

Dell EMC Ready Bundle for Red Hat OpenStack NFV Platform

PowerEdge R-Series Hardware Deployment Guide

Version 10.0.1



Dell EMC Converged Platforms and Solutions

Contents

List of Tables.....	iv
Trademarks.....	5
Notes, Cautions, and Warnings.....	6
Chapter 1: Overview.....	7
Deployment Methodology.....	8
Intended Audience.....	8
Dependencies.....	8
Chapter 2: Hardware Setup.....	9
Unpacking and Installing the Equipment.....	10
Powering Up the Equipment.....	10
Verifying the Equipment.....	10
Tested BIOS and Firmware.....	10
Chapter 3: Configuring Your Network.....	12
Network Configuration Overview.....	13
Using the Workbook.....	13
Dell Networking Switches Default Solution Values.....	14
Optional: S6010-ON Switches.....	15
Splitting Ports on S6010-ON.....	15
Third-Party Network Hardware.....	15
Chapter 4: Configuring PowerEdge Hardware.....	17
Configuring the SAH Node.....	18
IPMI Configuration.....	18
Open Source Hardware Configuration Toolkit.....	18
Configuring Overcloud Nodes.....	20
Configuring Server Network Settings.....	20
Repurposing Servers.....	21
Validating Server IPMI Configuration.....	21
Chapter 5: Dell Storage PS Series Storage Group.....	22
Dell Storage PS Series Configuration Information.....	23
Chapter 6: Dell Storage SC Series Storage Arrays.....	25
Dell Storage SC Series Configuration Information.....	26
Appendix A: References.....	27

To Learn More.....28

List of Tables

Table 1: Dell EMC PowerEdge R630/Dell EMC PowerEdge R730xd Tested BIOS and Firmware Versions.....	11
Table 2: Dell Storage Tested Software and Firmware Versions.....	11
Table 3: Dell Networking Tested Firmware Versions.....	11
Table 4: Example VLAN Assignments.....	13
Table 5: Switch Port Defaults.....	14
Table 6: iDRAC Specification for SAH Nodes.....	19
Table 7: SAH BIOS Specification.....	19
Table 8: PS Series Information Needed from Configuration.....	23
Table 9: SC Series Information Needed from Configuration.....	26

Trademarks

Copyright © 2014-2018 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.

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

Overview

Topics:

- [Deployment Methodology](#)
- [Intended Audience](#)
- [Dependencies](#)

This guide provides information necessary to deploy the Dell EMC Ready Bundle for Red Hat OpenStack NFV Platform, on Dell EMC PowerEdge R630 and Dell EMC PowerEdge R730xd servers with the Dell EMC PowerEdge H730 disk controller; and the network with Dell Networking S4048T-ON and S4048-ON switches.

Deployment Methodology

To perform a deployment of the Dell EMC Ready Bundle for Red Hat OpenStack NFV Platform:

1. Use the [Dell EMC Ready Bundle for Red Hat OpenStack NFV Platform Hardware Deployment Guide](#).
2. Then, depending on the methodology that you prefer, use the [Dell EMC Ready Bundle for Red Hat OpenStack NFV Platform Software Deployment Guide](#) to perform either:
 - a. An automated deployment using scripts and methods developed and validated by Dell EMC
 - b. A manual deployment using methods developed and validated by Dell EMC

Intended Audience

This guide assumes the reader is familiar with:

- OpenStack
- Dell EMC PowerEdge R630 and Dell EMC PowerEdge R730xd RAID configuration
- Red Hat Enterprise Linux (RHEL)
- Red Hat OpenStack Platform (RHOSP) documentation
- Network Configuration
- The concepts and procedures in [Red Hat's Red Hat OpenStack Platform update/upgrade documentation](#)

Dependencies

For customers performing a self-installation, these files are available on request from Dell EMC. Please contact your account representative, or email openstack@dell.com for instructions.

[Dell EMC Ready Bundle for Red Hat OpenStack NFV Platform Hardware Deployment Guide](#) dependencies and prerequisites include:

- IPMI Utilities - used for validating server IPMI configuration
 - For Linux® - <http://sourceforge.net/projects/ipmitool/>
 - For Windows® - <http://ipmiutil.sourceforge.net/>
- The Dell EMC Open Source Hardware Configuration Toolkit (OS-HCTK)
 - A copy of the archive: JetPack-manual-10.0.1.tgz, which contains:
 - The OS-HCTK ISO
 - Sample switch configs
 - Solution Workbook
 - Creating the OS-HCTK ISO requires a Linux® environment that has the following programs installed:
 - `cpio`
 - `rpm`
 - `rpm2cpio`

Chapter 2

Hardware Setup

Topics:

- *Unpacking and Installing the Equipment*
- *Powering Up the Equipment*
- *Verifying the Equipment*
- *Tested BIOS and Firmware*

These procedures ensure that your hardware is installed correctly prior to installing the Dell EMC Ready Bundle for Red Hat OpenStack NFV Platform.

Unpacking and Installing the Equipment

Before you proceed you must perform the following procedures following all standard industry safety procedures:

1. Unpack and install the racks.
2. Unpack and install the server hardware.
3. Unpack and install the switch hardware.
4. Unpack and install the network cabling.
5. Connect each individual machine to both power bus installations.
6. Apply power to the racks.



Note: This is usually performed by the Dell EMC EDT Team.

Powering Up the Equipment

To perform the power on test:



Note: This is usually performed by the Dell EMC EDT Team.

1. Power on each server node, individually.
2. Wait for internal system diagnostic procedures to complete.
3. Power on the network switches.
4. Wait for the switches' internal system diagnostic procedures to complete.

Verifying the Equipment

The cluster hardware should be verified before physical installation begins. After installation, the final functional tests should be run.

Recommended validation steps:

1. All power on tests complete successfully.
2. All drives should be powered on, verify that the hardware diagnostic LEDs and system console does not report any errors.
3. All nodes should be checked for correct memory size.
4. All network ports and cables should be checked for connections.

Tested BIOS and Firmware

[Table 1: Dell EMC PowerEdge R630/Dell EMC PowerEdge R730xd Tested BIOS and Firmware Versions](#) on page 11 lists the server BIOS and firmware versions that were tested for the Dell EMC Ready Bundle for Red Hat OpenStack NFV Platform.

[Table 2: Dell Storage Tested Software and Firmware Versions](#) on page 11 lists the Dell Storage Center, PS Series, and SC Series software and firmware versions that were tested for the Dell EMC Ready Bundle for Red Hat OpenStack NFV Platform.

[Table 3: Dell Networking Tested Firmware Versions](#) on page 11 lists the S4048T-ON, S4048-ON, and S6010-ON switch firmware versions that were tested for the Dell EMC Ready Bundle for Red Hat OpenStack NFV Platform.



Caution: You must ensure that the firmware on all servers, storage devices, and switches is up to date. Otherwise, unexpected results may occur.

Table 1: Dell EMC PowerEdge R630/Dell EMC PowerEdge R730xd Tested BIOS and Firmware Versions

Product	Version
BIOS	2.4.3
iDRAC Firmware	2.41.40.40
Lifecycle Controller	2.41.40.40
Intel® X520 10G NIC	18.0.17
PERC H730 RAID Controller	25.5.2.0001

Table 2: Dell Storage Tested Software and Firmware Versions

Product	Version
Dell Storage Center Software	2016 R2 Build 16.2.1.228
PS Series Firmware	9.0.5
SC Series Firmware	6.5.20.17

Table 3: Dell Networking Tested Firmware Versions

Product	Version
S4048T-ON Firmware	9.11(0.0)
S4048-ON Firmware	9.11(0.0P2)
S6010-ON Firmware	9.11(0.0P2)

Chapter

3

Configuring Your Network

Topics:

- [Network Configuration Overview](#)
- [Using the Workbook](#)
- [Dell Networking Switches Default Solution Values](#)
- [Optional: S6010-ON Switches](#)
- [Splitting Ports on S6010-ON](#)
- [Third-Party Network Hardware](#)

This topic describes the procedures required to configure the Dell EMC Ready Bundle for Red Hat OpenStack NFV Platform network.

Network Configuration Overview

Complete the following steps to set up your network:

1. Determine the VLANs to be used and how they are used.
2. Determine the IP ranges and associate them to a VLAN.
3. Fill in the Solution Workbook. See [Using the Workbook](#) on page 13.
4. Ensure that the cabling matches your Workbook.
5. Configure your switches per the Workbook:

- [Dell Networking Switches Default Solution Values](#) on page 14

 **Note:** To enable OVS-DPDK, refer to OVS-DPDK Appendix in [Dell EMC Ready Bundle for Red Hat OpenStack NFV Platform Version 10.0.1 Software Deployment Guide](#).

 **Note:** To enable SR-IOV, refer to [Dell EMC Ready Bundle for Red Hat OpenStack NFV Platform Version 10.0.1 SR-IOV User Guide](#).

Using the Workbook

 **Note:** The Solution Workbook was designed for Dell EMC and Red Hat deployment services to use as a general worksheet for planning a deployment. It is available upon request, as explained in [Dependencies](#) on page 8, for customers performing a self-deployment. For the Network Configuration portion of the deployment procedure, only the worksheets on Dell Networking are required. The rest of the worksheets are optional.

The *Solution Workbook* is set up to enable the installers to use the same information to:

- Program the switches
- Build the configuration for each node

Using the worksheet you must complete the following tasks for each connection from Server to Switch.

On Any Page:

- If the information is indicated as *required*, then it is needed in order to ensure a successful install.
- We provide a table for you to complete, similar to [Table 4: Example VLAN Assignments](#) on page 13.

 **Note:** When configuring your networks they must be aligned so that the tenant networks can be expanded by adding a VLAN for each OpenStack virtual network. This is accomplished by assigning all the other networks to lower VLAN numbers than that of the internal tenant network. The solution, as tested, uses a single external network for tenants. A service motion can be arranged to allow dedicated external tenant networks.

Table 4: Example VLAN Assignments

VLAN Name	Recommended VLAN
Management/Out of Band (OOB) Network	110
Provisioning Network VLAN	120
Tenant Tunneling Network VLAN	130
Private API Network VLAN	140
Storage Network VLAN	170

VLAN Name	Recommended VLAN
Storage Clustering Network VLAN	180
Public API Network VLAN (API, GUI)	190
External Network VLAN for Tenants (tenants' floating IP addresses)	191
Internal Networks VLAN for Tenants	200+

 **Note:** Management/Out of Band network ports Management 1/1 is used by the VLT for the heartbeat, and must terminate on the same VLAN in the same IP address range.

On the General Configuration Page:

- Fill in the information needed as required by the site.

On the Switch Configuration Pages:

- **Name** - The Port name on the switch.
- **Connector** - The type of connector used.
- **Device Name** - The name of the server/device connected to the switch. It is helpful to use a name that describes the server usage and associated bond.
- **Port** - The Network Interface Card (NIC) name and port number as the OS would generate.
- **Untagged** and **Tagged** - Used to indicate what VLAN(s) the port should be used, and how.
- **Port-channel number** - The unique number on the switch that represents the port-channel. Each pair in a server bond must have their own unique port-channel number, which is the same on both switches.
- **Mode** - The mode the port-channel will use *802.3ad* for all nodes.

Working switch configurations based on the validated solution are included within the Open Source Hardware Configuration Toolkit (OS-HCTK) `tgz` file.

Dell Networking Switches Default Solution Values

 **Note:** The VLT ports differ, and are defined in the example.

Table 5: Switch Port Defaults

CLI Command	Recommended Default	Example
<code>ip address</code>	No IP address assigned	<code>no ip address</code>
<code>portmode</code>	Hybrid mode	<code>portmode hybrid</code>
<code>switchport</code>	Enabled - after portmode hybrid executed	<code>switchport</code>
<code>mtu</code>	9216 on S4048T-ON; 9216 on S4048-ON	<code>mtu 9216</code> or <code>mtu 9216</code>
<code>flowcontrol</code>	Flow control receive on, transmit off	<code>flowcontrol rx on tx off</code>
<code>spanning-tree</code>	Rapid Spanning Tree set to Edge-Port going to servers.	<code>spanning-tree rstp edge-port</code>

Optional: S6010-ON Switches

The *Dell EMC Ready Bundle for Red Hat OpenStack NFV Platform Reference Architecture* is prescriptive in its hardware, software, networking, and installation definitions. This ensures a consistent experience when you create your OpenStack environment.

The Reference Architecture uses Dell Networking™ S4048-ON switches. However, when the solution is expanded, or initially built beyond ten (10) systems, it can benefit from using Dell EMC's newest aggregation and Top-of-Rack (TOR) switches - the Dell Networking S6010-ON switches. The S6010-ON provides the following benefits:

- Leverages a non-blocking switching architecture
- Delivers line-rate L2 and L3 forwarding capacity
- Provides up to 96 ports of 10Gbe, and eight (8) additional ports of 40Gbe

To utilize the S6010-ON in the solution:

1. Replace up to four S4048-ONs with two S6010-ONs.
2. Mount the two switches in separate racks, so that they draw power from different Power Distribution Units (PDUs).
3. Using the 40G to 10G splitter cables, wire the servers so that each of the two 10Gb ports is connected to a port on a separate switch (e.g., no single NIC has both ports connected to the same switch).

Splitting Ports on S6010-ON

There are several considerations you must take into account when splitting ports on Dell Networking S6010-ON switches:

- Splitting a single 40G port into four 10G ports is supported only on a standalone unit.
- Split ports cannot be used as stack-link to stack an S6010-ON.
- Split ports cannot be a part of any stacked system.

To split a port:

1. Log onto the switch.
2. Execute the following command:

```
stack-unit 0 port 48 portmode quad
```

Third-Party Network Hardware

If you are not using Dell Networking S4048T-ON and/or S4048-ON switches, you must program the switches to support your cloud instantiation. Your switches are expected to support the following:

- Support for IEEE 802.1Q VLAN traffic and port tagging
- Support for using one untagged and multiple tagged VLANs on the same port
- The ability to provide a minimum of 170 Gigabit Ethernet ports, in a non-blocking configuration, within the Provisioning VLAN
 - Configuration can be a single switch or a combination of stacked switches, to meet the additional requirements
- The ability to create link aggregation groups (LAGs) with a minimum of two physical links in each LAG

- If multiple switches are stacked:
 - The ability to create a LAG across stacked switches
 - Full-bisection bandwidth
 - Support for VLANs to be available across all switches in the stack
- 250,000 packets-per-second capability per switch
- A managed switch that supports both SSH and serial line configuration
- SNMP v3 support

Chapter 4

Configuring PowerEdge Hardware

Topics:

- [Configuring the SAH Node](#)
- [Configuring Overcloud Nodes](#)

This section describes manually configuring PowerEdge server hardware for the Dell EMC Ready Bundle for Red Hat OpenStack NFV Platform with Red Hat OpenStack Platform:

- IPMI Configuration
- BIOS Configuration
- RAID Configuration

Configuring the SAH Node

The SAH is configured using the Open Source Hardware Configuration Toolkit (OS-HCTK) tool.

IPMI Configuration

The servers' iDRACs must be configured correctly for Dell EMC Ready Bundle for Red Hat OpenStack NFV Platform deployment. Dell EMC recommends that you use the OS-HCTK to configure the iDRAC on the SAH node. See [iDRAC Default Settings](#) on page 19 for further details. Overcloud nodes' iDRACs should be configured correctly by the factory, with the possible exception of network settings. See [Configuring Server Network Settings](#) on page 20 for details on configuring iDRAC network settings.

Open Source Hardware Configuration Toolkit

The Open Source Hardware Configuration Toolkit (OS-HCTK) is a configuration utility with sample scripts and configuration files that is used to automate the setup and configuration of BIOS and RAID settings for Dell EMC servers used for OpenStack and Hadoop open source software solutions.

 **Note:** The OS-HCTK ISO is run **only** on the system that will be configured as the SAH.

The OS-HCTK enables you to create a USB key from which you can boot a Dell EMC PowerEdge R630 or Dell EMC PowerEdge R730xd, and apply the BIOS and RAID settings.

Topics discussed include:

- [Creating the Open Source Hardware Configuration Toolkit ISO](#) on page 18
- [Customizing the ISO](#) on page 19
- [iDRAC Default Settings](#) on page 19
- [SAH BIOS Specification](#) on page 19
- [Running the Open Source Hardware Configuration Toolkit ISO](#) on page 20

Creating the Open Source Hardware Configuration Toolkit ISO

To create the OS-HCTK ISO:

1. Ensure you are running in a Linux[®] environment that has the following programs installed:
 - `cpio`
 - `rpm`
 - `rpm2cpio`
2. Ensure the user you are logged in as has passwordless `sudo` rights.
3. Copy the `bootimg.iso` included in the OS-HCTK `tgz` file to your home directory.
4. Run the following command to ensure that the `loop` module is loaded, so that the script has access to loopback devices:

```
$ sudo modprobe loop
```

5. Copy the ISO onto a USB key using the following command:

```
$ sudo dd if=~/.bootimg.iso of=/dev/sdx bs=2048
```

Where: `/dev/sdx` is the device that was created when the USB key was inserted into the system. The USB key must be at least 512MB in size.

Customizing the ISO

This utility creates a CentOS 6.6 Live CD ISO that has the Dell EMC OS-HCTK installed. It uses `syscfg`, `raidcfg` and `racadm` to configure the system. The main customization point is `bootimg/node-config.sh`, which automatically runs when the ISO finishes booting.

iDRAC Default Settings

[Table 6: iDRAC Specification for SAH Nodes](#) on page 19 lists and describes iDRAC default `racadm` settings that will be set by the OS-HCTK.

Table 6: iDRAC Specification for SAH Nodes

Menu Choice	iDRAC Setting
iDRAC.IPMLan.Enable	Enabled
iDRAC.IPMLan.PrivLimit	4
iDRAC.IPv4.Enable	Enabled
iDRAC.Users.2.Enable	Enabled
iDRAC.Users.2.IpmiLanPrivilege	4
iDRAC.Users.2.Privilege	0x1ff
iDRAC.WebServer.Enable	Enabled

SAH BIOS Specification

[Table 7: SAH BIOS Specification](#) on page 19 lists and describes the default BIOS settings for the OpenStack servers that will be set by the OS-HCTK.

Table 7: SAH BIOS Specification

Display Name	Attribute	Settings
Boot Mode	BootMode	BIOS
Boot Sequence Retry	BootSeqRetry	Enabled
DCU IP Prefetcher	DculpPrefetcher	Enabled
DCU Streamer Prefetcher	DcuStreamerPrefetcher	Enable
Logical Processor Idling	DynamicCoreAllocation	Disabled
Integrated RAID Controller	IntegratedRaid	Enabled
Internal SD Card	InternalSdCard	Off
I/OAT DMA Engine	IoatEngine	Enabled
Logical Processor	LogicalProc	Enabled
Memory Operating Mode	MemOpMode	OptimizerMode
System Memory Testing	MemTest	Disabled
Node Interleaving	NodeInterleave	Disabled
OS Watchdog Timer	OsWatchdogTimer	Disabled
Adjacent Cache Line Prefetch	ProcAdjCacheLine	Enabled
Number of Cores per Processor	ProcCores	all

Display Name	Attribute	Settings
Execute Disable	ProcExecuteDisable	Enabled
Hardware Prefetcher	ProcHwPrefetcher	Enabled
CPU Power Management	ProcPwrPerf	MaxPerf
Turbo Mode	ProcTurboMode	Enabled
Virtualization Technology	ProcVirtualization	Enabled
QPI Speed	QpiSpeed	MaxDataRate
Alternate RTID (Requestor Transaction ID) Setting	RtidSetting	Disabled
SR-IOV Global Enable	SriovGlobalEnable	Enabled
System Profile	SysProfile	PerfOptimized

Running the Open Source Hardware Configuration Toolkit ISO

 **Note:** The OS-HCTK ISO is run **only** on the system that will be configured as the SAH.

To run the OS-HCTK ISO:

1. Boot the SAH to be configured from the USB key.
2. Once the SAH finishes booting it displays the detected RAID controllers, along with their current configuration, and offers you a choice of system configurations.
 - a. Select **OpenStack SAH**.
3. The OS-HCTK automatically configures the RAID and BIOS settings. Once finished, it will ask you to provide basic iDRAC connectivity information, including:
 - a. Whether the iDRAC should use DHCP
 - b. Or, basic IPv4 settings if not using DHCP
4. The system configures the iDRAC with some default settings plus the network settings from Step 3 above. Once it finishes applying those settings, it prompts you to reboot the system.
5. Assuming there were no errors, remove the USB key and then reboot the system.

The changes are applied, and the system is configured for its role.

Configuring Overcloud Nodes

This topic describes procedures you will use to configure Overcloud nodes. Procedures described, in the order they should be performed, include:

1. [Configuring Server Network Settings](#) on page 20
2. [Repurposing Servers](#) on page 21
3. [IPMI Configuration](#) on page 18
4. [Validating Server IPMI Configuration](#) on page 21

Configuring Server Network Settings

1. Set the iDRAC IP address source:
 - a. If the Overcloud nodes were ordered with the iDRACs configured for DHCP, or are currently configured for DHCP, then no further configuration is necessary.

- b. If you wish to use static IP addresses, then configure the Overcloud nodes' DRAC IP address, subnet mask, default gateway IP, and default VLAN (ID = 110, if required) using the iDRAC GUI.

Repurposing Servers

In addition to configuring the network, ensure that the following settings are configured as indicated:

1. Set the iDRAC NIC mode to *Dedicated*.
2. Configure the IPMI over LAN Setting to *Enabled*.

Configuring Server User Information

1. Set credentials for the root user, including changing the password based upon good password standards.
2. Set privileges for the user to the Admin level, including over LAN.
3. Enable the user, if disabled.

Validating Server IPMI Configuration

Validating that remote commands can be executed is an essential part of the IPMI setup.

1. Install the IPMI Utilities to your workstation from SourceForge:
 - a. For Linux® - <http://sourceforge.net/projects/ipmitool/>
 - b. For Windows® - <http://ipmiutil.sourceforge.net/>
 - c. Validate that you have all the requirements, and that it will run.
2. Plug your Ethernet port into a switch port that is on the same VLAN as your iDRACs.
3. Configure your NIC to use an IP address in the iDRAC network range.
4. Execute the following IPMI command, replacing "username" and "password" with the credentials for the iDRACs:

```
for i in $(seq 162 170); do ipmitool -P "password" -U "username" -I
lanplus -H 192.168.200.$i power status; done
```

This will perform a simple, non-destructive poll of the power status of the iDRAC from 192.168.200.162 to 192.168.200.170.

 **Note:** You will need to replace the IP address range and subnet with correct information for the deployed iDRAC subnet.

- a. You can replace the keyword status with *reset*, *off*, or *on*.

 **Note:** These may change the power state of the nodes.

5. Ensure that all machines return responses to the command.

Chapter 5

Dell Storage PS Series Storage Group

Topics:

- [Dell Storage PS Series Configuration Information](#)

The PS Series Storage Group can consist of one or more storage arrays with one or more storage groups.



Note: The configuration of the arrays is beyond the scope of this document. Please refer to the Dell Storage PS Series Support Website (<https://eqsupport.dell.com/secure/login.aspx>) for the latest guides, whitepapers, and best practices on how to setup your Storage Group for your application.

Dell Storage PS Series Configuration Information

Once the Storage Group(s) are setup, the information contained in [Table 8: PS Series Information Needed from Configuration](#) on page 23 must be collected to configure your storage backend.

 **Note:** To reduce volume initialization time, Dell EMC recommends that `san_thin_provision` be set to `true`.

More information can be found at <https://access.redhat.com/documentation/en/red-hat-openstack-platform/8/dell-equallogic-back-end-guide/dell-equallogic-back-end-guide>.

Table 8: PS Series Information Needed from Configuration

[DEFAULT]	Description
<code>volume_driver = cinder.volume.drivers.eqlx.DellEQLSanISCSIDriver</code>	Dell Storage PS Series volume driver
<code>san_ip = <IP_address_of_EQLX></code>	IP address used to reach the PS Series Group through SSH
<code>san_login = <user_name></code>	User name to login to the Group manager via SSH at the <code>san_ip</code>
<code>san_password = <password></code>	Password to login to the Group manager via SSH at the <code>san_ip</code> (not used when <code>san_private_key</code> is set)
<code>san_thin_provision = <true false></code>	Enable/disable creation of thin-provisioned volumes
<code>san_ssh_port = 22</code>	Port used for SSH
<code>ssh_conn_timeout = 30</code>	Timeout value, in seconds, used by CLI commands over SSH
<code>san_private_key = <filename></code>	Filename of the private key used for SSH authentication
<code>ssh_min_pool_conn = 1</code>	Minimum number of SSH connections in the pool
<code>ssh_max_pool_conn = 5</code>	Maximum number of SSH connections in the pool
<code>eqlx_chap_login = admin</code>	Existing CHAP account name
<code>eqlx_chap_password = password</code>	Password for specified CHAP account name
<code>eqlx_cli_max_retries = 5</code>	Maximum retry count for reconnection
<code>eqlx_cli_timeout = 30</code>	Timeout for the Group Manager CLI command execution
<code>eqlx_group_name = group-0</code>	Group name to use for creating volumes

[DEFAULT]	Description
eqlx_pool = default	Pool in which volumes will be created
eqlx_use_chap = False	Use CHAP authentication for targets?

Chapter 6

Dell Storage SC Series Storage Arrays

Topics:

- [Dell Storage SC Series Configuration Information](#)

The SC Series can consist of one or more Dell Storage Centers with Dell Storage Enterprise Manager platform.



Note: The configuration of the cluster is beyond the scope of this document. Please refer to the Dell Storage SC Series support website, <http://www.dell.com/support/contents/us/en/19/article/Product-Support/Dell-Subsidiaries/compellent>, for the latest guides, white papers, and best practices on how to setup your storage cluster.

Dell Storage SC Series Configuration Information

Once the Dell Storage SC Series with the Dell Storage Enterprise Manager platform is setup according to the [Dell EMC Ready Bundle for Red Hat OpenStack NFV Platform Reference Architecture](#), the information contained in [Table 9: SC Series Information Needed from Configuration](#) on page 26 must be collected to configure your storage backend.

More information can be found at <https://access.redhat.com/documentation/en/red-hat-openstack-platform/version-8/dell-storage-center-back-end-guide/>

Table 9: SC Series Information Needed from Configuration

[DEFAULT]	Description
Required Values	
<code>volume_backend_name = delliscsi</code>	Name given to the storage backend
<code>volume_driver = cinder.volume.drivers.dell.dell_storagecenter_iscsi.DellStorageCenterISCSIDriver</code>	Dell Storage SC Series iSCSI volume driver
<code>san_ip = <IP_address></code>	IP address of Enterprise Manager
<code>san_login = <user_name></code>	User name to log into Enterprise Manager at the <code>san_ip</code>
<code>san_password = <password></code>	Password to log into the Enterprise Manager at the <code>san_ip</code>
<code>iscsi_ip_address = <IP_address></code>	The Storage Center iSCSI IP address
<code>dell_sc_ssn = <serial_number></code>	The Storage Center serial number to use
Optional Vaules	
<code>dell_sc_api_port = <port_to_use></code>	Configured Enterprise Manager API port, default is 3033
<code>dell_sc_server_folder = <folder_name></code>	Server folder in which to place new server definitions
<code>dell_sc_volume_folder = <folder_name></code>	Volume folder in which to place created volumes
<code>iscsi_port = <port_number></code>	iSCSI port to use, if you do not wish to use the default port number 3260

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 NFV Platform visit <http://www.dell.com/learn/us/en/04/solutions/red-hat-openstack>.

Copyright © 2014-2018 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.