <idrac_ip> -n ~/newnodes.json
```

For example:

```
$ ./config_idrac.py 192.168.110.50 -n ~/newnodes.json
```

7. Execute the following command to register the new nodes:

```
$ ./import_nodes.py -n ~/newnodes.json
```

8. Identify the new nodes by UUID, and start the introspection process. For example:



Note: Your node UUIDs will differ.


```
$ ironic node-list
+-----+-----+-----+-----+
+-----+-----+-----+-----+
+-----+
| UUID                                | Name | Instance UUID |
| Power State | Provisioning State | Maintenance |
+-----+-----+-----+-----+
+-----+
| df221970-94d8-4ed0-9f68-835bd8e166d2 | None |                | |
| 855240f6-4943-4348-830f-2436b585d7e4 | power on | active          |
| False                                |         |                 |
| af5ab2bf-6b82-4587-92ba-cfb6f5e7a759 | None | bceb5b21-a068-44b4-9444- |
| a89f1f4blace | power on | active          | False          |
| d07bb1b2-13be-4762-8a99-4cef9ec26faf | None | df64c706-      |
| b4d5-4cb7-8eab-65ec4ca3b810 | power on | active          | False          |
| 40a6b131-c17b-44ed-ac32-76e6f51d3aea | None | d6f05ce1-92ca-4ee8-adcd- |
| fe7e2e423d69 | power on | active          | False          |
| f27bfbda-b69a-4d23-be07-93284d1c33f5 | None | 309a49d2-      |
| c50a-46e5-89b1-9bd3906e8cc7 | power on | active          | False          |
| f5787449-2e37-4737-93d0-84b18117f649 | None | aed26822-b803-43fc-b293- |
| f66e98b851ab | power on | active          | False          |
| f119bc65-9e51-4d0d-ad45-58916f7e8097 | None | ce43d39d-24f7-4e19-      |
| ac2e-11fa5417b46a | power on | active          | False          |
| 70c68e66-6465-4fd0-a91c-c4210045641e | None | ca0ca01c-8312-4752-8590- |
| ffdfa14c7192 | power on | active          | False          |
| f06aed59-b987-4634-83d4-5bed8d35b0c9 | None | None            |
| None |         | available      |                |
+-----+-----+-----+-----+
+-----+-----+-----+-----+
+-----+
```

9. After registering the new node, launch the introspection process:

```
$ openstack baremetal node manage <NODE UUID>
```

```
$ openstack overcloud node introspect <NODE UUID> --provide
```

10. Repeat the commands in Step 9 for each new node.

 **Caution:** You must not proceed until introspection is complete on each node.

11. Each node must then be assigned a role within the cluster. Use the `assign_role.py` script to assign roles to node:

```
$ ./assign_role.py -n ~/newnodes.json <IP> <role>
```

a. For `<IP>`, substitute the IP address of the node's iDRAC.

b. For `<role>`, substitute the role that the node will perform (*compute* or *storage*).

For example:

```
$ ./assign_role.py -n ~/newnodes.json 192.168.110.50 compute-3
```

12. Repeat the command in Step 11 for each new node, using its iDRAC IP address and node type.

13. The boot images used during deployment must be set for the new nodes. To get the UUIDs of the images:

```
$ glance image-list
```

The display returns a list similar to this:

```
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| ID                               | Name                               | Disk |
| Format | Container Format | Size          | Status |
+-----+-----+-----+-----+-----+
| 5f4bfcd7-24d3-436f-9c59-161ad862597e | bm-deploy-kernel               | aki  |
| aki | 5153184 | active |
| 1bf9ed8c-82c1-4ab0-ae54-f0067cc2abb9 | bm-deploy-ramdisk              | ari  |
| ari | 341547616 | active |
| 7fca165f-cd6b-4173-86a7-47b147c00320 | overcloud-full                 | qcow2 |
| bare | 1019496448 | active |
| 94a365a2-bc82-4fa9-a2ec-13670050d454 | overcloud-full-initrd          | ari  |
| ari | 40323665 | active |
| c136dbel-2960-492c-a1eb-3f7048b08b5c | overcloud-full-vmlinuz         | aki  |
| aki | 5153184 | active |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
```

14. Set these UUIDs for each new nodes' `deploy_kernel` and `deploy_ramdisk` settings:

```
$ ironic node-update <UUID> add \
driver_info/deploy_kernel='<kernel_image_ID>'
$ ironic node-update <UUID> add \
driver_info/deploy_ramdisk='<ramdisk_image_ID>'
```

For example:

```
$ ironic node-update f06aed59-b987-4634-83d4-5bed8d35b0c9 add \
driver_info/deploy_kernel='5f4bfcd7-24d3-436f-9c59-161ad862597e'
$ ironic node-update f06aed59-b987-4634-83d4-5bed8d35b0c9 add \
driver_info/deploy_ramdisk='1bf9ed8c-82c1-4ab0-ae54-f0067cc2abb9'
```

15. Repeat the commands in Step 14 for each new node.

16. The new nodes are ready to be deployed. Specify the new total number of Compute or Storage nodes to the `deploy_overcloud.py` script, to have the cluster updated and the new nodes added and configured:

```
$ ./deploy_overcloud.py --computes 4 --storage 3 --vlans 201:250
```

- a. Ensure that:

- You include all environment files and options from your initial Overcloud creation. This includes the same scale parameters for non-Compute nodes.
- If the `--static_ip` option was used in the original deployment, make sure to update the `~/pilot/templates/static-ip-environment.yaml` file to add the new static IP address for the added node.

- b. To extract the original invocation of the `deployovercloud.py` script, execute the following command on the Solution Admin Host:

```
# grep \deploy-overcloud.py $(ls -tr /auto_results/ \
  deployer.log.* | tail -1) | tail -1 | sed -e \
  's/ ^.*running : //'
```



Note: Running the `deploy_overcloud.py` script will take some time, as the OS will be installed and the OpenStack software deployed and configured on the additional node(s).

17. If a Storage node was added, attach the new Storage node to Red Hat Storage Console:

- a. Execute the Red Hat Storage Console initialization scripts that enables you to use Red Hat Storage Console services:

```
# ./config_rhscn.py <rhscn_node_ip> <root_password>
```

- b. Log into the Red Hat Storage Console web interface at http://<storage_console_hostname>/skyring.
c. Add the new Storage node to the Red Hat Storage Console group.

The new Storage node is now ready for use. The newly-added OSDs will automatically be used by existing Ceph Storage pools that were created in the initial solution deployment.

18. If a Compute node was added to the cluster, add the Compute node to Instance HA:

- a. Execute the `install-instanceHA.py` script to add the Compute node to the Instance HA configuration:

```
# ./install-instanceHA.py --compute <compute_node_provisioning_IP>
```

19. If a Controller node was added to the cluster, add the Controller node to Instance HA:

- a. Execute the `install-instanceHA.py` script to add the Controller node to the InstanceHA configuration:

```
# ./install-instanceHA.py --controller <compute_node_provisioning_IP>
```

The new nodes, and reconfigured cluster, are now ready for use.

Chapter

3

Removing Nodes

Topics:

- [Removing a Compute Node](#)
- [Removing a Storage Node](#)


This section describes prerequisites and procedures to remove a Compute or Storage node to a Dell EMC Ready Bundle for Red Hat OpenStack Platform cluster.

Removing a Compute Node

Prerequisite:

- Instances must be migrated from the Compute node being removed to another active Compute node.

To remove a Compute node from a Dell EMC Ready Bundle for Red Hat OpenStack Platform cluster:

 **Note:** If you remove a Compute node, **and** your existing installation uses Instance HA technology, you will need to do a manual uninstall of Instance HA of the Compute node from the cluster. See the instructions in the *Scripted HA Installation* section of the [Dell EMC Ready Bundle for Red Hat OpenStack Platform Deployment Guide](#)).

- Log onto the Director Node as the *admin_user* user (or user as configured in *Director Node Configuration Parameters*, in the [Dell EMC Ready Bundle for Red Hat OpenStack Platform Deployment Guide](#)).
- Source your *overcloudrc* file.

```
# source ~/overcloudrc
```

- Identify the Compute node that will be removed, by executing the following command to get a list of the host names:

```
# nova service-list
```

- Disable the Compute nodes service so that no new instances can be started from that node:

```
# nova service-disable <compute hostname> nova-compute
```

- Source the Undercloud *stackrc* file:

```
# source ~/stackrc
```

- Identify the ID of the Overcloud stack:

```
# openstack stack list
```

Example truncated output:

ID	Stack Name	...
4aefad19-5bc6-476b-9582-80dca79f2825	overcloud	...

- Identify the Nova ID of the Compute node to delete:

```
# nova list
```

Example truncated output:

ID	Name	...
2e19c09b-4fc1-456b-8d23-91dff107f5c3	overcloud-cephstorage-0	...
901bf060-4cdf-4021-ac52-32c97001659f	overcloud-cephstorage-1	...
2cc1159d-50bc-404e-9ec0-767ffdd1446b	overcloud-cephstorage-2	...
99c8def0-35ce-4162-885b-23f85426f41c	overcloud-cephstorage-3	...
52f28960-1a91-4955-847b-9969b84067d2	overcloud-compute-0	...
eef7edf3-ae98-4921-88ff-7fc296e8bef7	overcloud-compute-1	...

cb68faf8-ba6c-405d-81ab-c8582138f536	overcloud-compute-2	...
143b488d-6e2a-44d0-aca3-c9a7c50b8035	overcloud-controller-0	...
ea5378d6-8778-4d7a-950e-748b450710dd	overcloud-controller-1	...
fbcd5322-d79b-485d-80ad-86552bddbafa	overcloud-controller-2	...
+-----+-----+		...

8. Identify the Ironic ID associated with the nova ID of the Compute node to delete, by executing the following command:



Note: There is a single space between the quotation marks in the command below.

```
# ironic node-list | grep [Nova_ID] | cut -d " " -f 2
```

Example truncated output:

```
59ad7668-8ba2-4972-99a8-e080dd06803f
```

9. Gather the path to the templates and environmental files used for the initial deployment:

```
# cd ~/pilot
# cat overcloud_deploy_cmd.log
```

10. Execute the following command to delete the Compute node and update the deployment:

```
# openstack overcloud node delete --stack [STACK_ID] \
--templates [TEMPLATE_DIRECTORY] \
-e [ENVIRONMENT_FILE]... \
[Nova_ID]
```

For example:

```
# openstack overcloud node delete --stack
4aefad19-5bc6-476b-9582-80dca79f2825
--templates ~/pilot/templates/overcloud
-e ~/pilot/templates/overcloud/environments/network-isolation.yaml
-e ~/pilot/templates/network-environment.yaml
-e ~/pilot/templates/static-ip-environment.yaml
-e ~/pilot/templates/static-vip-environment.yaml
-e ~/pilot/templates/node-placement.yaml
-e ~/pilot/templates/overcloud/environments/storage-environment.yaml
-e ~/pilot/templates/dell-environment.yaml
-e ~/pilot/templates/overcloud/environments/puppet-pacemaker.yaml
-e ~/pilot/templates/dell-cinder-backends.yaml
52f28960-1a91-4955-847b-9969b84067d2
```

Example output:

```
Deleting the following nodes from stack R11b:
- d28140b0-2acf-4ace-b427-844baaf3d656
Started Mistral Workflow. Execution ID: 04221bcf-542a-48c4-bf76-
f35808b139a6
```

11. Execute the `nova list` command repeatedly until you no longer see the Compute node in the list.
12. Source your `overcloudrc` file:

```
# source ~/overcloudrc
```

13. Identify the ID of the removed Compute node:

```
# nova service-list
```

14. Remove the Compute node from the OpenStack service:

```
# nova service-delete [Id]
```

15. Remove the Compute nodes Open vSwitch agent:

```
# neutron agent-list
# neutron agent-delete [ Open vSwitch agent id]
```

16. Delete the Compute node from Ironic by executing the following command:

```
# ironic node-delete [Ironic_ID]
```

You can now remove the Compute node from the environment.

Removing a Storage Node

Prerequisite:

- There must be at least 3 Storage nodes AFTER the removal of one of the Storage nodes. This ensures that the minimum data redundancy is maintained.

To remove a Storage node from a Dell EMC Ready Bundle for Red Hat OpenStack Platform cluster:

1. Log onto the Director Node as the *admin_user* user (or user as configured in *Director Node Configuration Parameters*, in the [Dell EMC Ready Bundle for Red Hat OpenStack Platform Deployment Guide](#)).
2. Source the undercloud `stackrc` file:

```
# source ~/stackrc
```

3. Identify the Storage node that will be removed, by executing the following command to get a list of the host names:

```
# nova list
```

4. Initiate an `ssh` session to the Storage node to be removed.

- a. Change to the *root* user:

```
$ sudo -i
```

5. Identify the OSDs that need to be removed by executing the following command:

```
#root@cephstorage-3 [/] ceph osd tree
```

To locate the OSDs associated with the Storage node to be removed, two essential pieces of information are required:

- The node number, indicated by the number in the first column
- The OSD associated with the node, indicated by *osd.n* in the third column

Example truncated output:

0	1.07999	osd.0	up	1.00000	1.00000
4	1.07999	osd.4	up	1.00000	1.00000
7	1.07999	osd.7	up	1.00000	1.00000



Note: Although the OSDs are listed sequentially per node, the numeric identifiers may or may not be contiguous.

6. Temporarily disable scrubbing by executing the following commands:

```
# ceph osd set noscrub
# ceph osd set nodeep-scrub
```

7. Remove every OSD on that storage node from the cluster.



Note: OSDs should be removed one at a time, in order to minimize impact on the rest of the cluster.

- a. Place the OSD in the *out* state. This must be performed for **all** OSDs identified in Step 5:

```
#root@cephstorage-3 [/] ceph osd out osd.0
```

- b. Wait for the data to migrate off the OSD:

- a. Execute the following command to monitor the migration and rebalancing activity:

```
#root@cephstorage-3 [/] ceph -w
```

- b. Watch the output until the status is **active+clean**:

```
2016-10-03 16:11:23.045340 mon.0 [INF] from='client.?
192.168.170.18:0/1163991'
entity='client.admin' cmd=[{"prefix": "osd out", "ids":
["osd.0"]}]: dispatch

2016-10-03 16:11:23.240194 mon.0 [INF] pgmap v34180: 3648 pgs: 3648
active+clean;
45660 kB data, 4013 MB used, 53471 GB / 53475 GB avail

2016-10-03 16:11:23.245291 mon.0 [INF] from='client.?
192.168.170.18:0/1163991'
entity='client.admin' cmd=[{"prefix": "osd out", "ids":
["osd.0"]}]: finished

2016-10-03 16:11:40.279607 mon.0 [INF] pgmap v34190: 3648 pgs:
3648 active+clean;
45660 kB data, 3934 MB used, 52922 GB / 52926 GB avail
```

- c. To exit the utility, press **[Ctrl-C]**.

- c. Stop the OSD by executing the following commands:

```
#root@cephstorage-3 [/] systemctl disable ceph-osd@0
#root@cephstorage-3 [/] systemctl stop ceph-osd@0
```

A message similar to the following is displayed:

```
=== osd.0 ===
Stopping Ceph osd.0 on cephstorage-3...kill 66546...kill 66546...done
```

After running the stop command, the OSDs will be down. Now they need to be removed from the CRUSH map.

- d. Remove the OSD from the storage node to be removed by executing the following command:

```
#root@cephstorage-3 [/] ceph osd crush remove osd.0
```

A message similar to the following is displayed:

```
removed item id 0 name 'osd.0' from crush map
```



- e. Delete the OSD's authentication token by executing the following command:

```
#root@cephstorage-3 [/] ceph auth del osd.0
```

The following message is displayed:

```
updated
```

- f. Delete the OSD by executing the following command:

 **Note:** The syntax for specifying the OSD is just the numeric identifier (e.g., *N*) and not *osd.N*.

```
#root@cephstorage-3 [/] ceph osd rm 0
```

A message similar to the following is displayed:

```
removed osd.0
```

- g. After removing an OSD check to verify the cluster is not nearing the near-full ratio:

```
# ceph -s
# ceph df
```

- h. Repeat from Step 7a until all OSDs on the Storage node have been removed from the cluster.

8. Remove the Storage node from the CRUSH map by executing the following command:

```
#root@cephstorage-3 [/] ceph osd crush remove <storage-node-to-be-removed>
```

9. Determine the ID of the Overcloud:

- a. Log onto the Director Node as the *admin_user* (or the user configured in *Director Node Configuration Parameters*, in the Dell EMC Ready Bundle for Red Hat OpenStack Platform Deployment Guide).
- b. Source the *undercloud stackrc* file:

```
# source ~/stackrc
```

- c. Identify the ID of the Overcloud stack by executing the following command:

```
# openstack stack list
```

Example truncated output:

```
+-----+-----+-----+
| ID                               | Stack Name | ...
+-----+-----+-----+
| 4aefad19-5bc6-476b-9582-80dca79f2825 | overcloud  | ...
+-----+-----+-----+
```

10. Identify the Nova ID of the Storage node to delete, by executing the following command:

```
# nova list
```

Example truncated output:

 **Note:** IDs, names, and networks will appear differently in actual output.

```
+-----+-----+-----+
| ID                               | Name       | ...
+-----+-----+-----+
| 2e19c09b-4fc1-456b-8d23-91dff107f5c3 | overcloud-cephstorage-0 | ...
+-----+-----+-----+
```

901bf060-4cdf-4021-ac52-32c97001659f	overcloud-cephstorage-1	...
2cc1159d-50bc-404e-9ec0-767ffdd1446b	overcloud-cephstorage-2	...
99c8def0-35ce-4162-885b-23f85426f41c	overcloud-cephstorage-3	...
52f28960-1a91-4955-847b-9969b84067d2	overcloud-compute-0	...
eef7edf3-ae98-4921-88ff-7fc296e8bef7	overcloud-compute-1	...
cb68faf8-ba6c-405d-81ab-c8582138f536	overcloud-compute-2	...
143b488d-6e2a-44d0-aca3-c9a7c50b8035	overcloud-controller-0	...
ea5378d6-8778-4d7a-950e-748b450710dd	overcloud-controller-1	...
fbcd5322-d79b-485d-80ad-86552bddbafa	overcloud-controller-2	...
+-----+-----+-----+		...

11. Identify the Ironic ID associated with the nova ID of the Storage node to delete, by executing the following command:



Note: There is a single space between the quotation marks in the command below.

```
# ironic node-list | grep [Nova_ID] | cut -d " " -f 2
```

Example truncated output:

```
e4b5ca84-823a-42ec-a0d5-1ce091952c11
```

12. Gather the path to the templates and the environmental files used for the initial deployment.

```
# cd ~/pilot
# cat overcloud_deploy_cmd.log
```

13. Execute the following command to delete the Storage node and update the deployment:

```
# openstack overcloud node delete --stack [STACK_ID] \
--templates [TEMPLATE_DIRECTORY] \
-e [ENVIRONMENT_FILE]... \
[Nova_ID]
```

For example:

```
openstack overcloud node delete --stack
4aefad19-5bc6-476b-9582-80dca79f2825
--templates ~/pilot/templates/overcloud
-e ~/pilot/templates/overcloud/environments/network-isolation.yaml
-e ~/pilot/templates/network-environment.yaml
-e ~/pilot/templates/static-ip-environment.yaml
-e ~/pilot/templates/static-vip-environment.yaml
-e ~/pilot/templates/node-placement.yaml
-e ~/pilot/templates/overcloud/environments/storage-environment.yaml
-e ~/pilot/templates/dell-environment.yaml
-e ~/pilot/templates/overcloud/environments/puppet-pacemaker.yaml
-e ~/pilot/templates/dell-cinder-backends.yaml
-e ~/pilot/templates/overcloud/environments/ceph-radosgw.yaml
99c8def0-35ce-4162-885b-23f85426f41c
```

Example output:

```
Deleting the following nodes from stack R11b:
- d28140b0-2acf-4ace-b427-844baaf3d656
Started Mistral Workflow. Execution ID: 04221bcf-542a-48c4-bf76-
f35808b139a6
```

14. Execute the `nova list` command repeatedly until you no longer see the Storage node in the list.

15. Validate proper Red Hat Ceph Storage operation after node removal:

- a. Close the `ssh` session to the Storage node, returning to the session on the Director Node.

- b. Change to the *admin_user*'s home directory.
- c. Source your *overcloudrc* file:

```
$ source ~/overcloudrc
```

- d. Create a new Red Hat Ceph Storage object by executing the following commands:

```
$ swift post foo_test
$ swift list
foo_test
```

- e. Upload a large file into the new Red Hat Ceph Storage container by executing the following commands:

```
$ ls -alpS | head -2 | tail -1
-rw-rw-r--. 1 stack stack 13200896 Sep  8 2014 cirros-0.3.3-x86_64-disk.img
$ swift upload foo_test cirros-0.3.3-x86_64-disk.img
cirros-0.3.3-x86_64-disk.img
```

- f. Verify that the large file exists in the container by executing the following commands:

```
$ swift list
foo_test
$ swift list foo_test
cirros-0.3.3-x86_64-disk.img
```

- g. Validate proper storage of the object by downloading the test file to a different directory and comparing the contents:

```
$ swift list
$ cd /tmp
$ swift download foo_test cirros-0.3.3-x86_64-disk.img
cirros-0.3.3-x86_64-disk.img [auth 0.255s, headers 0.279s, total 0.342s,
150.484 MB/s]
$ diff cirros-0.3.3-x86_64-disk.img ~/cirros-0.3.3-x86_64-disk.img; echo
$?
0
```

- 16. Delete the Storage node from Ironic by executing the following command:

```
# ironic node-delete [Ironic_ID]
```

- 17. Initiate an *ssh* session to any existing Storage node.

- 18. Enable scrubbing by executing the following commands:

```
[heat-admin@r11b-cephstorage-0 ~]$ sudo -i
[root@r11b-cephstorage-0 ~]# ceph osd unset noscrub
unset noscrub
[root@r11b-cephstorage-0 ~]# ceph osd unset nodeep-scrub
unset nodeep-scrub
```

- 19. Verify that the Ceph Storage health status is *HEALTH_OK*:

```
root@r11b-cephstorage-0 ~]# ceph -s
cluster 68bc488a-2333-11e7-ab43-52540032fd4d
health HEALTH_OK
monmap e2: 3 mons at {r11b-controller-0=192.168.170.11:6789/0,r11b-controller-1=192.168.170.12:6789/0,r11b-controller-2=192.168.170.13:6789/0}
election epoch 20, quorum 0,1,2 r11b-controller-0,r11b-controller-1,r11b-controller-2
osdmap e337: 36 osds: 36 up, 36 in
```

```
flags sortbitwise
pgmap v2798: 2816 pgs, 16 pools, 12845 kB data, 189 objects
            1821 MB used, 130 TB / 130 TB avail
            2816 active+clean
```

You can now remove the Storage node from the environment.

Appendix

A

Getting Help

Topics:

- [Contacting Dell EMC](#)
- [References](#)

This appendix details contact and reference information for the Dell EMC Ready Bundle for Red Hat OpenStack Platform.

Contacting Dell EMC

For customers in the United States, call 800-WWW-DELL (800-999-3355).



Note: If you do not have an active Internet connection, you can find contact information on your purchase invoice, packing slip, bill, or Dell EMC product catalog.

Dell EMC provides several online and telephone-based support and service options. Availability varies by country and product, and some services may not be available in your area. To contact Dell EMC for sales, technical support, or customer service issues:

1. Visit dell.com/support.
2. Click your country/region at the bottom of the page. For a full listing of country/region, click **All**.
3. Click **All Support** from the **Support** menu.
4. Select the appropriate service or support link based on your need.
5. Choose the method of contacting Dell EMC that is convenient for you.

References

Additional information can be obtained at <http://www.dell.com/en-us/work/learn/openstack-cloud> or by e-mailing openstack@dell.com.

If you need additional services or implementation help, please contact your Dell EMC sales representative.

To Learn More

For more information on the Dell EMC Ready Bundle for Red Hat OpenStack Platform visit <http://www.dell.com/learn/us/en/04/solutions/red-hat-openstack>.

Copyright © 2014-2017 Dell Inc. or its subsidiaries. All rights reserved. Trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Specifications are correct at date of publication but are subject to availability or change without notice at any time. Dell EMC and its affiliates cannot be responsible for errors or omissions in typography or photography. Dell EMC's Terms and Conditions of Sales and Service apply and are available on request. Dell EMC service offerings do not affect consumer's statutory rights.

Dell EMC, the DELL EMC logo, the DELL EMC badge, and PowerEdge are trademarks of Dell Inc.