Dell EMC Ready Bundle for Red Hat OpenStack Platform

Operations Guide

Version 10.1

Dell EMC Converged Platforms and Solutions

Contents

List of Figures....v

List of Tables....vii

Trademarks....ix

Glossary....x

Notes, Cautions, and Warnings....xiv

Chapter 1: Executive Summary....15
- Intended Audience....16

Chapter 2: Updating BIOS and Firmware....17
- Overview....18
  - Summary....18
  - Prerequisites....18
- BIOS and Firmware Update Procedures....18
  - Download Files....18
  - Update the SAH Node's BIOS and Firmware....19
  - Update the Controller Nodes' BIOS and Firmware....23
  - Update the Ceph Storage Nodes' BIOS and Firmware....27
  - Update the Compute Nodes' BIOS and Firmware....31

Chapter 3: Adding and Removing Nodes....35
- Overview....36
- Adding Nodes....36
  - Prerequisites....37
  - Adding a Node....37
- Removing Nodes....41
  - Prerequisites....42
  - Removing a Compute Node....42
  - Removing a Storage Node....44

Chapter 4: Updating the Solution....51
- Overview....52
  - Dependencies....52
- Dell EMC Ready Bundle for Red Hat OpenStack Platform v10.0.1 Minor Undercloud/Overcloud Update....52
  - Minor Update Procedure....52
  - Prerequistes....52
  - Caveats and Impacts for Minor Update....53
  - Deploy the Update/Upgrade Tools....53
  - Remove Instance High Availability....53
  - Run the Update Script....55
  - Reboot the Director Node....55
  - Logging and Lock Files....56

Finalize the Update....56
Install Instance High Availability....56

Appendix A: References....59
To Learn More....60

List of Figures

Figure 1: Lifecycle Controller....20
Figure 2: Select Update Repository....21
Figure 3: Enter Access Details....21
Figure 4: Select Updates....22
Figure 5: Lifecycle Controller....24
Figure 6: Select Update Repository....25
Figure 7: Enter Access Details....25
Figure 8: Select Updates....26
Figure 9: Lifecycle Controller....28
Figure 10: Select Update Repository....29
Figure 11: Enter Access Details....29
Figure 12: Select Updates....30
Figure 13: Lifecycle Controller....32
Figure 14: Select Update Repository....33
Figure 15: Enter Access Details....33
Figure 16: Select Updates....34

List of Tables

Table 1: Logging Levels.....54

Trademarks

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.

Glossary

API

Application Programing Interface is a specification that defines how software components can interact.

BMC/IDRAC Enterprise

Baseboard management controller. An on-board microcontroller that monitors the system for critical events by communicating with various sensors on the system board, and sends alerts and log events when certain parameters exceed their preset thresholds.

BOSS

The Boot Optimized Storage Solution (BOSS) enables customers to segregate operating system and data on server-internal storage. This is helpful in the Hyper-Converged Infrastructure (HCI) and Software Defined Storage (SDS) arenas, to separate operating system drives from data drives, and implement hardware RAID mirroring (RAID1) for OS drives.

Bundle

A customer-orderable solution that consists of:

- All server, network, and storage hardware needed to install and operate the solution as outlined
- All necessary solution software licenses needed to install and operate the solution as outlined

CDH

Cloudera Distribution for Apache Hadoop

Cloud Computing

See *http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf*

Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.

Cluster

A set of servers dedicated to OpenStack that can be attached to multiple distribution switches.

Compute Node

The hardware configuration that best supports the hypervisor server or Nova compute roles.

DevOps

Development Operations (DevOps) is an operational model for managing data centers using improved automated deployments, shortened lead times between fixes, and faster mean time to recovery. See *https://en.wikipedia.org/wiki/DevOps*.

DIMM

Dual In-line Memory Module

DNS

The domain name system (DNS) defines how Internet domain names are located, and translated into Internet Protocol (IP) addresses.

FQDD

A fully qualified device descriptor (FQDD) is a method used to describe a particular component within a system or subsystem, and is used for system management and other purposes.

FQDN

A fully qualified domain name (FQDN) is the portion of an Internet Uniform Resource Locator (URL) that fully identifies the server to which an Internet request is addressed. The FQDN includes the second-level domain name, such as "dell.com", and any other levels as required.

GUI

Graphical User Interface - A visual interface for human interaction with the software, taking inputs and generating easy to understand visual outputs.

Hypervisor

Software that runs virtual machines (VMs).

IaaS

Infrastructure as a Service.

Infrastructure Node

Systems that handle the control plane and deployment functions.

ISV

Independent Software Vendor.

JBOD

Just a Bunch of Disks

LAG

Link Aggregation Group.

LOM

LAN on motherboard.

LVM

Logical Volume Management

ML2

The Modular Layer 2 plug-in is a framework that allows OpenStack to utilize different layer 2 networking technologies.

NFS

The Network File System (NFS) is a distributed filesystem that allows a computer user to access, manipulate, and store files on a remote computer, as though they resided on a local file directory.

NIC

Network Interface Card

Node

One of the servers in the cluster.

NUMA

Non-Uniform Memory Access

Overcloud

The functional cloud that is available to run guest VMs and workloads.

Pod

An installation comprised of three racks, and consisting of servers, storage, and networking.

REST

REST - Representational State Transfer (also ReST). Relies upon stateless, client-server, cacheable communications protocol to access the API.

RHOSP

Red Hat OpenStack Platform

RPC

Remote Procedure Call

SAH

The Solution Admin Host (SAH) is a physical server that supports VMs for the Undercloud machines needed for the cluster to be deployed and operated.

SDS

Software-defined storage (SDS) is an approach to computer data storage in which software is used to manage policy-based provisioning and management of data storage, independent of the underlying hardware.

SDN

Software-defined Network (SDN) is where the software will define, create, use and destroy different networks as needed.

Storage Node

The hardware configuration that best supports SDS functions such as Red Hat Ceph Storage.

ToR

Top-of-rack switch/router.

U

U used in the definition of the size of server, example 1U or 2U. A "U" is a unit of measure equal to 1.75 inches in height.

Undercloud

The Undercloud is the system used to control, deploy, and monitor the Overcloud - it is a single node OpenStack deployment completely under the administrators control. The Undercloud is *not* HA configured.

VLT

A Virtual Link Trunk (VLT) is the combined port channel between an attached device (ToR switch) and the VLT peer switches.

VLTi

A Virtual Link Trunk Interconnect (VLTi) is an interconnect used to synchronize states between the VLT peer switches. Both endpoints must be the same speed, i.e. 40Gb → 40Gb; 1G interfaces are not supported.

VM

Virtual Machine - a simulation of a computer system.

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*

This document provides instructions for performing key operational tasks for the Dell EMC Ready Bundle for Red Hat OpenStack Platform. The information is arranged according to the specific operational area covered:

- *Updating BIOS and Firmware* on page 17
- *Adding and Removing Nodes* on page 35
- *Updating the Solution* on page 51

Intended Audience

This guide assumes the reader is familiar with:

- OpenStack
- Dell EMC PowerEdge R640 and Dell EMC PowerEdge R740xd RAID and BIOS configuration or the Dell EMC PowerEdge FX with the Dell EMC PowerEdge FC630 and Dell EMC PowerEdge FD332
- Red Hat Enterprise Linux (RHEL)
- Red Hat OpenStack Platform (RHOSP) documentation
- Network Configuration

Chapter 2

Updating BIOS and Firmware

Topics:

- *Overview*
- *BIOS and Firmware Update Procedures*

This chapter provides an overview and instructions for updating the BIOS and firmware on the following servers, as deployed in the Dell EMC Ready Bundle for Red Hat OpenStack Platform:

- Dell EMC PowerEdge R640
- Dell EMC PowerEdge R740xd

Overview

Keeping the Dell EMC Ready Bundle for Red Hat OpenStack Platform environment's BIOS and firmware versions current is key to ensure that the data center has the most recent security and ongoing operational fixes

Summary

Ensuring that your Dell EMC Ready Bundle for Red Hat OpenStack Platform environment's BIOS and firmware are updated to the latest revisions is of utmost importance, in order to maintain the highest security and peak performance. This guide describes the procedures required to update your cluster's BIOS and firmware.

CAUTION: The procedures in this guide must be completed with an active OpenStack deployment *prior to installing updated processors and memory*.

Prerequisites

Prerequisites for updating your Dell EMC Ready Bundle for Red Hat OpenStack Platform cluster's BIOS and firmware include:

- Network access to *http://www.dell.com/support*
- All nodes have access to their respective iDRACs

BIOS and Firmware Update Procedures

Perform the following procedures to update your Dell EMC Ready Bundle for Red Hat OpenStack Platform environment's BIOS and firmware versions:

1. *Download Files* on page 18
2. *Update the SAH Node's BIOS and Firmware* on page 19
3. *Update the Controller Nodes' BIOS and Firmware* on page 23
4. *Update the Ceph Storage Nodes' BIOS and Firmware* on page 27
5. *Update the Compute Nodes' BIOS and Firmware* on page 31

Download Files

Download the files required to update your Dell EMC Ready Bundle for Red Hat OpenStack Platform servers' BIOS and/or firmware.

1. Using a Web browser, navigate to *http://www.dell.com/support*.
2. Enter one of the systems' *service tag* in the **Enter a Service Tag or other product ID** text box, and then click on the **Submit** button.

   Note: You do not need to enter the service tag for every Dell EMC PowerEdge R640 or Dell EMC PowerEdge R740xd in the cluster, since they are all in the same server family and use the same firmware versions.
3. On the next screen, select the **Drivers & downloads** link.
4. On the same Web page, scroll down to **Systems Management**.
   a. Select the **down arrow** to expand the list, and then look for *Dell Server Update Utility, Windows 64 bit Format, v.x.x.x*.

      Note: The latest version will be displayed.
   b. Click on the **Download** button, and save the file to a local machine.
5. Copy this file to a system that has access to the iDRAC IP addresses of the systems that need to be updated.

Note: The iDRAC must be able to access this file and load them to perform the updates.

6. Proceed to *Update the SAH Node's BIOS and Firmware* on page 19.

Update the SAH Node's BIOS and Firmware

Update your SAH node's system BIOS and firmware.

1. Log into the SAH node's iDRAC using your **iDRAC username and password**.
2. On the right hand side in the *Virtual Console Preview* box, click on the **Launch** button.
   a) The dialogue window asks for a few Java verifications, and then takes you to the system's Virtual Console.
3. On the top bar of the *Virtual Console* window, click on the **Virtual Media** button.
4. Click on the **Connect Virtual Media** button.
5. Once connected, click on the **Virtual Media** button.
6. Click on the **Map CD/DVD…** button.
7. Browse to the **ISO file** downloaded in *Download Files* on page 18, and then click on the **Map Device** button.
8. On the top bar of the *Virtual Console* window, click on the **Next Boot** button.
9. Click on the **Lifecycle Controller** button.
10. Reboot the SAH node:

```
$ sudo reboot
```

11. On the *Lifecycle Controller* screen, click on the **Firmware Update** button.

Figure 1: Lifecycle Controller

12. Click on the **Launch Firmware Update** button.

13. On the next screen, select **Local Drive (CD or DVD or USB)** and then click on the **Next** button.

Figure 2: Select Update Repository

14. You are directed to a screen that should display the correct ISO image mapped.

Figure 3: Enter Access Details

15. Click on the **Next** button to display a screen that allows you to select updates for different components in the system.

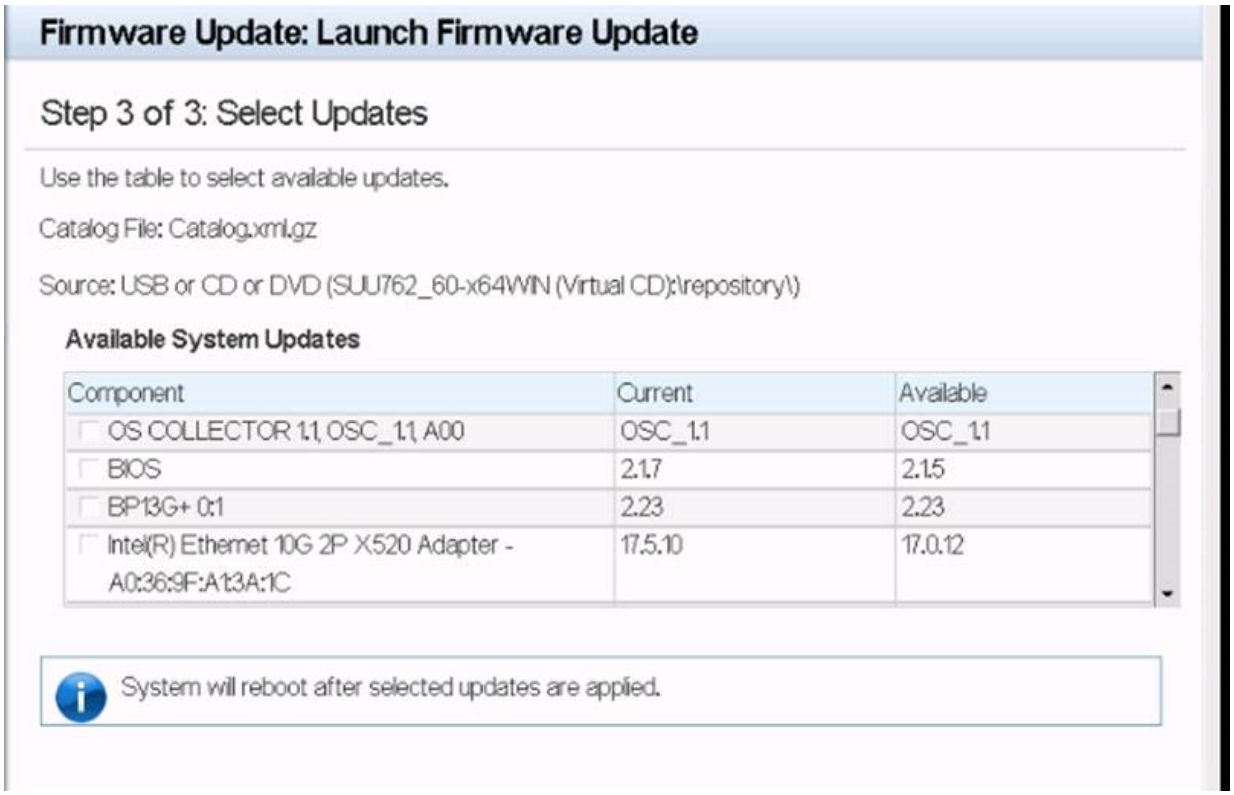

Figure 4: Select Updates

16. Click on the **Apply** button to launch the firmware update.
17. When complete, the system will reboot. You are returned to the Lifecycle Controller's *Firmware Update* screen.
18. Repeat step *11* on page 19 to step *16* on page 22 to ensure that everything was updated. Sometimes the NIC drivers update must be repeated.
    a) If that is the case, click on the **Apply** button again, and let the process go through its cycle.
19. Once everything has been updated successfully, disconnect the Virtual Media:
    a) On the top bar of the *Virtual Console* window, click on the **Virtual Media** button.
    b) Click on the **Disconnect Virtual Media** button.
20. After the SAH has been updated, log into the Director Node VM.
21. Ensure that everything is running properly:

```
$ sudo systemctl list-units "openstack*" "neutron*" "openvswitch*"
```

Note: It may take approximately 10 minutes for the `openstack-nova-compute` to become active after a reboot.

22. Verify the existence of your Overcloud and its nodes:

```
$ source ~/stackrc
$ openstack server list
$ openstack baremetal node list
$ openstack stack list
```

23. Proceed to *Update the Controller Nodes' BIOS and Firmware* on page 23.

Update the Controller Nodes' BIOS and Firmware

Update your Controller nodes' system BIOS and firmware.

CAUTION: You can only update the BIOS and firmware on *one Controller node at a time*, in order to ensure proper High Availability (HA) operation.

1. Log into the first Controller node's iDRAC using your **iDRAC username and password**.
2. On the right hand side in the *Virtual Console Preview* box, click on the **Launch** button.
   a) The dialogue window asks for a few Java verifications, and then takes you to the system's Virtual Console.
3. On the top bar of the *Virtual Console* window, click on the **Virtual Media** button.
4. Click on the **Connect Virtual Media** button.
5. Once connected, click on the **Virtual Media** button.
6. Click on the **Map CD/DVD…** button.
7. Browse to the **ISO file** downloaded in *Download Files* on page 18, and then click on the **Map Device** button.
8. On the top bar of the *Virtual Console* window, click on the **Next Boot** button.
9. Click on the **Lifecycle Controller** button.
10. Reboot the Controller node:

```
$ sudo reboot
```

11. On the *Lifecycle Controller* screen, click on the **Firmware Update** button.

Figure 5: Lifecycle Controller

12. Click on the **Launch Firmware Update** button.

13. On the next screen, select **Local Drive (CD or DVD or USB)** and then click on the **Next** button.

Figure 6: Select Update Repository

14. You are directed to a screen that should display the correct ISO image mapped.

Figure 7: Enter Access Details

15. Click on the **Next** button to display a screen that allows you to select updates for different components in the system.

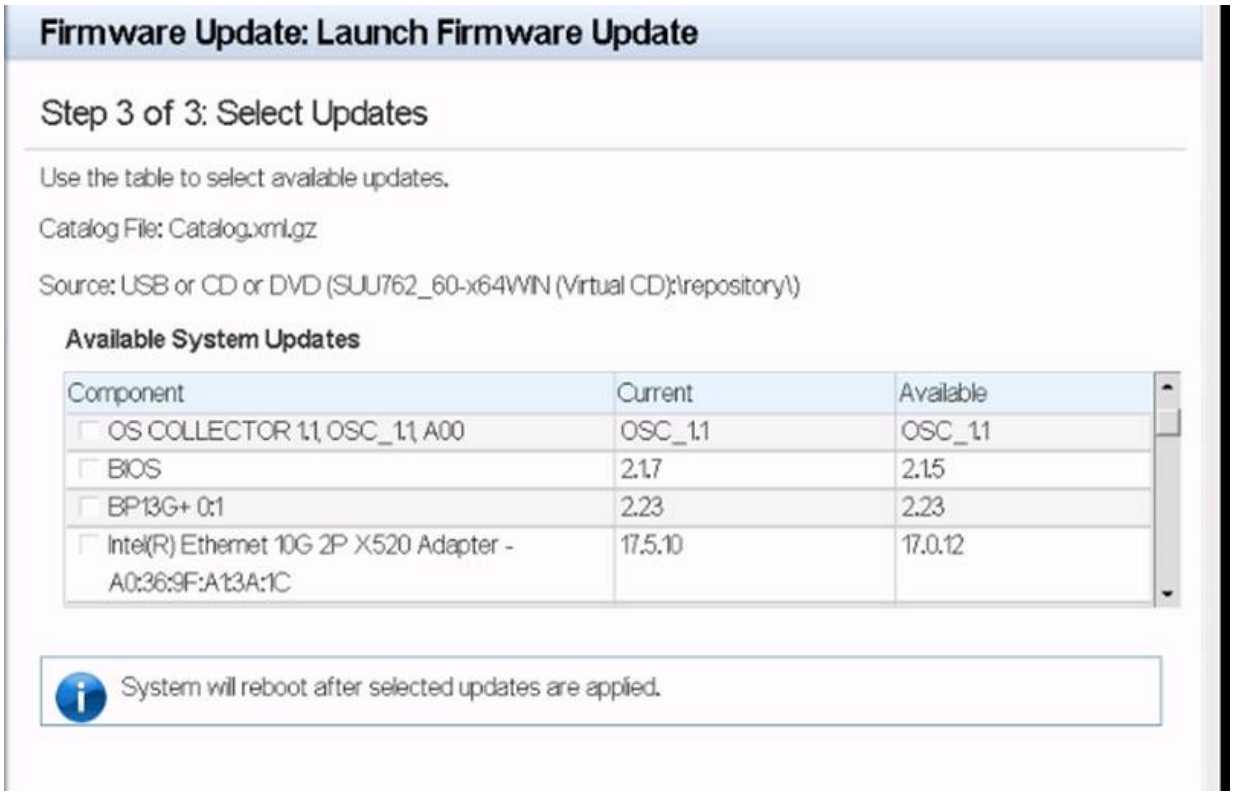

Figure 8: Select Updates

16. Click on the **Apply** button to launch the firmware update.
17. When complete, the system will reboot. You are returned to the Lifecycle Controller's *Firmware Update* screen.
18. Repeat step *11* on page 23 to step *16* on page 26 to ensure that everything was updated. Sometimes the NIC drivers update must be repeated.
    a) If that is the case, click on the **Apply** button again, and let the process go through its cycle.
19. Once everything has been updated successfully, disconnect the Virtual Media:
    a) On the top bar of the *Virtual Console* window, click on the **Virtual Media** button.
    b) Click on the **Disconnect Virtual Media** button.
20. Download the update files (see *Download Files* on page 18).
21. Reboot the node:

```
$ sudo reboot
```

22. After the Controller node has been updated, log into it.
23. Execute the following command:

```
$ sudo pcs status
```

    a) If any services fail after the reboot, execute the following command to clean the errors, and set the state of each resource to `Started`:

```
$ sudo pcs resource cleanup
```

    b) If any errors persist, contact Support to request guidance and assistance.

24. Check that all `systemd` services on the Controller node are active:

```
$ sudo systemctl list-units "openstack*" "neutron*" "openvswitch*"
```

Note: It may take approximately 10 minutes for the `openstack-nova-compute` to become active after a reboot.

25. Log out of the Controller node.

26. Repeat this procedure, sequentially, for the remainder of the Controller nodes.

27. Proceed to *Update the Ceph Storage Nodes' BIOS and Firmware* on page 27.

Update the Ceph Storage Nodes' BIOS and Firmware

Update your Ceph Storage nodes' system BIOS and firmware.

CAUTION: You can only update the BIOS and firmware on *one Ceph Storage node at a time*.

1. Log into the first Ceph Storage nodes' iDRAC using your **iDRAC username and password**.
2. On the right hand side in the *Virtual Console Preview* box, click on the **Launch** button.
   a) The dialogue window asks for a few Java verifications, and then takes you to the system's Virtual Console.
3. On the top bar of the *Virtual Console* window, click on the **Virtual Media** button.
4. Click on the **Connect Virtual Media** button.
5. Once connected, click on the **Virtual Media** button.
6. Click on the **Map CD/DVD…** button.
7. Browse to the **ISO file** downloaded in *Download Files* on page 18, and then click on the **Map Device** button.
8. On the top bar of the *Virtual Console* window, click on the **Next Boot** button.
9. Click on the **Lifecycle Controller** button.
10. Log into a Ceph Storage MON or Controller node
11. Download the update files (see *Download Files* on page 18).
12. Temporarily disable Ceph Storage cluster rebalancing:

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

13. Select the first Ceph Storage node to reboot, and then log into it.
14. Reboot the Ceph Storage node:

```
$ sudo reboot
```

15. On the *Lifecycle Controller* screen, click on the **Firmware Update** button.

Figure 9: Lifecycle Controller

16. Click on the **Launch Firmware Update** button.

17. On the next screen, select **Local Drive (CD or DVD or USB)** and then click on the **Next** button.

Figure 10: Select Update Repository

18. You are directed to a screen that should display the correct ISO image mapped.

Figure 11: Enter Access Details

19. Click on the **Next** button to display a screen that allows you to select updates for different components in the system.

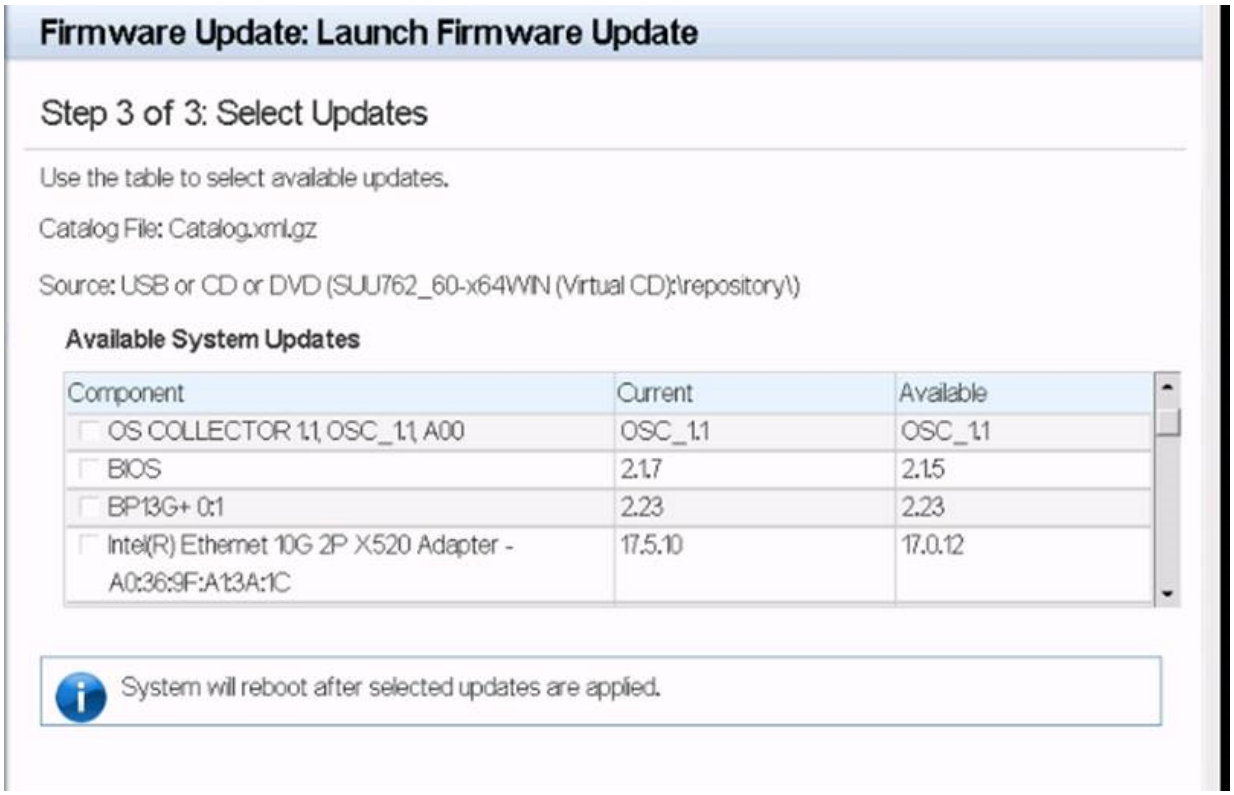

Figure 12: Select Updates

20. Click on the **Apply** button to launch the firmware update.
21. When complete, the system will reboot. You are returned to the Lifecycle Controller's *Firmware Update* screen.
22. Repeat step *15* on page 27 to step *20* on page 30 to ensure that everything was updated. Sometimes the NIC drivers update must be repeated.
    a) If that is the case, click on the **Apply** button again, and let the process go through its cycle.
23. Once the node has rebooted, log into it again.
24. Check the cluster status:

```
$ sudo ceph -s
```

25. Check that the `pgmap` reports all placement groups as normal `(active+clean)`.
26. Log out of the Ceph Storage node.
27. Repeat this procedure, sequentially, for the remainder of the Ceph Storage nodes.
28. When complete, log into a Ceph Storage MON or Controller node.
29. Re-enable cluster rebalancing:

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

30. Perform a final status check to verify the cluster reports `HEALTH_OK`:

```
$ sudo ceph status
```

31. Proceed to *Update the Compute Nodes' BIOS and Firmware* on page 31.

Update the Compute Nodes' BIOS and Firmware

Update your Compute nodes' system BIOS and firmware.

CAUTION: You can only update the BIOS and firmware on *one Compute node at a time*.

1. Log into the first Compute node's iDRAC using your **iDRAC username and password**.
2. On the right hand side in the *Virtual Console Preview* box, click on the **Launch** button.
   a) The dialogue window asks for a few Java verifications, and then takes you to the system's Virtual Console.
3. On the top bar of the *Virtual Console* window, click on the **Virtual Media** button.
4. Click on the **Connect Virtual Media** button.
5. Once connected, click on the **Virtual Media** button.
6. Click on the **Map CD/DVD…** button.
7. Browse to the **ISO file** downloaded in *Download Files* on page 18, and then click on the **Map Device** button.
8. On the top bar of the *Virtual Console* window, click on the **Next Boot** button.
9. Click on the **Lifecycle Controller** button.
10. List all Compute nodes and their UUIDs:

```
$ nova list | grep "compute"
```

11. Select a Compute node to reboot, and then log into it.
12. Download the update files:

```
$ source ~/overcloudrc
$ openstack compute service list
$ openstack compute service set [hostname] nova-compute --disable
```

13. List all instances on the Compute node:

```
$ openstack server list --host [hostname] --all-projects
```

14. Migrate each instance from the disabled host to another Compute node by using one of the following methods:
    a) Migrate the instance to a specific host of your choice:

```
$ openstack server migrate [instance-id] --live [target-host]--wait
```

    b) Let `nova-scheduler` automatically select the target host:

```
$ nova live-migration [instance-id]
```

15. Once the migration completes, confirm the instance has migrated from the Compute node:

```
openstack server list --host [hostname] --all-projects
```

16. Repeat step *15* on page 31 until you have migrated all instances from the Compute Node.
17. Reboot the Compute node:

```
$ sudo reboot
```

18. On the *Lifecycle Controller* screen, click on the **Firmware Update** button.

Figure 13: Lifecycle Controller

19. Click on the **Launch Firmware Update** button.

20. On the next screen, select **Local Drive (CD or DVD or USB)** and then click on the **Next** button.

Figure 14: Select Update Repository

21. You are directed to a screen that should display the correct ISO image mapped.

Figure 15: Enter Access Details

22. Click on the **Next** button to display a screen that allows you to select updates for different components in the system.

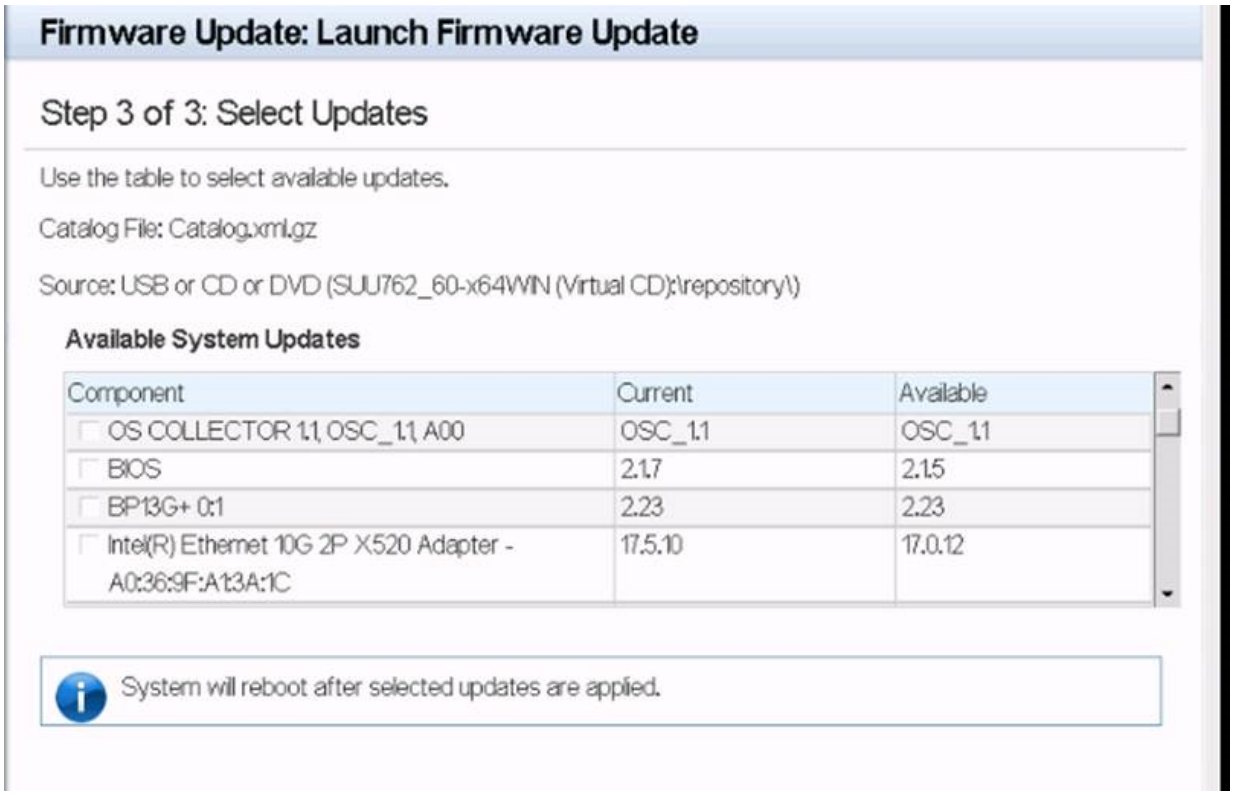

Figure 16: Select Updates

23. Click on the **Apply** button to launch the firmware update.

24. When complete, the system will reboot. You are returned to the Lifecycle Controller's *Firmware Update* screen.

25. Repeat step *18* on page 31 to step *23* on page 34 to ensure that everything was updated. Sometimes the NIC drivers update must be repeated.

a) If that is the case, click on the **Apply** button again, and let the process go through its cycle.

26. Once everything has been updated successfully, disconnect the Virtual Media:

a) On the top bar of the *Virtual Console* window, click on the **Virtual Media** button.

b) Click on the **Disconnect Virtual Media** button.

27. Once the Compute node has rebooted, re-enable it:

$ source ~/overcloudrc $ openstack compute service set [hostname] nova-compute --enable

28. Repeat this procedure, sequentially, for the remainder of the Compute nodes.

Chapter 3

Adding and Removing Nodes

Topics:

- *Overview*
- *Adding Nodes*
- *Removing Nodes*

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.

Overview

This chapter 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 R640
- Storage node:
  - Dell EMC PowerEdge R740xd

Note: The steps outilned in this document are to add a Dell EMC PowerEdge R640/Dell EMC PowerEdge R740xd server to an existing Dell EMC PowerEdge R640/Dell EMC PowerEdge R740xd deployment. The steps to add one of these servers to a deployment that consists of Dell EMC PowerEdge R630/Dell EMC PowerEdge R730xd needs to have had the initial deployment using the JetPack deployment 10.1 toolset. Previous versions of the toolset do not have support for Dell EMC PowerEdge R640/Dell EMC PowerEdge R740xd servers.

Compute nodes are used for hosting VMs in the Dell EMC Ready Bundle for Red Hat OpenStack Platform version 10.1.

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 10.1.

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

https://access.redhat.com/documentation/en-us/red_hat_openstack_platform/10/html/director_installation_and_usage/sect-scaling_the_overcloud

https://access.redhat.com/documentation/en-us/red_hat_ceph_storage/2/html/administration_guide/managing_cluster_size

Adding Nodes

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.

When adding a storage node to the cluster, consideration must be given to the Ceph Storage OSD configuration for the node being added.

If a 14G storage node with an HBA330 controller card is being added, then running `assign_role.py` will cause the OSD configuration to be automatically generated for the node. The generated configuration should be correct for most installations. The configuration is generated as follows:

- If the node has a mix of spinning drives and SSDs then the SSDs will be used as journaling devices and the spinning drives will be used as OSDs. The spinning drives/OSDs will be evenly distributed across the SSDs/journals.
- If the node has all spinning drives or all SSDs then the journals will be collocated with the OSDs, so no drives will be dedicated to journals.

The generated configuration may be customized if needed.

Adding a storage node of any other configuration requires the OSD configuration to be manually specified. If there are already storage nodes that do not have an HBA330 controller in the cluster then the OSD configuration for these nodes was already specified in the past and the OSD configuration for the node being added must match that existing configuration. The existing configuration may be found in `~/pilot/templates/dell-environment.yaml`. See the `ceph::profile::params::osds` attribute.

If adding the node to a cluster where all of the existing storage nodes are 14G with HBA330 controllers then the OSD configuration for the node being added must be specified and all future non-HBA330 storage nodes that are added later must have the same OSD configuration.

Prerequisites

The following prerequisites must be met:

- Dell EMC Ready Bundle for Red Hat OpenStack Platform version 10.1 installed
- Dell EMC Ready Bundle for Red Hat OpenStack Platform Software 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 Architecture Guide 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 Software 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 Software 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. Execute the following command to source the appropriate undercloud file:

```
$ source ~/stackrc
```

4. Run the following command specifying the iDRAC credentials (replace example below with actual values) and the IP address of the node being added:

Note: If the default password of the iDRAC is set on a new node but requires one that matches the rest of the nodes in the deployment, use the default password in this step. The password can be changed in a later step to match the rest of the nodes.

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

5. 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
```

6. 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": [
    {
      "model": "PowerEdge R640",
      "pm_addr": "192.168.110.122",
      "pm_password": "calvin",
      "pm_type": "pxe_drac",
      "pm_user": "root",
      "service_tag": "JR0BXK2"
    }
  ]
```

```
}
```

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.

7. If the new node being added is a Dell EMC PowerEdge R640/Dell EMC PowerEdge R740xd and the rest of the nodes in the deployment are the same models, proceed to step *10* on page 38.
8. If the new node being added is a Dell EMC PowerEdge R640/Dell EMC PowerEdge R740xd and the rest of the nodes in the deployment are Dell EMC PowerEdge R630/Dell EMC PowerEdge R730xd nodes, proceed to step *9* on page 38.

   Note: You must modify the NIC configuration in either the `compute.yaml` or `ceph-storage.yaml` file, since the NICs in Dell EMC PowerEdge R640/Dell EMC PowerEdge R740xd nodes are in different slots than in Dell EMC PowerEdge R630/Dell EMC PowerEdge R730xd nodes.
9. Navigate to the directory that contains the `compute.yaml` and `ceph-storage.yaml` files.

```
$ cd ~/pilot/templates/nic-configs
```

10. Identify the slot number of the network adaptor that is in a PCI slot.
    a) Edit the `compute.yaml` and/or `ceph-storage.yaml` file and modify the bonded NICs in the `ovs_bridge` section. For example, change *p3p1* on *bond0* and *p3p2* on *bond1* to *p2p1* for *bond0* and *p2p2* for *bond1*.
    b) Leave the *em1* and *em2* NICs as they are.
    c) Save the file, and proceed to step *11* on page 38.
11. 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.

    Note: If you want to change the root password on the iDRAC, use the -c option to set it to the new one.

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

For example:

```
$ ./config_idrac.py 192.168.110.122 -c <new_password> -n ~/newnodes.json
```

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

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

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

```
$ ./introspect_node.py <IP>
```

    a) For `<IP>`, substitute the IP address of the node's iDRAC.
14. Repeat the command in step *13* on page 38 for each new node.

    CAUTION: You must not proceed until introspection is complete on each node.
15. Each node must then be assigned a role within the cluster. Use the `assign_role.py` script to assign roles to each 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:

Note: When adding a new storage node, the size of the volume on which to install the OS must be set to 223 GB's by using the -o option.

```
$ ./assign_role.py  -n ~/newnodes.json 192.168.110.122 compute-3
or
$ ./assign_role.py  -n ~/newnodes.json 192.168.110.123 storage-3 -o 223
```

16. Repeat the command in step *15* on page 38 for each new node, using its iDRAC IP address and node type.
17. When adding a 14G storage node with an HBA330, if there is a need to change the generated OSD configuration for the node then edit `ceph-osd-config.yaml` and make any changes desired to the configuration for the new node.

Note: The generated OSD configuration for the node is identified by the system ID of the node. The system ID of the node may be obtained by running the following command:

```
$ probe-idrac -u root -p <password> -m DCIM_SystemView <idrac_ip>|grep
 UUID
```

CAUTION: When customizing the configuration for the node, be sure to only modify the configuration for the node being added.

18. When adding a storage node without an HBA330 controller, verify the storage configuration of the node matches the OSD configuration specified in the `ceph::profile::params::osds` attribute in `~/pilot/templates/dell-environment.yaml`. If no OSD configuration is specified there then add the OSD configuration for the node being added to the cluster following the instructions in the Dell EMC Ready Bundle for Red Hat OpenStack Platform Software Deployment Guide.
19. The boot images used during deployment must be set for the new nodes. To get the UUIDs of the nodes:

Note: Your node UUIDs will differ.

```
$ ironic node-list
```

The display returns a list similar to this:

```
+--------------------------------------+------
+--------------------------------------+-------------+--------------------
+-------------+
| 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-
a89f1f4b1ace | 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       |
+--------------------------------------+------
+--------------------------------------+-------------+-------------------
+-------------+
```

20. To get the UUIDs of the images:

```
$ openstack 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 |
| c136dbe1-2960-492c-a1eb-3f7048b08b5c | overcloud-full-vmlinuz | aki
    | aki              | 5153184    | active |
+--------------------------------------+-----------------------
+-------------+-----------------+------------+--------+
```

21. 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'
```

22. Repeat the commands in step *21* on page 40 for each new node.

23. The new nodes are ready to be deployed. Specify the new total number of Compute or Storage nodes to be deployed by modifying the original run time script:

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.**

24. Make a copy of the `overcloud_deploy_cmd.log` file, since the original one should be maintained for reference. This file contains the commands used for creating the initial deployment. Execute the following commands:

```
$ cd ~/pilot

$ cp overcloud_deploy_cmd.log overcloud_deploy_cmd.log.update

$ vi overcloud_deploy_cmd.log.update (use the editor of your choice)

Within this file, look for the line --compute-scale or
--ceph-storage-scale and increase the number to the new number
of deployed nodes, e.g., to go from 3 computes to 4, change
--compute-scale 3 to --compute-scale 4

If the original deployment was an Upgrade from a previous version of
OSP, e.g., OSP9, an additional environment line needs to be added to the
overcloud_deploy_cmd.log.update file.

$ vi overcloud_deploy_cmd.log.update

Within this file, add -e ~/pilot/templates/ceph-osd-config.yaml \
after the line -e ~/pilot/templates/network-environment.yaml \
```

25. Execute the following command to initiate the deployment:

```
$ /bin/bash overcloud_deploy_cmd.log.update
```

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

26. 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_rhscon.py <rhscon_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.

27. 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>
```

28. Update the `~/.ssh/config` file to allow SSH access to the new overcloud node:
    a) Execute the `update_ssh_config.py` script to add the new node to the deployment:

```
# ~/pilot/update_ssh_config.py
```

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

Removing Nodes

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.

Prerequisites

The following prerequisites must be met:

- For a Compute node, instances must be migrated from the Compute node being removed to another active Compute node.
- For a Storage cluster, there must be at least 3 Storage nodes AFTER the removal of one of the Storage nodes. This ensures that there is a quorum maintained.

Removing a 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 Software 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 Software Deployment Guide).
2. Source your `overcloudrc` file:

```
# source ~/overcloudrc
```

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

```
# nova service-list
```

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

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

5. Source the Undercloud `stackrc` file:

```
# source ~/stackrc
```

6. Identify the ID of the Overcloud stack:

```
# openstack stack list
```

Example truncated output:

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

7. 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-compute-0     |  ...
| 52f28960-1a91-4955-847b-9969b84067d2 | overcloud-compute-1     |  ...
| eef7edf3-ae98-4921-88ff-7fc296e8bef7 | overcloud-compute-2     |  ...
| cb68faf8-ba6c-405d-81ab-c8582138f536 | overcloud-compute-3     |  ...
| 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. From the output above, `overcloud-compute-3`'s ID is `cb68faf8-ba6c-405d-81ab-c8582138f536`, and will be used in step *11* on page 43.
9. Prior to removing the node, identify the Ironic ID associated with the nova ID of the Compute node to delete, by executing the following command. This will be used in step *17* on page 44:

   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:

   ```
   d031095e-a673-4532-849a-9c5d263944ed
   ```

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

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

11. 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 cb68faf8-ba6c-405d-81ab-
    c8582138f536
    ```

    Example output:

    ```
    Deleting the following nodes from stack R11b:
    - cb68faf8-ba6c-405d-81ab-c8582138f536
    Started Mistral Workflow. Execution ID: 04221bcf-542a-48c4-bf76-
    f35808b139a6
    ```

12. Execute the `nova list` command repeatedly until you no longer see the Compute node in the list.

13. Source your `overcloudrc` file:

```
# source ~/overcloudrc
```

14. Identify the ID of the removed Compute node:

```
# nova service-list
```

15. Remove the Compute node from the OpenStack service:

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

16. Remove the Compute nodes Open vSwitch agent:

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

17. Using the output gathered from step *9* on page 43, 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

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 Software 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* on page 44:

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

8. 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) Exit the utility by pressing *[Ctrl-C]*.

d) 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.

e) 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
```

f) 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
```

g) 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
```

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

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

i) Repeat from step *7.a* on page 45 until all OSDs on the Storage node have been removed from the cluster.

9. 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>
```

10. 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 Software 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  |  ...
+--------------------------------------+------------+- ...
```

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

12. From the output above, `overcloud-cephstorage-3`'s ID is `99c8def0-35ce-4162-885b-23f85426f41c`, and will be used in step *15* on page 47.
13. Prior to removing the node, identify the Ironic ID associated with the nova ID of the Storage node to delete, by executing the following command. This will be used in step *18* on page 48:

    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:

    ```
    b330c67c-505f-4e97-a145-2861edb063a7
    ```

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

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

15. 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:
 - 99c8def0-35ce-4162-885b-23f85426f41c
 Started Mistral Workflow. Execution ID: 04221bcf-542a-48c4-bf76-
f35808b139a6
```

16. Execute the `nova list` command repeatedly until you no longer see the Storage node in the list.
17. 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  28  2017 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
```

18. Using the output gathered from step *13* on page 47, delete the Storage node from Ironic by executing the following command:

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

19. Initiate a `ssh` session to any existing Storage node.
20. 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
```

21. 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.

Chapter 4

Updating the Solution

Topics:

- *Overview*
- *Dell EMC Ready Bundle for Red Hat OpenStack Platform v10.0.1 Minor Undercloud/Overcloud Update*

The ability to keep the Red Hat OpenStack Platform up to date with the latest patches will arise from time to time. In this guide, this scenario will be described for manual updates. To facilitate fewer touch points, Dell EMC has created a set of tools to support the process and automate many of the steps.

Overview

The *Dell EMC Ready Bundle for Red Hat OpenStack Platform v10.0.1 Minor Undercloud/Overcloud Update* on page 52 raises the solution:

- **From** — the Dell EMC Ready Bundle for Red Hat OpenStack Platform v10.0.1 on RHOSP 10
- **To** — the Dell EMC Ready Bundle for Red Hat OpenStack Platform v10.1 on RHOSP 10, Minor Update

Dependencies

Dell EMC Ready Bundle for Red Hat OpenStack Platform update/upgrade dependencies and prerequisites include:

- *Archive Files* on page 52
- *Red Hat Subscriptions* on page 52

Archive Files

For customers performing a self-installation, these files are available on request from Dell EMC:

- `Update-upgrade.tgz` — update/upgrade scripts archive

Please contact your account representative, or email *openstack@dell.com* for instructions.

Red Hat Subscriptions

Before starting an update, you must have purchased enough subscriptions for all of the nodes in the Overcloud. Each node requires all Red Hat Ceph Storage and OpenStack subscriptions.

Dell EMC Ready Bundle for Red Hat OpenStack Platform v10.0.1 Minor Undercloud/Overcloud Update

This topic describes how to perform a minor RHOSP Undercloud/Overcloud update:

- **From** — the Dell EMC Ready Bundle for Red Hat OpenStack Platform v10.0.1 on RHOSP 10
- **To** — the Dell EMC Ready Bundle for Red Hat OpenStack Platform v10.1 on RHOSP 10, Minor Update

Minor Update Procedure

At some point after deploying the Dell EMC Ready Bundle for Red Hat OpenStack Platform version 10.0.1 using locked package versions, a software update will be required. The following procedure enables you to perform a **minor update**.

- A minor update is when the OpenStack version **stays the same** (i.e., updating within Newton)

Prerequistes

Prerequisites for updating your Dell EMC Ready Bundle for Red Hat OpenStack Platform include:

- *Enable Red Hat Storage Console Imports* on page 52

Enable Red Hat Storage Console Imports

You must perform the following procedure **before** you begin the the update, in order to successfully import clusters via the Red Hat Storage Console after the update is complete .

To enable Red Hat Storage Console imports:

1. Access the Red Hat Storage Console UI at <*public_api_ip*>/*skyring*

   Note: You can find the Red Hat Storage Console (`rhscon`) *public_api_ip* in your *.properties* file on your SAH node.

2. Log in with the following credentials:
   - **Username** — admin
   - **Password** — admin
3. Click on the **Import Cluster** button.
4. Select **controller-0** from the list (this is the monitor host.)
   a) Click on the **Continue** button.

   The summary page appears. It will show the Controller nodes as monitors, and the Storage nodes as OSDs.
5. Click on the **Import** button.

   The import operation will take several minutes to complete.

When complete, you will see a green check mark at the left of the screen. You can now monitor the storage from this UI.

Caveats and Impacts for Minor Update

This topic discusses caveats and impacts for the minor update procedure.

Minor Update Caveats

Caveats include:

- You must have installed the Dell EMC Ready Bundle for Red Hat OpenStack Platform v10.0.1 on RHOSP 10 with Lock deployment files. See *Updating the Environment*.

CAUTION: If you have established another cluster (Horizon) password since your original Dell EMC Ready Bundle for Red Hat OpenStack Platform Version 10.0.1 installation, the update process may reset the password in the database to the **original**, auto-generated password. You must log into Horizon with this original password to reset it to your new password.

Minor Update Impacts

Impacts include:

- **Service downtime** — The update will involve some service interruptions. Contact Red Hat Support for more information.
- **Total system downtime** — Reboots may be required for kernel updates. Contact Red Hat Support for more information.
- **Data loss** — None.
- **Backup, rollback, and/or recovery procedures** — See *Important Pre-Upgrade Notes*.

Deploy the Update/Upgrade Tools

To deploy the update/upgrade tools:

1. Upload the `Update-upgrade.tgz` file to the OpenStack Administrator's home directory on the Director Node.
2. Log into the Director Node as the *stack* user, and then extract the tar file contents into the *home* directory:

```
$ cd /home/stack
$ tar zxvf ~/Update-upgrade.tgz
```

Remove Instance High Availability

If Instance HA is installed on your Dell EMC Ready Bundle for Red Hat OpenStack Platform cluster, you must uninstall it before proceeding with the update. Otherwise, proceed to *Run the Update Script* on page 55.

To check whether or not Instance HA is installed:

1. Log into one of the Controller nodes as the *root* user.
2. Execute the following command:

```
# pcs status
```

3. Look for any of the following resources in the output:

```
ipmilan-<compute_node_name>-N  (stonith:fence_ipmilan):
fence-nova     (stonith:fence_compute):
<compute_node_name>-N (ocf::pacemaker:remote):
```

To uninstall Instance HA:

1. Log into the Red Hat OpenStack Platform Director Node as the RHOSP *admin* user.
2. Source the `stackrc` file:

```
$ source ~/stackrc
```

3. Change into the *~/update_upgrade* directory:

```
$ cd ~/update_upgrade
```

4. Execute the following command:

```
$ ./uninstall-instanceHA-10.py | --logging-level <logging_level>
```

   Optional parameters include:

   - `--logging-level` — controls the logging level output. It can be ignored unless you require detailed logging for debug efforts. Logging levels conform to standard `syslog` levels:

Table 1: Logging Levels

Logging Level	Value	Description
EMERG	0	System is unusable
ALERT	1	Action must be taken immediately
CRIT	2	Critical conditions
ERR	3	Error conditions
WARNING	4	Warning conditions
NOTICE	5	Normal but significant condition
INFO	6	Informational
DEBUG	7	Debug-level messages

Note: The `uninstall-instanceHA-10.py` script may take many minutes to execute, as it waits for the `pacemaker_remote` service to shut down on each Compute node.

5. All `nova` services should be in the *up* state. To check their states:

   a. Initiate an SSH session to one the Controller nodes as the *root* user.

b. Execute the following command:

```
# nova service-list
```

6. If some services are in the *down* state:

   a. Execute the following command:

   ```
   # systemctl restart <service>
   ```

   For example:

   ```
   # systemctl restart openstack-nova-api
            # systemctl restart openstack-nova-conductor
            # systemctl restart openstack-nova-consoleauth
            # systemctl restart openstack-nova-novncproxy
            # systemctl restart openstack-nova-schedule
   ```

7. Re-check the states of all `nova` services:

   ```
   # nova service-list
   ```

The Instance HA cluster configuration is now removed. Proceed to *Run the Update Script* on page 55.

Run the Update Script

To run the update script:

1. Ensure that the cluster is running cleanly.
2. Log into the Director Node as the *stack* user, and then execute the `update-js10.0.1.sh` script from the *home* directory:

   ```
   $ ~/update_upgrade/update-js10.0.1.sh <openstack_pool_id> <ceph_pool_id>
    <rh_subscription_id> <rh_subscription_pw> [ <stack_name> ]
   ```

   Required arguments include:

   - `openstack_pool_id` — The OpenStack pool ID. Requires a Red Hat OpenStack Platform subscription.
   - `ceph_pool_id` — The Ceph Storage pool ID. Requires a Red Hat Ceph Storage subscription.
   - `rh_subscription_id` — The CDN username.
   - `rh_subscription_pw` — The CDN user password.

   Optional argument includes:

   - `stack_name` — The Overcloud stack name. Defaults to *overcloud*.

Reboot the Director Node

After the first phase of updating the Director packages, the script will prompt to reboot the Director Node (by selecting *[Enter]*), to incorporate a new OS kernel and software. When it is finished rebooting:

1. Log into the Director Node as the *stack* user.
2. Execute the `update-js10.0.1.sh` script from the *home* directory, with the exact arguments as in *Run the Update Script* on page 55:

   ```
   $ ~/update_upgrade/update-js10.0.1.sh <openstack_pool_id> <ceph_pool_id>
    <rh_subscription_id> <rh_subscription_pw> [ <stack_name> ]
   ```

3. When prompted, reboot the Director Node by following the procedure outlined in *https://access.redhat.com/documentation/en-us/red_hat_openstack_platform/10/html/director_installation_and_usage/sect-rebooting_the_overcloud#sect-Rebooting-Director*

The update process will proceed to the next update phase.

Logging and Lock Files

The logging for the update process will be in `~/update-upgrade/update-js10.0.1.log`. As each phase of the update process completes, a lock file will be put in `~/update-upgrade/update-lockfiles`:

- `director-updated.lock`
- `overcloud-images-updated.lock`
- `overcloud-registered.lock`
- `overcloud-prepared.lock`
- `overcloud-updated.lock`

If an update phase fails, the lock file will not be created. When the error condition has been fixed, the script can be restarted (with the same arguments) and the process will continue, starting at the failed phase and continuing on from there.

Finalize the Update

To verify that the cluster's status is clean:

1. Once the update is complete, Dell EMC recommends that you perform a full reboot of all nodes. Follow the procedures outlined in *https://access.redhat.com/documentation/en-us/red_hat_openstack_platform/10/html/director_installation_and_usage/sect-rebooting_the_overcloud*. Virtual machines on the cluster **must be restarted after Compute nodes have been rebooted.**
2. Log into a Controller node as the *heat-admin* user:

```
$ ssh cntl0
```

3. Execute the following command:

```
$ sudo pcs status
```

The result should show that all resources have started.

4. Check `/var/log/messages` for errors.
5. Fix any error conditions that may be found.
6. If errors were found and fixed, execute the `update-js10.0.1.sh` script from the *home* directory, with the exact arguments as in *Run the Update Script* on page 55:

```
$ ~/update-js10.0.1.sh <openstack_pool_id> <ceph_pool_id>
 <rh_subscription_id> <rh_subscription_pw> [ <stack_name> ]
```

Install Instance High Availability

If Instance HA was installed on your Dell EMC Ready Bundle for Red Hat OpenStack Platform cluster prior to beginning the update procedure, you can reinstall it.

To reinstall Instance HA:

1. Log into the Red Hat OpenStack Platform Director Node as the RHOSP *admin* user.
2. Source the `stackrc` file:

```
$ source ~/stackrc
```

3. Change into the *update_upgrade* directory:

```
$ cd ~/update_upgrade
```

4. Execute the following command:

```
$ ./install-instanceHA.py | --logging-level <logging_level>
```

   Optional parameters include:

   - `--logging-level` — controls the logging level output. It can be ignored unless you require detailed logging for debug efforts.

   Note: The `install-instanceHA.py` script may take many minutes to execute, as it waits for the `pacemaker_remote` service to shut down on each Compute node.

5. Reboot the Director Node by following the procedure outlined in *https://access.redhat.com/documentation/en-us/red_hat_openstack_platform/10/html/director_installation_and_usage/sect-rebooting_the_overcloud#sect-Rebooting-Director*.

The Instance HA cluster configuration is now reinstalled.

Appendix A

References

Topics:

- *To Learn More*

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://dell.com/openstack*.

This document and all other related architecture and technical guides can be found in the Dell EMC TechCenter community at *http://en.community.dell.com/techcenter/cloud/w/wiki/12047.dell-emc-red-hat-openstack-cloud-solutions*