

Adding a Storage Node to a Ceph Cluster in the Dell Red Hat OpenStack Cloud Solution - Version 4.0



Contents

Trademarks.....	3
Notes, Cautions, and Warnings.....	4
Executive Summary.....	5
Intended Audience.....	5
Adding a PowerEdge R730xd Storage Node to a Ceph Cluster.....	6
Prerequisites.....	6
Node Definitions.....	6
Adding the Node.....	7
Getting Help.....	11
Contacting Dell.....	11
References.....	11
To Learn More.....	11

Trademarks

© 2014-2015 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is prohibited. For more information, contact Dell.

Trademarks used in this text: Dell™, the DELL logo, Dell Precision™, OptiPlex™, Latitude™, PowerEdge™, PowerVault™, OpenManage™, EqualLogic™, Dell Compellent™, KACE™, FlexAddress, Dell Networking™, and Vostro™ are trademarks of Dell Inc. Intel®, Pentium®, Xeon®, Core®, and Celeron® are registered trademarks of Intel Corporation in the U.S. and other countries. AMD® is a registered trademark and AMD Opteron™, AMD Phenom™, and AMD Sempron™ are trademarks of Advanced Micro Devices, Inc. Microsoft®, Windows®, Windows Server®, MS-DOS®, and Windows Vista® are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries.

Red Hat®, Red Hat® Enterprise Linux®, the Shadowman logo, and JBoss are 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. Novell® is a registered trademark and SUSE™ is a trademark of Novell Inc. in the United States and other countries. Oracle® is a registered trademark of Oracle Corporation and/or its affiliates. Citrix®, Xen®, XenServer® and XenMotion® are either registered trademarks or trademarks of Citrix Systems, Inc. in the United States and/or other countries. VMware®, Virtual SMP®, vMotion®, vCenter®, and vSphere® are registered trademarks or trademarks of VMware, Inc. in the United States or other countries.

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.

Other trademarks and trade names may be used in this publication to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

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.

Executive Summary

One of the common needs of a cloud platform is storage expansion for performance or capacity improvements.

This technical guide explains the process of adding a Dell™ PowerEdge™ R730xd storage node to an existing Ceph cluster in the Dell Red Hat® OpenStack Cloud Solution.

Red Hat Ceph Storage version 1.3 is used for block, image, and ephemeral storage in the Dell Red Hat OpenStack Cloud Solution version 4.0.

Intended Audience

This technical guide is written for OpenStack administrators or deployment engineers who are responsible for installation and ongoing operation of OpenStack clusters.

Adding a PowerEdge R730xd Storage Node to a Ceph Cluster

This section describes prerequisites and procedures to add a storage node to a Dell™ Red Hat® Cloud Solutions Red Hat Ceph Storage cluster. Topics discussed include:

- [Prerequisites](#) on page 6
- [Node Definitions](#) on page 6
- [Adding the Node](#) on page 7

Prerequisites

The following prerequisites must be met:

- Dell Red Hat OpenStack Cloud Solution version 4.0 installed
- *Dell Red Hat OpenStack Cloud Solution Deployment Guide* available for reference
- The new Storage node has been wired to the network in the same manner as the other Storage nodes, per the Reference Architecture
- The Dell DTK utility has been used to configure RAID/BIOS on the new Storage node
- `hammer` deployment scripts available on the OpenStack Foreman Installer Node (OFI Node)

Node Definitions

[Table 1: Example Node Definitions](#) on page 6 presents a listing of servers as used in this white paper.

 **Note:** These definitions are examples only. Replace as needed for your environment.

Table 1: Example Node Definitions

Function	IP Address	FQDN	Short Name	Notes
Ceph VM (Red Hat Ceph Storage Admin Node)	192.168.170.43	ceph.13g.rcbd.lab	ceph	Existing
Controller Node 1	192.168.170.44	cntl1.13g.rcbd.lab	cntl1	Existing
Controller Node 2	192.168.170.45	cntl2.13g.rcbd.lab	cntl2	Existing
Controller Node 3	192.168.170.46	cntl3.13g.rcbd.lab	cntl3	Existing
Compute Node 1	192.168.170.47	nova1.13g.rcbd.lab	nova1	Existing
Compute Node 2	192.168.170.48	nova2.13g.rcbd.lab	nova2	Existing
Compute Node 3	192.168.170.49	nova3.13g.rcbd.lab	nova3	Existing
Storage Node	192.168.170.50	ss1.13g.rcbd.lab	ss1	Existing
Storage Node	192.168.170.51	ss2.13g.rcbd.lab	ss2	Existing
Storage Node	192.168.170.52	ss3.13g.rcbd.lab	ss3	Existing
Storage Node	192.168.170.53	ss4.13g.rcbd.lab	ss4	New

The following code is an example listing of servers prior to running the `hammer host list` command, which is used in [Adding the Node](#) on page 7.

 **Note:** The ID in the first column will be used during configuration.

```
hammer host list

[root@fore pilot]# hammer host list

ID | NAME | OPERATING SYSTEM | HOST GROUP | IP
   | MAC
---|-----|-----|-----|---
 2 | cnt11.13g.rcbd.lab | RHEL Server 7.1 HA All In One Controller
   | 192.168.120.44 | ec:f4:bb:c7:97:b4
 3 | cnt12.13g.rcbd.lab | RHEL Server 7.1 HA All In One Controller
   | 192.168.120.45 | ec:f4:bb:c7:92:1c
 4 | cnt13.13g.rcbd.lab | RHEL Server 7.1 HA All In One Controller
   | 192.168.120.46 | ec:f4:bb:c7:93:4c
 1 | fore.13g.rcbd.lab | RHEL Server 7.1 10.148.44.42 52:54:00:85:09:7e
 5 | nova1.13g.rcbd.lab | RHEL Server 7.1 Compute (Neutron)192.168.170.47
   | ec:f4:bb:c7:90:78
 6 | nova2.13g.rcbd.lab | RHEL Server 7.1 Compute (Neutron)192.168.120.48
   | ec:f4:bb:c7:93:b4
 7 | nova3.13g.rcbd.lab | RHEL Server 7.1 Compute (Neutron)192.168.120.49
   | ec:f4:bb:c7:93:7c
 8 | ss1.13g.rcbd.lab | RHEL Server 7.1 192.168.120.50 | ec:f4:bb:c7:96:d4
 9 | ss2.13g.rcbd.lab | RHEL Server 7.1 192.168.120.51 | ec:f4:bb:c7:96:6c
10 | ss3.13g.rcbd.lab | RHEL Server 7.1 192.168.120.52 | ec:f4:bb:c7:92:a4
```

Adding the Node

To add a storage node to a Dell Red Hat OpenStack Cloud Solution Ceph cluster:

1. Log onto the OFI Node as the `root` user.
2. Navigate to the directory in which the hammer deployment scripts reside. The Dell Red Hat OpenStack Cloud Solution places those scripts into the `/root/pilot` directory.
3. Ensure that the `osp_config.sh` file has been previously configured to the existing installation. Bonds, partition tables, etc., must be pre-defined.
4. Execute the following command:

```
./hammer-deploy-storage.sh <new_hostname> <PXE_interface_MAC_address>
<IP_address_for_provisioning> <server_model>
```

5. Execute the following command to return the `host_id` for the new node.

 **Note:** In our examples, the new `host_id` is **11**. The `host_id` for the new node in your environment may be different.

```
hammer host list
```

Example output:

```
[root@fore pilot]# hammer host list

ID | NAME | OPERATING SYSTEM | HOST GROUP | IP
   | MAC
---|-----|-----|-----|---
 2 | cnt11.13g.rcbd.lab | RHEL Server 7.1 HA All In One Controller
   | 192.168.120.44 | ec:f4:bb:c7:97:b4
```

```

3 | cnt12.13g.rcbd.lab | RHEL Server 7.1 HA All In One Controller
192.168.120.45 | ec:f4:bb:c7:92:1c
4 | cnt13.13g.rcbd.lab | RHEL Server 7.1 HA All In One Controller
192.168.120.46 | ec:f4:bb:c7:93:4c
1 | fore.13g.rcbd.lab | RHEL Server 7.1 10.148.44.42 52:54:00:85:09:7e
5 | nova1.13g.rcbd.lab | RHEL Server 7.1 Compute (Neutron)192.168.170.47
ec:f4:bb:c7:90:78
6 | nova2.13g.rcbd.lab | RHEL Server 7.1 Compute (Neutron)192.168.120.48
ec:f4:bb:c7:93:b4
7 | nova3.13g.rcbd.lab | RHEL Server 7.1 Compute (Neutron)192.168.120.49
ec:f4:bb:c7:93:7c
8 | ss1.13g.rcbd.lab | RHEL Server 7.1 192.168.120.50 |
ec:f4:bb:c7:96:d4
9 | ss2.13g.rcbd.lab | RHEL Server 7.1 192.168.120.51 |
ec:f4:bb:c7:96:6c
10 | ss3.13g.rcbd.lab | RHEL Server 7.1 192.168.120.52 |
ec:f4:bb:c7:92:a4
11 | ss4.13g.rcbd.lab | RHEL Server 7.1 192.168.120.53 |
ec:f4:bb:ce:d2:ec

```

6. Make the RPM version locking file available during provisioning by executing the following command:

```
hammer host set-parameter --host-id <host_ID_from_hammer_host_list> --name <yum_versionlock_file> -value \ 'http://<IP_address_of_OFI_node>/ceph.vlock'
```

For example:

```
hammer host set-parameter --host-id 11 --name yum_versionlock_file --value 'http://192.168.120.42/ceph.vlock'
```

7. PXE boot the new node by powering on the Storage node and selecting *F12 - PXE Boot*.
- a. Ensure that the new node appears in the OFI UI.
8. Ensure that all settings are configured correctly on the newly-provisioned R730xd. These values will be set by the kickstart scripts that are a part of the solution. For example:
- Networking - Ensure the configured networks are communicating properly:
 - Ping the Red Hat Ceph Storage Admin Node (Ceph VM) on the 192.167.170.x network
 - Ping one of the other cluster Storage nodes on the 192.168.180.x network
 - SELinux=permissive
 - Firewalld is off and disabled
 - NetworkManager is off and disabled
 - NTPD is running
9. Log onto the Red Hat Ceph Storage Admin Node (Ceph VM) as the *root* user.
10. Open the */etc/hosts* file for editing:

```
vi /etc/hosts
```

11. Edit the file to include the new host. For example:

```

192.168.170.43 ceph.13g.rcbd.lab ceph
192.168.170.44 cnt11.13g.rcbd.lab cnt11
192.168.170.45 cnt12.13g.rcbd.lab cnt12
192.168.170.46 cnt13.13g.rcbd.lab cnt13
192.168.170.47 nova1.13g.rcbd.lab nova1
192.168.170.48 nova2.13g.rcbd.lab nova2
192.168.170.49 nova3.13g.rcbd.lab nova3
192.168.170.50 ss1.13g.rcbd.lab ss1
192.168.170.51 ss2.13g.rcbd.lab ss2

```

```
192.168.170.52 ss3.13g.rcbd.lab ss3
192.168.170.53 ss4.13g.rcbd.lab ss4 (New node.)
```

12. Save the file, then execute the following commands:

```
ssh-copy-id ss4
ssh ss4 'echo -e "<ip_address_of_calamari_node>
<hostname_of_calamari_node> ceph" >> /etc/hosts'; done
ssh root@ss4 'useradd -g adm -m ceph-user';
ssh root@ss4 'passwd ceph-user';
ssh root@ss4 'echo -e "ceph-user ALL = (root) NOPASSWD:ALL" > /etc/
sudoers.d/ceph-user';
ssh root@ss4 'echo -e "Defaults:ceph-user !requiretty" >> /etc/sudoers.d/
ceph-user';
ssh root@ss4 'chmod 0440 /etc/sudoers.d/ceph-user';
```

13. Switch to the *Ceph* user by executing the following command:

```
su - ceph-user
```

14. Execute the following commands:

```
ssh-copy-id ss4
ceph-deploy install ss4
cd cluster
```

a. Then perform **one** of the following two steps:

- Execute the `ceph-deploy disk list ss4` command, **or**
- Log into the new Storage node as the *root* user, and execute the `lsblk` command.

These commands will display the list of drives that are available to be used for OSDs, journals, and which one is used for the OS.

Example output on *ss4*:

NAME	MAJ:MIN	RM	SIZE	RO	TYPE	MOUNTPOINT
sda	8:0	0	186.3G	0	disk	
sdb	8:16	0	186.3G	0	disk	
sdc	8:32	0	186.3G	0	disk	
sdd	8:48	0	3.7T	0	disk	
sde	8:64	0	3.7T	0	disk	
sdf	8:80	0	3.7T	0	disk	
sdg	8:96	0	3.7T	0	disk	
sdh	8:112	0	3.7T	0	disk	
sdi	8:128	0	3.7T	0	disk	
sdj	8:144	0	3.7T	0	disk	
sdk	8:160	0	3.7T	0	disk	
sdl	8:176	0	3.7T	0	disk	
sdm	8:192	0	3.7T	0	disk	
sdn	8:208	0	3.7T	0	disk	
sdo	8:224	0	3.7T	0	disk	
sdp	8:240	0	3.7T	0	disk	
sdq	65:0	0	278.9G	0	disk	
##sdq1	65:1	0	2M	0	part	
##sdq2	65:2	0	1G	0	part	/boot
##sdq3	65:3	0	277.9G	0	part	
##VolGroup-lv_root	253:0	0	231.8G	0	lvm	/
##VolGroup-lv_swap	253:1	0	16G	0	lvm	[SWAP]
##VolGroup-lv_var	253:2	0	20G	0	lvm	/var
##VolGroup-lv_tmp	253:3	0	10G	0	lvm	/tmp

This output shows that:

- *sda*, *sdb*, and *sdc* will be used for the journals since they are smaller drives

- *sdd* thru *sdp* will be the OSD drives
- *sdq* has the Red Hat OS installed

15. Execute the following commands to ensure any existing data that might be on the drives are wiped clean:



Note: Ensure that the drive allocated for the Red Hat OS is **not** included in the list below!

```
for drv in a b c d e f g h i j k l m n o p ; do ceph-deploy disk zap\
ss4:/dev/sd$drv; done
```

16. Execute the following commands to create the OSDs, and assign a set of drives to each journal drive:

```
for drv in d e f g h; do ceph-deploy --overwrite-conf osd create\
ss4:/dev/sd$drv:/dev/sda; done
for drv in i j k l; do ceph-deploy --overwrite-conf osd create\
ss4:/dev/sd$drv:/dev/sdb; done
for drv in m n o p; do ceph-deploy --overwrite-conf osd create\
ss4:/dev/sd$drv:/dev/sdc; done
```

17. Attach the new storage node to Calamari:

```
ceph-deploy calamari connect --master
<FQDN_of_the_Red_Hat_Ceph_Storage_Admin_Node> ss4
```

18. Log into the Calamari web interface.

19. Add the new *ss4* to the Calamari group.

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

Getting Help

This appendix details contact and reference information for the Dell™ Red Hat® Cloud Solutions with Red Hat Enterprise Linux® OpenStack Platform.

Contacting Dell

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 product catalog.

Dell 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 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 that is convenient for you.

References

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

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

To Learn More

For more information on the Dell Red Hat OpenStack Cloud Solution with Red Hat Enterprise Linux™ OpenStack Platform visit <http://www.dell.com/openstack>.

© 2014-2015 Dell Inc. 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 and its affiliates cannot be responsible for errors or omissions in typography or photography. Dell's Terms and Conditions of Sales and Service apply and are available on request. Dell service offerings do not affect consumer's statutory rights.

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