Dell EMC VxFlex Ready Node

Operating System Installation and Configuration Guide for Linux

Rev 01

October 2019

DELL EMC

Copyright © 2019 Dell Inc. or its subsidiaries All rights reserved.

Dell believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS-IS." DELL MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. USE, COPYING, AND DISTRIBUTION OF ANY DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA.

Dell EMC
Hopkinton, Massachusetts 01748-9103
1-508-435-1000 In North America 1-866-464-7381
www.DellEMC.com

Contents

Figures **5**

Tables **7**

Chapter 1 VxFlex Ready Node deployment overview **9**
Typographical conventions 10
Hardware and operating systems 10
Supported operating systems and requirements 10
VxFlex OS packages 12
VxFlex OS component requirements 13
VxFlex OS cluster components 13
VxFlex OS Gateway server requirements—VxFlex Ready Node 13
General prerequisites 14
Java requirements 14
Supported Internet browsers 14
Connectivity requirements 14
Disk prerequisites 15
iDRAC Service Module 15
openIPMI (Linux servers only) 15
Required packages for NVDIMM 15
Supported hardware configurations 15
Additional equipment and network resource requirements 15

Chapter 2 Configuring the hardware **17**
Set up the iDRAC IP address and BIOS 18
Verify the status of the system hardware and drives 20
Log in to the KVM console 21
Updating the BIOS, firmware and settings 22
Update the hardware using the iDRAC virtual console 22

Chapter 3 Installing the Linux operating system **25**
Linux system requirements 26
Rebuild the RAID 1 boot device using the replaced M.2 module 26
Map the Linux ISO file on a VxFlex Ready Node server 27
Optimize CPU performance on RHEL systems 27
Optimize CPU performance on SUSE systems 28

Chapter 4 Network architecture and physical connectivity for Linux servers **29**
Networking connectivity architecture and cabling best practice for Linux deployments 30
IP addresses for R640 and R740xd servers 30
IP addresses for R840 servers 32

Chapter 5 Configuring network ports on Linux servers **35**
Linux ports overview 36
Configure network ports on R640 or R740xd servers 36
Determine the 2U1N node type 36
R640 and R740xd server NIC and iDRAC port locations 36
VxFlex Ready Node R640 port designations - Linux 38
VxFlex Ready Node R740xd port designations - Linux 38
Correlate the PCI slot locations and interface names on R640/R740xd servers 39
Configure network ports on R840 servers 45

R840 server NIC and iDRAC port locations....45
VxFlex Ready Node R840 Linux server port designations.... 46

Chapter 6 Installing the Drivers....53
Install the VxFlex OS drivers on a Linux server....54

Chapter 7 Preparing Disks.... 55
Verify the disk controller type....56
Enable PERCCLI 56
Create virtual drives with PERCCLI....57
Verify the storage and disk controller status 57
Create virtual drives for an HDD using PERCCLI 59
Create virtual drives for SSD using PERCCLI 59
Ensure virtual drive creation with PERCCLI....60
Retrieve device paths on the server.... 61
Retrieve device paths in a Linux server with an HBA330 controller.... 61
Retrieving device paths in a Linux server with an H730p/H740p controller.... 62
Retrieving device paths in a Linux server with NVMe devices configured.... 63

Chapter 8 Additional configurations.... 65
Install OpenManage Enterprise.... 66
Disable Smartmontool error messages....66
iDRAC Service Module....67
Install the iDRAC Service Module on a Linux server.... 67
Prepare the DAX devices.... 67
NVDIMM information table.... 67
Prepare an NVDIMM as a DAX device....68
iDRAC Service Module....71
iDRAC Service Module....72

Chapter 9 Next steps....73
Next steps....74

Chapter 10 Reference material....75
DTK - Hardware Update Bootable ISO.... 76
Dell EMC OpenManage DRAC Tools (RACADM).... 76
Recommended BIOS and firmware settings....78
Troubleshooting the Hardware ISO....80
Additional resources.... 82

Chapter 11 Getting help.... 85
Contacting Dell EMC.... 86
Secure Remote Services....86
Recycling or End-of-Life service information....86

Figures

1 R640 and R740xd PCI slots, integrated NICs and BMC port locations....37
2 Correlating physical slot and Ethernet network names....41
3 R840 PCI slots, integrated NICs, and iDRAC port locations....45
4 Correlating physical slot and Ethernet network names.... 48

Tables

1 Linux.....10
2 Windows, except Core systems (see note).....11
3 Linux.....11
4 Windows.....12
5 VxFlex OS management IP network.....30
6 VxFlex OS data IP network for Subnet #1.....30
7 VxFlex OS data IP network for Subnet #2.....31
8 VxFlex OS IP network for 2-layer configuration (Data 3 + Data 4).....31
9 VxFlex OS management IP network.....32
10 VxFlex OS data IP network for Subnet #1.....32
11 VxFlex OS data IP network for Subnet #2.....32
12 VxFlex OS IP network for 2-layer configuration (Data 3 + Data 4).....33
13 R640 single-node server configuration.....38
14 R740xd single-node server configuration for 10/25GbE, 2CPU, SFP+/SFP28.....39
15 R840 server configuration for 2x10GbE+2GbE rNDC, 2x10/25GbE PCIe slot 3, 2x10/25GbE PCIe slot 5.....46
16 NVDIMM information table.....67
17 Hardware ISO configuration settings.....79

CHAPTER 1

VxFlex Ready Node deployment overview

This guide explains how to install the operating system and perform system configurations such as network connectivity and ports, on VxFlex Ready Node servers.

Use this guide after installing a new or replacement VxFlex Ready Node server at the customer site, before deploying the system. Perform the procedures in this guide in the order they are presented.

- Typographical conventions 10
- Hardware and operating systems 10
- VxFlex OS component requirements 13
- General prerequisites 14
- Supported hardware configurations 15
- Additional equipment and network resource requirements 15

Typographical conventions

Dell EMC uses the following type style conventions in this document:

Bold	Used for names of interface elements, such as names of windows, dialog boxes, buttons, fields, tab names, key names, and menu paths (what the user specifically selects or clicks)
Italic	Used for full titles of publications referenced in text
`Monospace`	Used for: • System code • System output, such as an error message or script • Pathnames, filenames, prompts, and syntax • Commands and options
Monospace italic	Used for variables
`Monospace bold`	Used for user input
[ ]	Square brackets enclose optional values
|	Vertical bar indicates alternate selections - the bar means “or”
{ }	Braces enclose content that the user must specify, such as x or y or z
...	Ellipses indicate nonessential information omitted from the example

Hardware and operating systems

Ensure that all hardware is assembled as described in the *VxFlex Ready Node Hardware Installation Guide*.

Unless otherwise specified in this guide, install the operating system in the server's BOSS device and configure the IP addresses.

Supported operating systems and requirements

The following is a list of supported operating systems and additional requirements for this version of VxFlex OS on VxFlex Ready Node servers.

For the most updated information, see the *Dell EMC VxFlex Ready Node Firmware and Driver Matrix* at https://support.emc.com/products/42216_ScaleIO-Ready-Node--PowerEdge-14G.

VxFlex Ready Node R640 and R740xd

Table 1 Linux

Component	Requirement
Supported OS versions	• RHEL 7.3, 7.5, 7.6 • SLES 12.2, 12.3, 12.4 • Oracle Linux 7.5, 7.6 with Red Hat Kernels
Required packages, all Linux distributions	numactl libaio

Table 1 Linux (continued)

Component	Requirement
Additional requirements for specific operating systems	SLES 12.3, 12.4: • hwinfo • net-tools • pciutils • ethtool Hypervisor support: Red Hat KVM
Additional packages required for MDM components	bash-completion (for SCLI completion) Latest version of Python 2.X

Secure authentication mode	Ensure that OpenSSL 64-bit v1.0.1 or later is installed on all servers in the system:
	`libopenssl1_0_0-1.0.1g-0.40.1.x86_64.rpm`
	`openssl1-1.0.1g-0.40.1.x86_64.rpm`
	Note: (v1.1 is not supported)

LDAP	Ensure that OpenLDAP 2.4 is installed on all servers

Table 2 Windows, except Core systems (see note)

Component	Requirement
Supported OS versions	2016 + Hyper-V, 2019 + Hyper-V (GUI, SDC, and LIA) Including the Visual C++ redistributable 2010 package, 64-bit
Secure authentication mode	Ensure that the following are installed on all servers in the system: • OpenSSL 64-bit v1.0.1 or later (v1.1, however, is not supported) • Visual C++ redistributable 2010 package, 64-bit

Note: Windows OS is supported for the VxFlex OS Gateway and SDC, and is required for LIA. Windows OS is not supported in Core systems (MDM and SDS). For specific VxFlex OS Gateway requirements, see "VxFlex OS Gateway server requirements" .

VxFlex Ready Node R840

Table 3 Linux

Component	Requirement
Supported OS versions	• RHEL 7.6 • SLES 12.4 • CentOS 7.6 • Oracle Linux 7.5, 7.6 with Red Hat Kernels
Required packages, all Linux distributions	numactl libaio

Table 3 Linux (continued)

Component	Requirement
Additional requirements for specific operating systems	SLES 12.4: • hwinfo • net-tools • pciutils • ethtool Hypervisor support: Red Hat KVM
Additional packages required for MDM components	bash-completion (for SCLI completion) Latest version of Python 2.X

Secure authentication mode	Ensure that OpenSSL 64-bit v1.0.1 or later is installed on all servers in the system:
	`libopenssl1_0_0-1.0.1g-0.40.1.x86_64.rpm`
	`openssl1-1.0.1g-0.40.1.x86_64.rpm`
	Note: (v1.1 is not supported)

LDAP	Ensure that OpenLDAP 2.4 is installed on all servers

Table 4 Windows

Component	Requirement
Supported OS versions	2019 (GUI, SDC, and LIA) Including the Visual C++ redistributable 2010 package, 64-bit
Secure authentication mode	Ensure that the following are installed on all servers in the system: • OpenSSL 64-bit v1.0.1 or later (v1.1, however, is not supported) • Visual C++ redistributable 2010 package, 64-bit

Note: Windows OS is supported for the VxFlex OS Gateway and SDC, and is required for LIA. Windows OS is not supported in Core systems (MDM and SDS). For specific VxFlex OS Gateway requirements, see VxFlex OS component requirements on page 13.

VxFlex OS packages

Download the VxFlex OS complete software ZIP file and extract the folders for your operating system.

The files can be downloaded from the Support site. Refer to the EMC Simple Support Matrix (ESSM) for supported versions.

Note: You are also required to download and extract the VxFlex OS GUI for Windows Software Download. For Linux 2-layer/storage-only, the VxFlex OS Complete VMware Software Download is also required.

VxFlex OS component requirements

Components and servers in the VxFlex OS system must meet the following requirements:

VxFlex OS cluster components

The following is the list of required VxFlex OS servers:

VxFlex OS component servers

- 3-node cluster
  - One Master MDM
  - One Slave MDM
  - One Tie Breaker
  - Minimum of three SDSs (on the same servers as the above components, or on three different servers)
  - SDCs, up to the maximum allowed (on the same servers as the above components, or on different servers)
- 5-node cluster
  - One Master MDM
  - Two Slave MDMs
  - Two Tie Breakers
  - Minimum of three SDSs (on the same servers as the above components, or on three different servers)
  - SDCs, up to the maximum allowed (on the same servers as the above components, or on different servers)

VxFlex OS Gateway

Install the VxFlex OS Gateway server on a separate server, or together with an MDM or SDS. Do not install the VxFlex OS Gateway on an SDC server or on an SDS on which RFcache will be enabled.

VxFlex OS Gateway server requirements—VxFlex Ready Node

The following is a list of server requirements for the VxFlex Ready Node VxFlex OS Gateway:

Supported operating systems - VxFlex Ready Node R640 and R740xd

- Windows 2016 or 2019, including the Visual C++ redistributable 2010 package, 64-bit. Server Core editions are not supported.
- Linux:
  - CentOS 7.x
  - Oracle Linux 7.x
  - Red Hat 7.x
  - SUSE 12.x
- rpmlib version should be 4.6.0-1 or higher

Every server requires 2 cores and a minimum of 3 GB available RAM.

Supported operating systems - VxFlex Ready Node R840

- Windows 2019 , including the Visual C++ redistributable 2010 package, 64-bit. Server Core editions are not supported.
- Linux RHEL 7.x, SUSE 12.x

- "rpmlib" version should be 4.6.0-1 or higher

Every server requires 2 cores and a minimum of 3 GB available RAM.

Java

- v1.8 (64-bit), build 149 or earlier.

You can download previous versions from this link: http://www.oracle.com/technetwork/java/javase/downloads/java-archive-javase8-2177648.html

Connectivity

The VxFlex OS Gateway server must have connectivity to all the nodes that are being installed. If you are using separate networks for management and data, the server must be able to communicate with both networks.

The following TCP ports are not used by any other application, and are open in the local firewall of the server: 80 and 443 (or 8080 and 8443).

You can change the default ports. For more information, see "Communication Security Settings" in the *Dell EMC VxFlex OS Security Configuration Guide* .

Additional requirements - R640/R740xd only

For a VxFlex OS Gateway server on a Windows node, the Windows Management Instrumentation service must be enabled on the VxFlex OS Installer server and on all Windows VxFlex OS nodes.

General prerequisites

Before you start the procedures in this guide, ensure that the server and installation environment meet the following additional general prerequisites:

Java requirements

Both the VxFlex OS GUI and iDRAC management client (Java console) require the Java 1.8 latest update.

In your internet browser, ensure that pop-ups are enabled for Java console features.

Supported Internet browsers

The VxFlex OS GUI supports the following Internet browsers:

- Google Chrome 72
- Firefox 65.0.1
- Internet Explorer 11

Connectivity requirements

Ensure that you have the following information at hand:

General connectivity

- The IP address range, subnet, and gateway IP addresses for the management and data network for all servers.
- The IP address range, subnet, and gateway IP addresses for the iDRAC port for all servers.
- All root and administrator passwords that you plan to set or use on the server and iDRAC.

Connection via a service PC or laptop computer using SSH/RDP

Ensure that the computer has a functional network port and an available IP address, subnet, and gateway on the management network.

Console operations (KVM access)

Ensure that you have either a VGA tool kit to allow console connection from a laptop computer to a server, or a computer screen and keyboard connection to the rack.

Disk prerequisites

For R640 and R740xd systems only, if an H730p/H740p controller card is installed on the server, install the PERCCLI disk utility.

iDRAC Service Module

If an NVMe device is installed on the server, install the iDRAC Service Module (iSM), as described later in this guide.

openIPMI (Linux servers only)

For the iDRAC Service Module package to install properly on a Linux server, ensure that openIPMI is installed and operational on the node. You can install openIPMI via yum package manager.

Required packages for NVDIMM

When NVDIMM acceleration devices are installed in the system, the following Linux packages are required:

- ndctl – NVDIMM control utility, version 62 or above
- libpmem – persistent memory library, version 1.4.3
- jq – Command-line JSON processor, version 1.5

For nodes running on RHEL 7.6 or higher, ensure that `kmod-redhat-nfit-3.10.0_957-1.el7_6.x86_64.rpm` is installed.

Supported hardware configurations

Supported hardware configurations are listed in the EMC Simple Support Matrix (ESSM) for your software version.

Download the ESSM from https://elabnavigator.emc.com/eln/elnhome. Click **Simple Support Matrices** > **Storage** and select the desired product and version.

Additional equipment and network resource requirements

Ensure that you have the following equipment and network resources and available before installing VxFlex Ready Node servers.

Boot drive

All configurations have two 120/240 GB BOSS boot drives (configured as RAID 1).

Switch requirements

One 10/25 GbE switch is required. Two or more 10/25GbE switches may be used for high availability and load balancing purposes. Separate (one or more) 1 GbE switches may be used for the management network.

Before you begin connecting the switches, ensure that they are installed and configured.

CAUTION Do not connect the switches to your organization's network until all installation and configuration activities are complete.

Each node has at least four 10GbE or two 25 GbE ports. Two are used for internal network communications. The other two ports are used for client network traffic. On ESXi nodes, they are also used for VMware high availability and DRS.

In cases of only two 25 GbE ports, the ports are also used for other network traffic.

The switches must have sufficiently available network ports to accommodate the following:

- Data network 10/25 GbE switches: Two 10/25 GbE ports per node, per switch
- Management network switches: One 1/10 GbE and one iDRAC port per node.

Network requirements are described later in this document.

Cabling

Ensure that you have the correct quantity of cables. Calculate the number of cables needed, according to the following guidelines:

For RJ45 NIC ports on your switch, Cat6 6 m Ethernet cables are required, four cables per node. (These cables are included with the server nodes)

CHAPTER 2

Configuring the hardware

This section describes how to configure the hardware, set iDRAC IP addresses, and map the ISO for servers in a VxFlex Ready Node environment.

- Set up the iDRAC IP address and BIOS 18
- Verify the status of the system hardware and drives 20
- Log in to the KVM console 21
- Updating the BIOS, firmware and settings 22

Set up the iDRAC IP address and BIOS

Set up the iDRAC IP address and set up or validate the BIOS on the VxFlex Ready Node servers.

Before you begin

Ensure that you have access to, or have the details for:

- The KVM console
- The server iDRAC IP address
- The server iDRAC subnet mask
- The gateway IP address
- The VLAN ID of the iDRAC, if the VLAN is used

Note: The iDRAC IP address may be an IPv4 or an IPv6 address.

About this task

Defaults:

- iDRAC username - root
- iDRAC password - Scaleio123

Note: Dell EMC recommends that you change the iDRAC password as soon as possible, because leaving the default password may create a security risk.

During the iDRAC IP and BIOS setup, use the following keyboard operations:

- Use the arrow keys to navigate in the BIOS screens.
- Press + or - to change the option selection in the BIOS screens.
- Press the Spacebar or Enter to change the settings.
- Press Enter to open a list of possible values.

Procedure

1. Log in to the KVM console.

   For console operations (KVM access), ensure that you have either a VGA tool kit/Crash Cart to allow physical console connection from a laptop computer to a server, or a computer monitor and keyboard connection to the rack.

2. Press F2 to access the main menu.
3. From the **System Setup Main Menu** page, select the **iDRAC Settings** menu.
4. Configure the network settings:

   a. In the **iDRAC Settings** page, select **Network**.

   b. In the **iDRAC Settings** > **Network** pane, verify the following parameter values:

      - **Enable NIC** = **Enabled**
      - **NIC Selection** = **Dedicated**

   c. From the **IPv4 Settings** pane, configure the IPv4 parameter values for the iDRAC port:

      - **Enable IP IPv4** = **Enabled**
      - **Enable DHCP** = **Disabled**

- **Static IP Address** = Static IP address (customer-provided). The static IP address must be accessible by the remote computer that will be used for system setup.
- **Static Gateway** = Gateway IP address
- **Static Subnet Mask** = Subnet mask IP address

d. From the **IPv6 Settings** pane, configure the IPv6 parameter values for the iDRAC port.

e. From the **IPMI Settings** pane, verify the following parameter values:

- **Enable IPMI Over LAN** = **Enabled**
- **Channel Privilege Level Limit** = **Administrator**

f. If you are working in a VLAN setup, access the **VLAN Configuration** pane and configure the VLAN ID parameters.

g. When the parameter set up is complete, click **Back** to display the **iDRAC Settings** page.

5. From the **iDRAC Settings** page, click **Finish**, **Yes**, and then **OK** to return to the **System Setup Main Menu** page.
6. In the **System Setup Main Menu** page, select the **System BIOS** menu.
7. In the **System BIOS Settings** page, verify that the processor settings are correct.

   If the settings are incorrect, configure them as follows:

   a. Select **Processor Settings**.

   b. In the **System BIOS Settings** > **Processor Settings** pane, verify the following parameter values:

   - **Virtualization Technology** = **Enabled**
   - **Number of Cores Per Processor** = **All**

   c. Click **Back** to return to the **System BIOS Settings** page.
8. Configure boot settings:

   a. Select **Boot Settings**.

   b. In the **System Boot Settings** > **Boot Settings** pane, verify that **Boot Mode** is set to **UEFI**.

   c. Select the **BIOS Boot Settings** link.

   d. In the **System BIOS Settings** > **Boot Settings** > **BIOS Boot** pane, verify that:

   - In the **Boot Sequence** list, **Hard drive C:** appears as the first item.
   - In the **Hard-Disk Drive Sequence** list, the BOSS device appears as the first item.

   e. Click **Back** twice to return to the **System BIOS Settings** page.
9. Configure integrated devices (R740xd/R640 servers only):

   a. Select **Integrated Devices**.

   The **System BIOS Settings--Integrated Devices** screen appears.

   b. Set **SR-IOV Global Enable** to **Enabled**.

   c. Verify that the **Internal USB Port** parameter is set to **Off**.

   d. Click **Back** to return to the **System BIOS Settings** screen.
10. From the **System BIOS Settings** page, click **Finish**, **Yes**, and then **OK** to return to the **System Setup Main Menu** page.

11. Select **Finish** to exit the BIOS and apply all settings post boot.

Results

The iDRAC IP and server BIOS address configuration is complete.

Verify the status of the system hardware and drives

Verify the status of the system hardware and drives in a VxFlex Ready Node server.

Before you begin

Ensure that you know:

- The IP address of the iDRAC port
- The username and password for the iDRAC portal (default username and password are root and password)

Procedure

1. From a browser, go to http://*<iDRAC_IP_address>*.
2. In the **DELL Console Login** window, enter the user name and password, and then click **Login**.

   The **Dashboard** displays the high-level status of all hardware devices in the **System Health** pane.
   In an ideal scenario, all hardware sensors are green.
3. Navigate to **Maintenance** > **System event log.**

   The **System Event Log** pane is displayed with color-coded severity levels.
4. Ensure that any power supply- and fan-related events in the event log are non-repetitive. Repetitive events may be due to the intermittent nature of faults, such as poor physical connections.

   For repetitive events, it is recommended that you remove the relevant hardware module and replace it in its socket.
5. For R640/R740xd systems only, in the navigation pane, select **Storage** > **Overview** > **Controller**.

   A table displays the controller-related information, with the controller type in the **Name** column. For example:

   ```
   PERC H740P Mini (Embedded)
   ```

6. In the navigation pane, select **Storage** > **Physical Disks**.

   A table displays information regarding the physical drives, and an option to **Blink/Unblink** the selected drive.
7. Verify that no drive is in failed state.

   If any of the drives has failed, refer to the drive FRU procedure in the relevant *Dell EMC VxFlex Ready Node Field Replaceable Unit Guide*.
8. In the navigation pane, select **System**.
9. In the **Inventory** pane, verify that the server drivers and firmware in the **Firmware Inventory** list match the required versions, as published in the VxFlex Ready Node Driver and Firmware Matrix. To access the Driver and Firmware Matrix, go to https://support.emc.com/products/42216.

   If the driver and firmware versions do not match the matrix, you must update them using the DTK - Hardware Update Bootable ISO.

Log in to the KVM console

Use the following procedure to log in to iDRAC from the KVM console.

Before you begin

Ensure that:

- The system environment meets the prerequisites for using the KVM console.
- You know the IP address of the iDRAC port.
- You know the username and password for accessing the iDRAC (default username/password are root/Scaleio123).

Procedure

1. From your Internet browser, go to https://*<iDRAC_IP_address>*.
2. In the **DELL Console Login** window, type the user name and password, and click **Login**.
3. From the dashboard, click **Launch Virtual Console** to start a console session.
4. If a security warning appears, select **Accept**, and then click **Run**.

   If Java is being used, the **Java Console** window opens and provides you with console access to the server. If this is the first time that you are opening a console, additional warning and confirmation prompts may appear. Click **OK** to grant approvals in these prompts.

   The subsequent steps depend on your browser selection and how it is configured. If downloads run automatically, the console window appears. If not, follow the instructions given in the next step to open the console window.
5. Depending on your browser, perform the following steps:

Browser	Steps
Firefox	a. Click **Launch Virtual Console**. The **"What should Firefox do with this file?"** window appears, with the **Open with Java(™) Web Start Launcher** option selected. b. Click **OK**. c. Scroll through the successive pop-up windows, and then click **Run** to launch the console.
Google Chrome	a. Click **Show all downloads**. b. In the **Downloads** window, at the warning message, click **Keep**. c. In the downloads list, click the downloaded file URL. d. At the warning message, click **Keep** or **Keep anyway**. e. Click the `jviewer.jnlp` file. If a security message appears, click **Run**. The management console window is displayed.
Internet Explorer	Click **Open** to run the `jviewer.jnlp` file.

Results

The KVM console is open and ready for use.

Updating the BIOS, firmware and settings

VxFlex Ready Node deployments require specific versions of drivers, BIOS, and firmware that have been qualified by Dell EMC. If the servers do not have the correct versions, you must update them.

A variety of factors can influence a mismatch between the required versions and the versions installed on the servers, such as firmware updates post server shipment, or a FRU replacement with a different firmware version than in the warehouse. For example, if you have replaced the server system board, the FRU's BIOS and iDRAC firmware versions will be different.

You are required to verify that all server drivers, BIOS, and firmware meet the required versions, as published in the VxFlex Ready Node Driver and Firmware Matrix, located at https://support.emc.com/products/42216, before deploying a server in the VxFlex Ready Node environment.

To perform any updates needed to meet VxFlex Ready Node requirements, use the VxFlex Ready Node Hardware Update Bootable ISO ("Hardware ISO"). The Hardware ISO is based on the Dell OpenManage Deployment Toolkit (DTK). The DTK provides a framework of tools necessary for the configuration of VxFlex Ready Node servers. For VxFlex OS, a custom script has been injected, along with specific qualified BIOS/firmware update packages.

The Hardware ISO has been designed to make the firmware update process consistent and simple. You can use it to update firmware, BIOS, and configuration settings two different ways, depending on how many VxFlex Ready Node servers you are updating:

- To update the firmware, BIOS, and configuration settings on a single server, use the iDRAC Virtual KVM.
- If several servers in the VxFlex Ready Node deployment require updates, it is recommended that you use remote RACADM. RACADM allows for the simultaneous update of the versions on multiple VxFlex Ready Node servers, with minimal steps. For instructions, see Update the hardware using remote RACADM on page 76.

For additional information regarding the Hardware ISO, see DTK - Hardware Update Bootable ISO on page 76.

Update the hardware using the iDRAC virtual console

The iDRAC Virtual KVM console and Virtual CDROM features, provided by the iDRAC Enterprise license, eliminate the need for physical access to the VxFlex Ready Node servers. The Hardware ISO can be attached as a remote/Virtual CDROM image and configured to perform hardware updates to the VxFlex Ready Node firmware.

About this task

If you are performing updates on multiple servers, you can optionally use remote RACADM instead. When remote RACADM is used, you do not need to use multiple web browser windows and perform manual keyboard and mouse clicks for each server. To use remote RACADM, seeUpdate the hardware using remote RACADM on page 76.

For more information about the Hardware ISO, see DTK - Hardware Update Bootable ISO on page 76.

Procedure

1. Ensure that the laptop or server is configured with access to the iDRAC network.

   For details, see "Set up the iDRAC IP address and BIOS" in this publication.

2. Download the VxFlex Ready Node Hardware Update Bootable ISO from https://support.emc.com/products/42216 and make it accessible in the network share folder or other appropriate folder.
3. Connect to the iDRAC Virtual KVM console.
4. Attach the Hardware ISO to Virtual Media.
   a. From the **Virtual Media** menu, select **Connect Virtual Media**.
   b. From the **Virtual Media** menu, select **Map CD/DVD**.
   c. Click **Browse**, and select the Hardware ISO.

5. From the **Next Boot** menu, select **Virtual CD/DVD/ISO** and configure the Next Boot option so that the server will boot to the iDRAC Virtual CDROM.
6. Power-cycle the server using the appropriate **Power** menu option, and allow the server to boot to the virtual media.
7. Repeat steps 3 through 6 for each VxFlex Ready Node server that requires BIOS or firmware updates and configuration. All VxFlex Ready Node servers can be updated in parallel.
8. When the updates are completed, refresh the iDRAC browser screen, log in to the iDRAC, and re-launch the virtual console as needed.
9. Wait for the configuration and firmware updates to complete. The server console screen will indicate when the script is complete.

   WARNING Do not reboot the VxFlex Ready Node server while the update process is being performed!

   The iDRAC will be reset several times during the update process. This causes the iDRAC virtual console viewer to close, virtual media to disconnect, and the iDRAC browser window to be unavailable for several minutes during each reset. The hardware update scripts will continue to run from RAM on the server.

   The update script will generate a log indicating whether each configuration and firmware flash is successful.

10. (Optional) Check each VxFlex Ready Node server's log for successful competition:
    a. After the update script completes, press Alt+F2 to access a user console, and then Enter to log in.
    b. Check the log contents for errors:

    ```
    less /bundleapplicationlogs/apply_components.log
    ```

    For more information, see Troubleshooting the Hardware ISO on page 80 .

    c. Press q to exit the log viewer, and then Alt+F1 to access the original console screen.
11. Reboot the VxFlex Ready Node server and allow the update and configuration jobs to complete. Power-cycle the server using the appropriate **Power** menu option.
12. For each VxFlex Ready Node server, after the updates are finalized, clear the iDRAC job queue using the iDRAC GUI:
    a. From your Internet browser, go to https://*<iDRAC_IP_address>*.
    b. In the **DELL Console Login** window, type these credentials:
       - username: root
       - password: *<password>*
    c. Click **Login**.
    d. In the iDRAC GUI navigation pane, select the **Maintenance** tab, and then select the **Job Queue** tab.
    e. Ensure that all jobs have completed successfully. Any job failures may require re-running the bootable ISO, or further troubleshooting.
    f. Select the checkbox for all of items in the **Job Queue** list, and then click **Delete**.
13. Continue with OS installation and configuration, VxFlex OS deployment, and other required tasks.

CHAPTER 3

Installing the Linux operating system

This section describes the procedures for installing a Linux operating system on a VxFlex Ready Node server.

- Linux system requirements 26
- Rebuild the RAID 1 boot device using the replaced M.2 module 26
- Map the Linux ISO file on a VxFlex Ready Node server 27
- Optimize CPU performance on RHEL systems 27
- Optimize CPU performance on SUSE systems 28

Linux system requirements

To deploy on Linux nodes in the VxFlex Ready Node environment, the following prerequisites must be met:

Note: For operating system installation guidelines, refer to the operating system vendor's documentation, taking into account the VxFlex Ready Node prerequisites listed in this section.

- Configure the M.2 device on the BOSS as a RAID 1 device.
- During operating system installation, use Scaleio123 as the password for the user name root or administrator. (Alternately, provide a password that meets the local security criteria.)
- In a 2-Layer installation, Dell EMC supplies a VxFlex Ready Node image ISO.
  Follow the wizard installation steps. The default password is Scaleio123. If required, you can change the password by using the `passwd` command.
- Install the operating system on the BOSS device.

Rebuild the RAID 1 boot device using the replaced M.2 module

The BOSS device consists of two 120/240 GB M.2 expansion cards. If necessary, rebuild the M.2 cards as a RAID 1 device.

Procedure

1. Power on the node.
2. In the iDRAC Virtual KVM console, press F2 immediately after you see the message `F2 = System Setup` during system start up.
3. On the **System Setup Main Menu** screen, select **Device Settings** > **AHCI Controller in Slot 2: BOSS-S1 Configuration Utility**.
4. On the **BOSS-S1 Configuration Utility** screen, select **Physical/Virtual Disk Information** to verify the disk configuration.
5. On the **BOSS-S1 Configuration Utility > Operation Menu** screen, select **Physical Disk Info** as the Device Output Type, and then press Enter.

   The physical disk information should display two expansion card devices configured to the server.
6. Select **Virtual Disk Info** as the Device Output Type, and then press Enter.

   The configuration should show one virtual disk in a degraded RAID state configured on the server.
7. Click **Back**.
8. On the **BOSS-S1 Configuration Utility** screen, select **Rebuild Raid**.

   You can monitor the status of the rebuild in the **Storage** > **Virtual Disks** screen. During the rebuild, the status of the M.2 will be displayed as `Degraded`.

   When the M.2 array is rebuilt, its status will be similar to the following:

   ```
   VD_R1_1     Online
   ```

9. On the **Virtual Disk Information** screen, verify that the settings are configured as follows:
   - **RAID Level**: RAID1
   - **Stripe Size**: 64 K
   - **Virtual Disk Size**: 111/222 GB (The size will correspond to the M.2 device in use.)
   - **Name**: OS

- **Would you like to create this virtual disk**:
  - First-time configuration: Yes
  - M.2 card replacement: No

10. Exit System Setup and boot to the OS.

Results

The **BOSS-S1 Configuration Utility** > **Operation Menu** screen displays one virtual disk with the specified name.

Map the Linux ISO file on a VxFlex Ready Node server

Map the Linux ISO file on a VxFlex Ready Node server.

Before you begin

Ensure that you can access the KVM console.

Procedure

1. Open the KVM console, using the **Launch** link.
2. In the **DELL System Setup** screen, select **Virtual Media** > **Connect Virtual Media**.

   The **Initializing connection** screen is displayed with a connecting message.
   You may have to wait for some time for the connecting message to finish.
3. In the **DELL System Setup** screen, select **Virtual Media** > **Map CD/DVD**.
4. From the **Virtual Media - Map CD/DVD** screen, browse and set up the `ISO` file and then click **Map Device**.
5. To verify the `ISO` file selection, click **Virtual Media** and view the ISO selection in the drop down list.
6. Reset the server:
   a. In the main **DELL System Setup** console window, select **Power** > **Rest System (warm boot)** to display the list of keyboard keys with related functionality.
   b. Refer to the best practices of Dell EMC OS installation: How to Install an Operating System on Dell VxFlex Ready Nodes.

      Ensure that you select the BOSS device as the OS install target drive.

   The server boots using the OS ISO file and finishes the installation.

After you finish

Note: A Red Hat license is required for storage-only configurations with Red Hat OS image downloaded from the VxFlex OS support page.
Customers can choose a Red Hat licensing option listed on this page, or bring their own Red Hat license. Dell EMC is not responsible for enforcing OS licensing.

Optimize CPU performance on RHEL systems

To ensure maximum CPU performance on systems running RHEL, perform the following procedure on the server:

Procedure

1. Open the GRUB template for editing:

```
vim /etc/default/grub
```

2. Find the `GRUB_CMDLINE_LINUX` configuration option and append the following to the line to the kernel parameters:

```
intel_idle.max_cstate=0 processor.max_cstate=1 intel_pstate=disable
```

Example:

```
GRUB_CMDLINE_LINUX="crashkernel=auto rd.lvm.lv=rhel/root rd.lvm.lv=rhel/swap rhgb
intel_idle.max_cstate=0 processor.max_cstate=1 intel_pstate=disable quiet"
```

3. Compile the new GRUB:

```
grub2-mkconfig -o /boot/efi/EFI/redhat/grub.cfg
```

4. Stop and then disable the tuned service:

```
systemctl stop tuned
systemctl disable tuned
```

5. Reboot the server.

Optimize CPU performance on SUSE systems

To ensure maximum CPU performance on systems running SUSE 12.2.2 or higher, perform the following procedure:

Procedure

1. Open the GRUB template for editing:

```
vim /etc/default/grub
```

2. Find the `GRUB_CMDLINE_LINUX` configuration option and append the following to the line:

```
"intel_idle.max_cstate=0 processor.max_cstate=1 intel_pstate=disable"
```

Example:

```
GRUB_CMDLINE_LINUX="intel_idle.max_cstate=0 processor.max_cstate=1 intel_pstate=disable"
```

3. Compile the new GRUB:

```
mkgrub2-mkconfig -o /boot/efi/EFI/sles/grub.cfg
```

4. Reboot the server.

CHAPTER 4

Network architecture and physical connectivity for Linux servers

This section provides networking requirements and connectivity information for VxFlex Ready Node Linux servers.

- Networking connectivity architecture and cabling best practice for Linux deployments 30
- IP addresses for R640 and R740xd servers 30
- IP addresses for R840 servers 32

Networking connectivity architecture and cabling best practice for Linux deployments

The following information describes connectivity architecture, cabling best practice information, and cable connection examples from typical VxFlex OS configurations to help you plan your network.

Note: If you are not familiar with VxFlex Ready Node system architecture, refer to the "Architecture" chapter in the *Getting to Know VxFlex OS Guide*.

The best way to use VLANs with VxFlex OS is to configure the switch port to Mode Access, and assign it to one VLAN. A management switch that has only one use is recommended. Usually, this implies one IP subnet, but the switch can contain several subnets on a single VLAN.

Note: AMS (R640 and R740xd systems only) uses APIPA IP addresses (169.254.x.x/16) for broadcast and unicast, and this requires special configuration during deployment.

It is mandatory to set the management VLAN as either a Layer 2 VLAN without an IP address, or as VLAN 0 (meaning, no VLAN). This facilitates a seamless onsite "Add Node" operation.

IP addresses for R640 and R740xd servers

Prepare IP addresses in your network for the VxFlex Ready Node servers in Linux-based environments, based on the following calculations:

Note: In addition to the networking requirements below, ensure that you have prepared the items described in Additional equipment and network resource requirements on page 15 earlier in this guide.

Table 5 VxFlex OS management IP network

Item	Description
N	Number of nodes
IP address pools	The pools of IP addresses used for static allocation for the following groups: 1. Node_MGMT_IP = Management IP addresses of the VxFlex Ready Node servers 2. BMC_MGMT_IP = iDRAC IP address for each node
The formula to calculate IP address subnet pool or subnet needs is: N*Node_MGMT_IP+N*BMC_MGMT_IP + Optional (AMS server IP address)	

Note: Each pool can belong to a single subnet or to several subnets, as long as the AMS server has routable networking access to the Management IP address.

The VxFlex OS data IP network requires two separate subnets.

Table 6 VxFlex OS data IP network for Subnet #1

Item	Description	Comments
N	Number of nodes	

Table 6 VxFlex OS data IP network for Subnet #1 (continued)

Item	Description	Comments
IP address pool for Subnet #1	The pools of IP addresses used for static allocation for the following groups: 1. Node_DATA1_IP = internal (interconnect) IP addresses of the VxFlex Ready Node servers 2. MDM_Cluster_Virtual_IP_DATA1 = the Virtual IP addresses of the MDM cluster in the "Data1" network	For clarity, the first subnet is referred to as "Data1"
The formula to calculate the subnet pool or subnet needs is: N*Node_DATA1_IP+MDM_Cluster_Virtual_IP_DATA1		

Table 7 VxFlex OS data IP network for Subnet #2

Item	Description	Comments
N	Number of nodes	
IP address pool for Subnet #2	The pools of IP addresses used for static allocation for the following groups: 1. Node_DATA2_IP = internal (interconnect) IP addresses of the VxFlex Ready Node servers 2. MDM_Cluster_Virtual_IP_DATA2 = the Virtual IP addresses of the MDM cluster in the "Data2" network	For clarity, the second subnet is referred to as "Data2"
The formula to calculate the subnet pool or subnet needs is: N*Node_DATA2_IP+MDM_Cluster_Virtual_IP_DATA2		

Table 8 VxFlex OS IP network for 2-layer configuration (Data 3 + Data 4)

Item	Description	Comments
N	Number of nodes	
IP address pool for Data 3	The pool of IP addresses used for static allocation for the following group: • Node_DATA3_IP = internal (interconnect) IP addresses of the VxFlex Ready Node servers	For clarity, the second subnet is referred to as "Data3"
The formula to calculate the subnet pool or subnet needs is: N*Node_DATA3_IP		

Item	Description	Comments
N	Number of nodes	
IP address pool for Data 4	The pool of IP addresses used for static allocation for the following group: • Node_DATA4_IP = internal (interconnect) IP addresses of the VxFlex Ready Node servers	For clarity, the second subnet is referred to as "Data4"
The formula to calculate the subnet pool or subnet needs is: N*Node_DATA4_IP		

IP addresses for R840 servers

Prepare IP addresses in your network for the VxFlex Ready Node servers in Linux-based environments based on the following calculations:

Note: In addition to the networking requirements below, ensure that you have prepared the items described in Additional equipment and network resource requirements on page 15 earlier in this guide.

Table 9 VxFlex OS management IP network

Item	Description
N	Number of nodes
IP address pools	The pools of IP addresses used for static allocation for the following groups: 1. Node_MGMT_IP = Management IP addresses of the VxFlex Ready Node servers 2. BMC_MGMT_IP = iDRAC IP address for each node
The formula to calculate IP address subnet pool or subnet needs is: N*Node_MGMT_IP+N*BMC_MGMT_IP	

Note: Each pool can belong to a single subnet or to several subnets.

The VxFlex OS data IP network requires two separate subnets.

Table 10 VxFlex OS data IP network for Subnet #1

Item	Description	Comments
N	Number of nodes	
IP address pool for Subnet #1	The pools of IP addresses used for static allocation for the following groups: 1. Node_DATA1_IP = internal (interconnect) IP addresses of the VxFlex Ready Node servers 2. MDM_Cluster_Virtual_IP_DATA1 = the Virtual IP addresses of the MDM cluster in the "Data1" network	For clarity, the first subnet is referred to as "Data1"
The formula to calculate the subnet pool or subnet needs is: N*Node_DATA1_IP+MDM_Cluster_Virtual_IP_DATA1		

Table 11 VxFlex OS data IP network for Subnet #2

Item	Description	Comments
N	Number of nodes	
IP address pool for Subnet #2	The pools of IP addresses used for static allocation for the following groups: 1. Node_DATA2_IP = internal (interconnect) IP addresses of the VxFlex Ready Node servers 2. MDM_Cluster_Virtual_IP_DATA2 = the Virtual IP addresses of the MDM cluster in the "Data2" network	For clarity, the second subnet is referred to as "Data2"
The formula to calculate the subnet pool or subnet needs is: N*Node_DATA2_IP+MDM_Cluster_Virtual_IP_DATA2		

Table 12 VxFlex OS IP network for 2-layer configuration (Data 3 + Data 4)

Item	Description	Comments
N	Number of nodes	
IP address pool for Data 3	The pool of IP addresses used for static allocation for the following group: • Node_DATA3_IP = internal (interconnect) IP addresses of the VxFlex Ready Node servers	For clarity, the second subnet is referred to as "Data3"
The formula to calculate the subnet pool or subnet needs is: N*Node_DATA3_IP		

Item	Description	Comments
N	Number of nodes	
IP address pool for Data 4	The pool of IP addresses used for static allocation for the following group: • Node_DATA4_IP = internal (interconnect) IP addresses of the VxFlex Ready Node servers	For clarity, the second subnet is referred to as "Data4"
The formula to calculate the subnet pool or subnet needs is: N*Node_DATA4_IP		

CHAPTER 5

Configuring network ports on Linux servers

This section describes how to configure the ports on a VxFlex Ready Node Linux server.

- Linux ports overview 36
- Configure network ports on R640 or R740xd servers 36
- Configure network ports on R840 servers 45

Linux ports overview

Typically, VxFlex OS uses two network card data connections.

- If a CX4 card configured as an rNDC is installed on the node and VxFlex OS uses the card for the data connection, you are required to collect the BUS address and interface name for the rNDC. In this case, you must configure an additional data connection for a CX4 card in a PCIe slot and are required to collect the BUS address and interface name for that card, as well.
- If VxFlex OS uses two PCIe slot-based CX4 or CX5 cards for data collection, collect the BUS address and interface name for these two cards

Configure network ports on R640 or R740xd servers

Use the following information to configure network ports on R640 or R740xd servers.

Determine the 2U1N node type

For a 2U1N server, determine whether it is a converged or storage-only node.

Before you begin

Ensure that you have:

- The IP address of the iDRAC port
- The user name and password for accessing the iDRAC

About this task

The 2U1N server architecture can be converged (2CPU) or storage-only (1 CPU). You must determine the architecture to enable proper port configuration.

Procedure

1. From your Internet browser, go to http://*<iDRAC_IP_address>*.
2. In the **DELL Console Login** window, type the user name and password, and click **Login**.
3. In the navigation pane, select **System** > **Inventory** > **Hardware inventory**, and then select the **CPU** node.
4. Count the number of lines in the **Processor Information** table.

   If the node contains:

   - Two CPUs, it is for converged architecture
   - One CPU, it is for storage-only

Results

The 2U1N node architecture is determined.

R640 and R740xd server NIC and iDRAC port locations

NIC and iDRAC port designations for VxFlex Ready Node R640 and R740xd servers for various server configurations are described in this topic. NICs may be located onboard, or in PCI slots, depending on the server configuration.

The following figure shows possible connector locations on the rear panel of the servers. Actual NIC locations depend on server configurations. Some models do not contain NICs in all possible slots. Items marked 1 to 8 represent PCI slots.

Figure 1 R640 and R740xd PCI slots, integrated NICs and BMC port locations

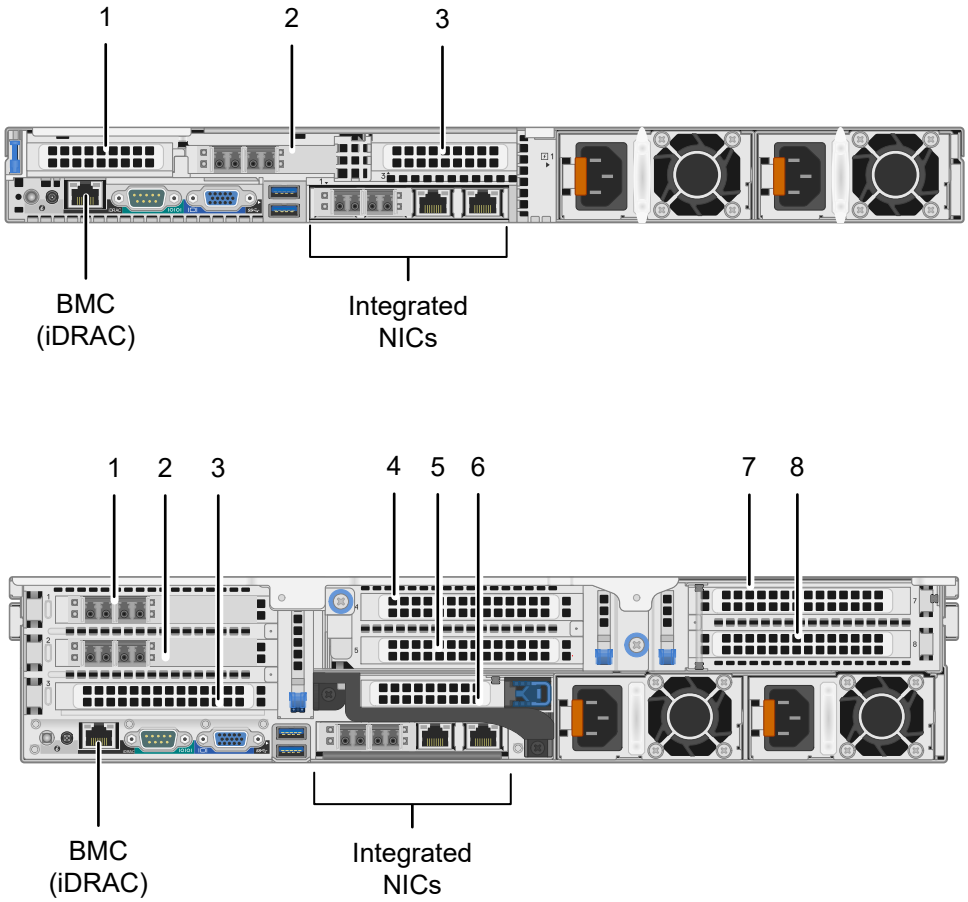

Data cables are connected to two switches for high availability, via the nodes' 10 GbE, 25 GbE, or 100 GbE ports. Management interfaces are connected to a switch on a separate management network using onboard and iDRAC ports.

The iDRAC IP address can reside on a different subnet and/or VLAN, and the only requirement from the iDRAC side is that the AMS (when used) must have a routable IP Path to the iDRAC IP address on the node.

The default configuration used in simple mode deployment follows the following rules of thumb:

- If there is a 1G NIC onboard, the left one of the pair is always used for the management network.
- If there are two PCI cards or more, the lowest slot number is used for allocation of data ports. For example, if slots 2, 3 and 5 are populated with NICs, AMS uses slots 2 and 3, and the left port on each PCI card is used for data.

  Note: You can always change the default using advanced deployment mode in AMS.
- If there is only one NIC, the left 10G onboard port is used for the data network.
- If there are only 2*25GbE onboard ports, the left port is used for both the management and data network, and the right port for is used for data.

  Note: The application\client network will have to run on those same two ports.
- If there are only 4*10G onboard ports:
  - The left port is used for management
  - The left port, and the third port from the left are used for VxFlex OS data

- The second port from the left, and the right port are used for application\client traffic

VxFlex Ready Node R640 port designations - Linux

For single-node VxFlex Ready Node servers running RHEL or SLES, connect the cables, and configure the ports as shown in the configuration tables according to the server type.

The following additional information should be taken into account:

- The logical bus address is dynamic and will be allocated by the operating system. Use the slot name to correlate to between the bus address and port name of the interface in the operating system.
- In storage-only configurations, the customer ports can be used as additional data ports.

Table 13 R640 single-node server configuration

Description	iDRAC Monitoring	VxFlex OS Data1	VxFlex OS Client Network Port1	VxFlex OS Client Network Port2	VxFlex OS Data2	VxFlex OS Management Network	Not in Use
Interface name (RHEL7)	N/A	emXX	emXX	pXpX	pXpX	emXX	emXX
Interface name (SLES)	N/A	emXX	emXX	pXpX	pXpX	emXX	emXX
Slot name	N/A	Integrated NIC1	Integrated NIC2	PCIe Slot3 Port 1	PCIe Slot3 Port 0	Integrated NIC 3	Integrated NIC 4
Logical bus address	N/A	Z:0	Z:1	Y:1	Y:0	X:0	X:1
Physical port on node	iDRAC	Onboard 10 GbE Left	Onboard 10 GbE Right	PCI 10/25 GbE Right	PCI 10/25 GbE Left	Onboard 1 GbE Left	Onboard 1 GbE Right
Speed	1 GbE	10 GbE	10 GbE	10/25 GbE	10/25 GbE	1 GbE	1 GbE
Physical switch	Mgmt Switch	Data Switch1	Data Switch1	Data Switch2	Data Switch2	Mgmt Switch	Not in use

Note: When using the onboard 10 GbE left and right ports (x710 Intel NICs) in the above configuration, contact Customer Support for guidance.

Note: You can also use a second NIC in Slot 2 for a 4*25 GbE option.

Note: Two or more PCIe slot-based NICs are supported.

VxFlex Ready Node R740xd port designations - Linux

For single-node VxFlex Ready Node servers running RHEL or SLES, connect the cables, and configure the ports as shown in the configuration tables according to the server type.

The following additional information should be taken into account:

- The logical bus address is dynamic and will be allocated by the operating system. Use the slot name to correlate to between the bus address and port name of the interface in the operating system.
- In storage-only configurations, the customer ports can be used as additional data ports.

Note: When using the onboard 10 GbE left and right ports (x710 Intel NICs) in the above configuration, contact Customer Support for guidance.

Note: You can also use a second NIC in Slot 2 for a 4*25 GbE option.

Table 14 R740xd single-node server configuration for 10/25GbE, 2CPU, SFP+/SFP28

Description	iDRAC Monitoring	VxFlex OS Data1	VxFlex OS Client Network Port1	VxFlex OS Client Network Port2	VxFlex OS Data2	VxFlex OS Management Network	Not in Use
Interface name (RHEL 7)	N/A	emXX	emXX	pXpX	pXpX	emXX	emXX
Interface name (SLES)	N/A	emXX	emXX	pXpX	pXpX	emXX	emXX
Slot name	N/A	Free PCI slot, left port	Free PCIe slot, right port	Free PCIe slot, right port	Free PCIe slot, left port	Integrated NIC 3	Integrated NIC 4
Logical bus address	N/A	Z:0	Z:1	Y:1	Y:0	X:0	X:1
Physical port on node	iDRAC	PCI 10/25 GbE Left	PCI 10/25 GbE Right	PCI 10/25 GbE Right	PCI 10/25 GbE Left	Onboard 1 GbE Left	Onboard 1G Right
Speed	1GbE	10/25 GbE	10/25 GbE	10/25 GbE	10/25 GbE	1 GbE	1 GbE
Physical Switch	Mgmt Switch	Data Switch1	Data Switch1	Data Switch2	Data Switch2	Mgmt Switch	Not in use

Note: Two or more PCIe slot-based NICs are supported.

Correlate the PCI slot locations and interface names on R640/R740xd servers

Gather and correlate PCI slot locations and interface names in order to build a table with the relevant interface names for the Linux nodes.

About this task

RHEL 7.x does not use the `70-persistent-net.rules` file for Ethernet persistency. Instead, in RHEL 7.x nodes, the interface names are in the following format: enXXX (for example: enp6s0f0 or eno3).
Other nodes (such as SLES) interface names may be in the following format: ethX (for example: eth0 or eth4).

Note: In the following procedure, the commands are being run on a RHEL 7.x node, so the example outputs include interface names in the format enXXX.

Procedure

1. Find the correlation between physical slot names and interface names so that you can discover the correlation between PCI slot locations and interface names in the operating system:

   Note: For each of the following commands, write down the output and set it aside for use later in this task:

   a. Correlate the physical PCIe slot to a logical bus address:

   ```
   dmidecode -t slot |egrep "Designation|Current Usage|Bus Address"
   ```

   This command displays the PCI location of each PCIe card. In the following sample output, the logical bus address in the last three lines is "0000:05:00.0," which is shown as `Current Usage: In Use` and correlates to physical PCIe card slot location "PCIe Slot 6."

In this case, the 10GB PCI NIC is presented on slot 6, and the logical bus address of its ports is 5:0 and 5:1. (See the image in substep c, in which the correlation between PCIe slot and the logical bus address is highlighted in yellow.)

```
        Designation: PCIe Slot 1
        Current Usage: Available
        Designation: PCIe Slot 2
        Current Usage: Available
        Designation: PCIe Slot 3
        Current Usage: Available
        Designation: PCIe Slot 4
        Current Usage: Available
        Designation: PCIe Slot 5
        Current Usage: Available
        Designation: PCIe Slot 6
        Current Usage: In Use
        Bus Address: 0000:05:00.0
```

b. Correlate the onboard device name (NIC1 to NIC4) to a logical bus address:

```
dmidecode -t baseboard | sed -n -e '/NIC/,/Bus Address/p' | egrep "NIC|Bus Address"
```

This command displays the onboard location of each NIC port, In the following sample output, the logical bus address in the first two lines is "0000:01:00.0" which correlates to physical NIC `Integrated NIC 1`.

```
        Reference Designation: Integrated NIC 1
        Bus Address: 0000:01:00.0
        Reference Designation: Integrated NIC 2
        Bus Address: 0000:01:00.1
        Reference Designation: Integrated NIC 3
        Bus Address: 0000:09:00.0
        Reference Designation: Integrated NIC 4
        Bus Address: 0000:09:00.1
```

c. Correlate the list of OS interface names to a logical bus address:

```
ip a | grep ": " | awk '{print $2}' | tr -d ":" | grep -v ^"lo"$ | xargs -I '{}' sh -
c 'echo {}; ethtool -i {} | grep bus-info'
```

This command displays the Ethernet network names (enXXX) of the interfaces and the correlating logical bus address (without the loopback). For example:

```
        em1
        bus-info: 0000:19:00.0
        p1p1
        bus-info: 0000:3b:00.0
        em2
        bus-info: 0000:19:00.1
        em3
        bus-info: 0000:19:00.2
        em4
        bus-info: 0000:19:00.3
        p1p2
        bus-info: 0000:3b:00.1
```

Figure 2 Correlating physical slot and Ethernet network names

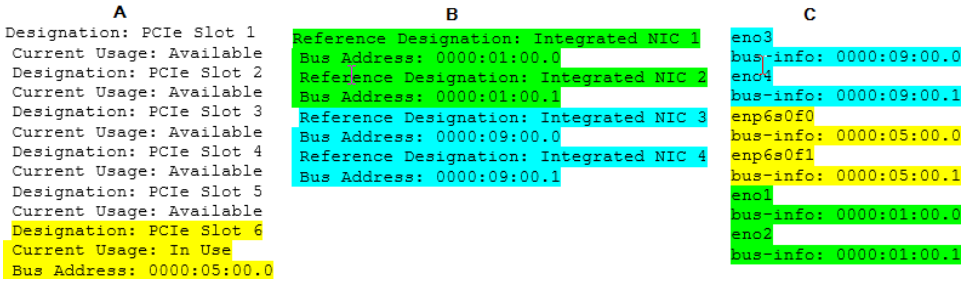

2. Create a table for your own use, similar to the example table shown below:
   - Base the table on the RHEL port definitions described for the relevant server model (R640 or R740xd) and type. Refer to the server port designation topic for the server model you are configuring.
   - In the **Port** row, record the enXXX names that correspond with the PCIe slots.
   - Map each line in turn from the output in substep a to the corresponding line in the output from substep b.

   Note: In 1U1N and 2U1N platforms, the 1 GB port on the right is used for management, and the left 10 GB ports on both NICs are used for data.

VxFlex OS role	...	Optional VxFlex OS Data2 2nd port
Port		enp6s0f0
Slot name		PCIe Slot 6 Port 0
Logical bus address		0000:05:00.0
Physical port on node		PCI 10G Left
Speed		0 GB

Configure the ports on RHEL 7.x nodes - R640/R740xd servers

Configure ports on Linux-based (RHEL 7.x-type) nodes according to the following procedure.

About this task

As part of this procedure, you will assign the traffic management and two data NICs the following names:

- Traffic management NIC: `ifcfg-sio_mgmt`
- VxFlex OS Data NIC 1: `ifcfg-sio_d_1`
- VxFlex OS Data NIC 2: `ifcfg-sio_d_2`

Procedure

1. Navigate to the `/etc/sysconfig/network-scripts/` folder:

```
cd /etc/sysconfig/network-scripts/
```

2. Create three files that correspond to the names you are assigning the NICS (`ifcfg-sio_mgmt`, `ifcfg-sio_d_1`, and `ifcfg-sio_d_2`):

```
cp /etc/sysconfig/network-scripts/ifcfg-eth0 /etc/sysconfig/network-scripts/ifcfg-
sio_mgmt
cp /etc/sysconfig/network-scripts/ifcfg-sio_mgmt /etc/sysconfig/network-scripts/ifcfg-
sio_d_1
cp /etc/sysconfig/network-scripts/ifcfg-sio_mgmt /etc/sysconfig/network-scripts/ifcfg-
sio_d_2
```

3. Configure each of the files you created with the correct management or data IP addresses.

   a. For the `ifcfg-sio_mgmt` file, run:

```
echo DEVICE= sio_mgmt > /etc/sysconfig/network-scripts/ifcfg-sio_mgmt
    echo STARTMODE=onboot >> /etc/sysconfig/network-scripts/ifcfg-sio_mgmt
    echo USERCONTROL=no >> /etc/sysconfig/network-scripts/ifcfg-sio_mgmt
    echo BOOTPROTO=static >> /etc/sysconfig/network-scripts/ifcfg-sio_mgmt
    echo NETMASK=X.X.X.X >> /etc/sysconfig/network-scripts/ifcfg-sio_mgmt
    echo IPADDR=X.X.X.X >> /etc/sysconfig/network-scripts/ifcfg-sio_mgmt
    echo 1 >/sys/bus/pci/rescan
```

   b. For the `ifcfg-sio_d_1` file, run:

```
echo DEVICE= sio_d_2 > /etc/sysconfig/network-scripts/ifcfg-sio_d_1
echo STARTMODE=onboot >> /etc/sysconfig/network-scripts/ifcfg-sio_d_1
echo USERCONTROL=no >> /etc/sysconfig/network-scripts/ifcfg-sio_d_1
echo BOOTPROTO=static >> /etc/sysconfig/network-scripts/ifcfg-sio_d_1
echo NETMASK=X.X.X.X >> /etc/sysconfig/network-scripts/ifcfg-sio_d_1
echo IPADDR=X.X.X.X >> /etc/sysconfig/network-scripts/ifcfg-sio_d_1
echo 1 >/sys/bus/pci/rescan
```

   c. For the `ifcfg-sio_d_2` file, run:

```
echo DEVICE= sio_d_2 > /etc/sysconfig/network-scripts/ifcfg-sio_d_2
echo STARTMODE=onboot >> /etc/sysconfig/network-scripts/ifcfg-sio_d_2
echo USERCONTROL=no >> /etc/sysconfig/network-scripts/ifcfg-sio_d_2
echo BOOTPROTO=static >> /etc/sysconfig/network-scripts/ifcfg-sio_d_2
echo NETMASK=X.X.X.X >> /etc/sysconfig/network-scripts/ifcfg-sio_d_2
echo IPADDR=X.X.X.X >> /etc/sysconfig/network-scripts/ifcfg-sio_d_2
echo 1 >/sys/bus/pci/rescan
```

4. Configure the gateway:

```
echo “NETWORKING=yes” >> /etc/sysconfig/network
echo “GATEWAY=XXX.XXX.XXX.XXX” >> /etc/sysconfig/network
```

   where *XXX.XXX.XXX.XXX* is the gateway IP address.

5. Create the `/etc/udev/rules.d/70-persistent-ipoib.rules` file:

```
touch /etc/udev/rules.d/70-persistent-ipoib.rules
```

6. Change the permissions of the file you just created to allow the root user access:

```
chmod a+x /etc/udev/rules.d/70-persistent-ipoib.rules
```

7. Find the MAC address for each NIC:

```
ethtool -p <NIC_name>
```

Output similar to the following should appear:

```
ethtool -P ens192
Permanent address: 00:50:56:a7:b0:19
```

8. For each NIC, edit the `70-persistent-ipoib.rules` file according to the information in the output in the previous step:

```
    echo "ACTION=="add", SUBSYSTEM=="net", DRIVERS=="?*", ATTR{type}=="32",
ATTR{address}=="<address>", NAME="<name>"" >> /etc/udev/rules.d/70-persistent-ipoib.rules
```

9. Reboot to the host.
10. After the network comes up, log in using SSH to the management IP address, and ping all data IP addresses of another node to verify that you have the correct connectivity on each node.

Results

Network configuration is complete.

Configure the ports on SLES nodes - R640/R740xd servers

Configure the ports on Linux-based (SLES-type) nodes according to the following procedure:

Procedure

1. Configure the `/etc/udev/rules.d/70-persistent-net.rules` file to match the following:
   - The 1 GB port is em3 and is used for management.
   - The 10 GB data ports are em1 (Data 1) and p1p1 (Data 2) .
2. Use `vi` to edit the `/etc/udev/rules.d/70-persistent-net.rules` file and replace all the Eth names with the correct allocation:
   - The 1 GB port is em3 and em4.
   - The 10 GB ports are em1, em2, p2p1, p2p2.

   The following example shows how the `/etc/udev/rules.d/70-persistent-net.rules` file will look after the eth ports are configured:

```
# PCI device 0x8086:0x10fb (ixgbe)
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="2c:60:0c:c0:9b:a1", ATTR{dev_id}=="0x0"
=="eth*", NAME="eth1"

# PCI device 0x8086:0x10fb (ixgbe)
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="90:e2:ba:0b:9b:bd", ATTR{dev_id}=="0x0"
=="eth*", NAME="eth3"

# PCI device 0x8086:0x1521 (igb)
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="2c:60:0c:ca:ff:49", ATTR{dev_id}=="0x0"
=="eth*", NAME="eth5"

# PCI device 0x8086:0x10fb (ixgbe)
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="90:e2:ba:8b:9b:bc", ATTR{dev_id}=="0x0"
=="eth*", NAME="eth2"

# PCI device 0x8086:0x1521 (igb)
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="2c:60:0c:ca:ff:40", ATTR{dev_id}=="0x0"
=="eth*", NAME="eth4"

# PCI device 0x8086:0x10fb (ixgbe)
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="2c:60:0c:c0:9b:a3", ATTR{dev_id}=="0x0"
=="eth*", NAME="eth0"
```

3. After updating the `70-persistent-net.rules` files, reboot the node using the `reboot` command.
4. After the node comes up, log in again to the console.
5. Create the `ifcfg-emX` files for all interfaces.

   Note: A good way to create a file is using the `cp` and `echo` commands:

   Example:
   ```
   cp /etc/sysconfig/network/ifcfg-em3 /etc/sysconfig/network/ifcfg-em4;
   echo "" > /etc/sysconfig/network/ifcfg-em4;
   ```
6. Create the following files:
   - Eth configuration files for 1 G interface as ifcfg-em3
   - Eth Configuration files for 10 G interface as (first 10 G NIC) ifcfg-em1 & ifcfg-em2
   - Eth Configuration files for 10 G interface as (second 10 G NIC) ifcfg-p2p1 & ifcfg-p2p2

   Refer to the following commands as examples for how to populate each file. Replace each use of "X" with the relevant interface, subnet or IP address.
   Example:

   ```
   cp /etc/sysconfig/network/ifcfg-em3 /etc/sysconfig/network/ifcfg-em4
   echo DEVICE=em3 > /etc/sysconfig/network/ifcfg-emX
   echo STARTMODE=onboot >> /etc/sysconfig/network/ifcfg-emX
   echo USERCONTROL=no >> /etc/sysconfig/network/ifcfg-emX
   echo BOOTPROTO=static >> /etc/sysconfig/network/ifcfg-emX
   echo NETMASK=X.X.X.X >> /etc/sysconfig/network/ifcfg-emX
   echo IPADDR=X.X.X.X >> /etc/sysconfig/network/ifcfg-emX
   echo 1 >/sys/bus/pci/rescan
   sleep 2
   ifup emX
   ```
7. Configure the default gateway:

   ```
   echo "default <XXX.XXX.XXX.XXX> - em3 " >/etc/sysconfig/network/ifroute-emX
   ```

   For example:

   ```
   echo "default - em3" > /etc/sysconfig/network/ifroute-em3
   ```

Note: Only the one default gateway is required. Most likely this will be the management network interface. Create this file by using the following command:

```
touch /etc/sysconfig/network/ifroute-emX
```

8. Perform reboot to the host.
9. After the host is up, log in to the management IP address using SSH, and ping all data IP addresses of another node to make sure that you have the correct connectivity on each node.

Results

The network configuration on the SLES node is complete. Next perform disk configuration procedures.

Configure network ports on R840 servers

Use the following information to configure network ports on R840 servers.

R840 server NIC and iDRAC port locations

NIC and iDRAC port designations for VxFlex Ready Node R840 servers for various server configurations are described in this topic. NICs may be located onboard, or in PCI slots, depending on the server configuration.

The following figure shows possible connector locations on the rear panel of the servers. Actual NIC locations depend on server configurations. Some models do not contain NICs in all possible slots. Items marked 1 to 3 represent PCI slots.

Note: When a GPU is installed in a VxFlex Ready Node R840 server, not all the slots are available.

Figure 3 R840 PCI slots, integrated NICs, and iDRAC port locations

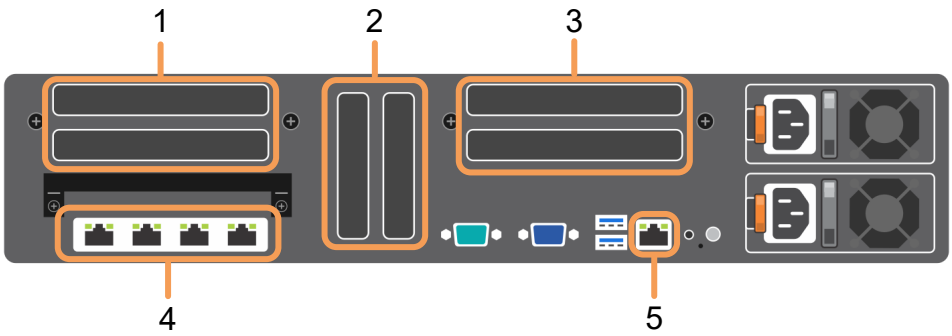

1. 2x FH PCIe slots
2. 2x LP PCIe slots
3. 2x FH PCIe slots
4. Integrated NIC/rNDC ports
5. iDRAC

Data cables are connected to two switches for high availability, via the nodes' 25 GbE ports (not supported for integrated NICs). Management interfaces are connected to a switch on a separate management network using onboard and iDRAC ports.

Additionally, integrated NICs can be used for data connections for up to 25 GbE support.

The iDRAC IP address can reside on a different subnet and/or VLAN.

The default configuration used in simple mode deployment follows the following rules of thumb:

- If there is a 1G NIC onboard, the left one of the pair is always used for the management network.
- If there is only one NIC, the left 10G onboard port is used for the data network.
- If there are only 2*25GbE onboard ports, the left port is used for both the management and data network, and the right port for is used for data.
  Note: The application/client network must run on these same two ports.
- If there are only 4*10G onboard ports:
  - The left port is used for management
  - The left port, and the third port from the left are used for VxFlex OS data
  - The second port from the left, and the right port are used for application\client traffic

VxFlex Ready Node R840 Linux server port designations

For single-node VxFlex Ready Node servers running a Linux distribution, connect the cables, and configure the ports as shown in the configuration tables according to the server type.

The following additional information should be taken into account:

- The logical bus address is dynamic and will be allocated by the operating system. Use the slot name to correlate to between the bus address and port name of the interface in the operating system.
- In storage-only configurations, the customer ports can be used as additional data ports.

Table 15 R840 server configuration for 2x10GbE+2GbE rNDC, 2x10/25GbE PCIe slot 3, 2x10/25GbE PCIe slot 5

Description	iDRAC Monitoring	VxFlex OS Data1	VxFlex OS Client Network Port1	VxFlex OS Client Network Port2	VxFlex OS Data2	VxFlex OS Management Network	Not in Use
Interface name (RHEL 7)	N/A	p3p1	p3p2	p5p2	p5p1	em3	em1,2,4
Interface name (SLES)	N/A	p3p1	p3p2	p5p2	p5p1	em3	emXX
Slot name	N/A	PCIe Slot3 Port 1	PCIe Slot3 Port 2	PCIe Slot5 Port 2	PCIe Slot5 Port 1	Integrated NIC 3	Integrated NIC 1,2,4
Logical bus address	N/A	33:00.0	33:00.1	5b:00.1	5b:00.0	01:00.1	01:00.X
Physical port on node	iDRAC	PCIe 10/25 GbE Bottom	PCIe 10/25 GbE Top	PCIe 10/25 GbE Right	PCI 10/25 GbE Left	Onboard 1 GbE Left	Onboard
Speed	1GbE	10/25 GbE	10/25 GbE	10/25 GbE	10/25 GbE	1 GbE	1 GbE
Physical Switch	Mgmt Switch	Data Switch1	Data Switch1	Data Switch2	Data Switch2	Mgmt Switch	Not in use

Note: Two or more PCIe slot-based NICs are supported.

Correlate the PCIe slot locations and interface names - R840 servers

Gather and correlate PCIe slot locations and interface names in order to build a table with the relevant interface names for the Linux nodes.

About this task

RHEL 7.x does not use the `70-persistent-net.rules` file for Ethernet persistency. Instead, in RHEL 7.x nodes, the interface names are in the following format: enXXX (for example: enp6s0f0 or eno3).

Other nodes (such as SLES) interface names may be in the following format: ethX (for example: eth0 or eth4).

Note: In the following procedure, the commands are being run on a RHEL 7.x node, so the example outputs include interface names in the format enXXX.

Procedure

1. Find the correlation between physical slot names and interface names so that you can discover the correlation between PCIe slot locations and interface names in the operating system:

   Note: For each of the following commands, record the output and set it aside for use later in this task:

   a. Correlate the physical PCIe slot to a logical bus address:

   ```
   dmidecode -t slot |egrep "Designation|Current Usage|Bus Address"
   ```

   This command displays the PCIe location of each PCIe card. In the following example output, the logical bus address in the last three lines is "0000:05:00.0," which is shown as `Current Usage: In Use` and correlates to physical PCIe card slot location "PCIe Slot 6."

   In this case, the 10GB PCIe NIC is presented on slot 6, and the logical bus address of its ports is 5:0 and 5:1. (See the figure in this topic, "Correlating physical slot and Ethernet network names", in which the correlation between PCIe slot and the logical bus address is highlighted in yellow.)

   ```
           Designation: PCIe Slot 1
           Current Usage: Available
           Designation: PCIe Slot 2
           Current Usage: Available
           Designation: PCIe Slot 3
           Current Usage: Available
           Designation: PCIe Slot 4
           Current Usage: Available
           Designation: PCIe Slot 5
           Current Usage: Available
           Designation: PCIe Slot 6
           Current Usage: In Use
           Bus Address: 0000:05:00.0
   ```

   b. Correlate the onboard device name (NIC1 to NIC4) to a logical bus address:

   ```
   dmidecode -t baseboard | sed -n -e '/NIC/,/Bus Address/p' | egrep "NIC|Bus Address"
   ```

   This command displays the onboard location of each network port. In the following example output, the logical bus address in the first two lines is "0000:01:00.0," which correlates to physical NIC `Integrated NIC 1`.

   ```
           Reference Designation: Integrated NIC 1
           Bus Address: 0000:01:00.0
           Reference Designation: Integrated NIC 2
   ```

```
        Bus Address: 0000:01:00.1
        Reference Designation: Integrated NIC 3
        Bus Address: 0000:09:00.0
        Reference Designation: Integrated NIC 4
        Bus Address: 0000:09:00.1
```

c. Correlate the list of OS/Ethernet interface names to a logical bus address:

```
ip a | grep ": " | awk '{print $2}' | tr -d ":" | grep -v ^"lo"$ | xargs -I '{}' sh -
c 'echo {}; ethtool -i {} | grep bus-info'
```

This command displays the Ethernet network names (enXXX) of the interfaces and the correlating logical bus address (without the loopback). For example:

```
        em1
        bus-info: 0000:19:00.0
        p1p1
        bus-info: 0000:3b:00.0
        em2
        bus-info: 0000:19:00.1
        em3
        bus-info: 0000:19:00.2
        em4
        bus-info: 0000:19:00.3
        p1p2
        bus-info: 0000:3b:00.1
```

Figure 4 Correlating physical slot and Ethernet network names

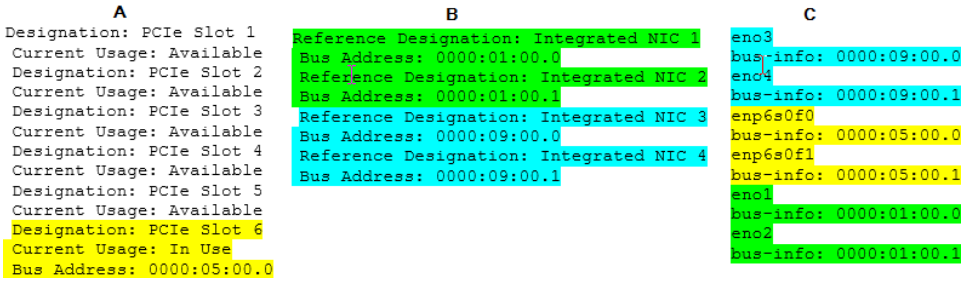

2. Create a table for your own use, similar to the example table shown below:
   - Base the table on the RHEL port definitions described for the relevant server type in the R840 port designation topic.
   - In the **Port** row, record the enXXX names that correspond with the PCIe slots.
   - Map each line in turn from the output in substep a to the corresponding line in the output from substep b.

VxFlex OS role	...	Optional VxFlex OS Data2 2nd port
Port		enp6s0f0
Slot name		PCIe Slot 6 Port 0
Logical bus address		0000:05:00.0
Physical port on node		PCI 10G Left
Speed		0 GB

Configure the ports on RHEL 7.x nodes - R840 servers

Configure ports on Linux-based (RHEL 7.x type) nodes according to the following procedure.

About this task

As part of this procedure, you will assign the traffic management and two data NICs the following names:

- Traffic management NIC: `ifcfg-sio_mgmt`
- VxFlex OS Data NIC 1: `ifcfg-sio_d_1`
- VxFlex OS Data NIC 2: `ifcfg-sio_d_2`

Procedure

1. Navigate to the `/etc/sysconfig/network-scripts/` folder:

```
cd /etc/sysconfig/network-scripts/
```

2. Create three files that correspond to the names you are assigning the NICS (`ifcfg-sio_mgmt`, `ifcfg-sio_d_1`, and `ifcfg-sio_d_2`):

```
cp /etc/sysconfig/network-scripts/ifcfg-eth0 /etc/sysconfig/network-scripts/ifcfg-
sio_mgmt
cp /etc/sysconfig/network-scripts/ifcfg-sio_mgmt /etc/sysconfig/network-scripts/ifcfg-
sio_d_1
cp /etc/sysconfig/network-scripts/ifcfg-sio_mgmt /etc/sysconfig/network-scripts/ifcfg-
sio_d_2
```

3. Configure each of the files you created with the correct management or data IP addresses.
   a. For the `ifcfg-sio_mgmt` file, run:

```
echo DEVICE= sio_mgmt > /etc/sysconfig/network-scripts/ifcfg-sio_mgmt
    echo STARTMODE=onboot >> /etc/sysconfig/network-scripts/ifcfg-sio_mgmt
    echo USERCONTROL=no >> /etc/sysconfig/network-scripts/ifcfg-sio_mgmt
    echo BOOTPROTO=static >> /etc/sysconfig/network-scripts/ifcfg-sio_mgmt
    echo NETMASK=X.X.X.X >> /etc/sysconfig/network-scripts/ifcfg-sio_mgmt
    echo IPADDR=X.X.X.X >> /etc/sysconfig/network-scripts/ifcfg-sio_mgmt
    echo 1 >/sys/bus/pci/rescan
```

   b. For the `ifcfg-sio_d_1` file, run:

```
echo DEVICE= sio_d_2 > /etc/sysconfig/network-scripts/ifcfg-sio_d_1
echo STARTMODE=onboot >> /etc/sysconfig/network-scripts/ifcfg-sio_d_1
echo USERCONTROL=no >> /etc/sysconfig/network-scripts/ifcfg-sio_d_1
echo BOOTPROTO=static >> /etc/sysconfig/network-scripts/ifcfg-sio_d_1
echo NETMASK=X.X.X.X >> /etc/sysconfig/network-scripts/ifcfg-sio_d_1
echo IPADDR=X.X.X.X >> /etc/sysconfig/network-scripts/ifcfg-sio_d_1
echo 1 >/sys/bus/pci/rescan
```

c. For the `ifcfg-sio_d_2` file, run:

```
echo DEVICE= sio_d_2 > /etc/sysconfig/network-scripts/ifcfg-sio_d_2
echo STARTMODE=onboot >> /etc/sysconfig/network-scripts/ifcfg-sio_d_2
echo USERCONTROL=no >> /etc/sysconfig/network-scripts/ifcfg-sio_d_2
echo BOOTPROTO=static >> /etc/sysconfig/network-scripts/ifcfg-sio_d_2
echo NETMASK=X.X.X.X >> /etc/sysconfig/network-scripts/ifcfg-sio_d_2
echo IPADDR=X.X.X.X >> /etc/sysconfig/network-scripts/ifcfg-sio_d_2
echo 1 >/sys/bus/pci/rescan
```

4. Configure the gateway:

```
echo "NETWORKING=yes" >> /etc/sysconfig/network
echo "GATEWAY=XXX.XXX.XXX.XXX" >> /etc/sysconfig/network
```

where *XXX.XXX.XXX.XXX* is the gateway IP address.

5. Create the `/etc/udev/rules.d/70-persistent-ipoib.rules` file:

```
touch /etc/udev/rules.d/70-persistent-ipoib.rules
```

6. Change the permissions of the file you just created to allow the root user access:

```
chmod a+x /etc/udev/rules.d/70-persistent-ipoib.rules
```

7. Find the MAC address for each NIC:

```
ethtool -p <NIC_name>
```

Output similar to the following should appear:

```
ethtool -P ens192
Permanent address: 00:50:56:a7:b0:19
```

8. For each NIC, edit the `70-persistent-ipoib.rules` file according to the information in the output in the previous step:

```
    echo "ACTION=="add", SUBSYSTEM=="net", DRIVERS=="?*", ATTR{type}=="32",
ATTR{address}=="<address>", NAME="<name>"" >> /etc/udev/rules.d/70-persistent-ipoib.rules
```

9. Reboot to the host.
10. After the network comes up, log in using SSH to the management IP address, and ping all data IP addresses of another node to verify that you have the correct connectivity on each node.

Results

Network configuration is complete.

Configure the ports on SLES nodes - R840 servers

Configure the ports on Linux-based (SLES type) nodes.

Procedure

1. Configure the `/etc/udev/rules.d/70-persistent-net.rules` file to match the following:
   - The 1 GB port is em3 and is used for management.
   - The 10 GB data ports are em1 (Data 1) and p1p1 (Data 2) .
2. Use `vi` to edit the `/etc/udev/rules.d/70-persistent-net.rules` file and replace all the Eth names with the correct allocation:
   - The 1 GB port is em3 and em4.
   - The 10 GB ports are em1, em2, p2p1, p2p2.

   The following example shows how the `/etc/udev/rules.d/70-persistent-net.rules` file will look after the eth ports are configured:

```
# PCI device 0x8086:0x10fb (ixgbe)
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="2c:60:0c:c0:9b:a1", ATTR{dev_id}=="0x0"
=="eth*", NAME="eth1"

# PCI device 0x8086:0x10fb (ixgbe)
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="90:e2:ba:8b:9b:bd", ATTR{dev_id}=="0x0"
=="eth*", NAME="eth3"

# PCI device 0x8086:0x1521 (igb)
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="2c:60:0c:ca:ff:49", ATTR{dev_id}=="0x0"
=="eth*", NAME="eth5"

# PCI device 0x8086:0x10fb (ixgbe)
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="90:e2:ba:8b:9b:bc", ATTR{dev_id}=="0x0"
=="eth*", NAME="eth2"

# PCI device 0x8086:0x1521 (igb)
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="2c:60:0c:ca:ff:40", ATTR{dev_id}=="0x0"
=="eth*", NAME="eth4"

# PCI device 0x8086:0x10fb (ixgbe)
SUBSYSTEM=="net", ACTION=="add", DRIVERS=="?*", ATTR{address}=="2c:60:0c:c0:9b:a3", ATTR{dev_id}=="0x0"
=="eth*", NAME="eth0"
```

3. After updating the `70-persistent-net.rules` files, reboot the node using the `reboot` command.
4. After the node comes up, log in again to the console.
5. Create the `ifcfg-emX` files for all interfaces.

   Note: A good way to create a file is using the `cp` and `echo` commands:

   Example:
   ```
   cp /etc/sysconfig/network/ifcfg-em3 /etc/sysconfig/network/ifcfg-em4;
   echo "" > /etc/sysconfig/network/ifcfg-em4;
   ```
6. Create the following files:
   - Eth configuration files for 1 G interface as ifcfg-em3
   - Eth Configuration files for 10 G interface as (first 10 G NIC) ifcfg-em1 & ifcfg-em2
   - Eth Configuration files for 10 G interface as (second 10 G NIC) ifcfg-p2p1 & ifcfg-p2p2

   Refer to the following commands as examples for how to populate each file. Replace each use of "X" with the relevant interface, subnet or IP address.

Example:

```
cp /etc/sysconfig/network/ifcfg-em3 /etc/sysconfig/network/ifcfg-em4
echo DEVICE=em3 > /etc/sysconfig/network/ifcfg-emX
echo STARTMODE=onboot >> /etc/sysconfig/network/ifcfg-emX
echo USERCONTROL=no >> /etc/sysconfig/network/ifcfg-emX
echo BOOTPROTO=static >> /etc/sysconfig/network/ifcfg-emX
echo NETMASK=X.X.X.X >> /etc/sysconfig/network/ifcfg-emX
echo IPADDR=X.X.X.X >> /etc/sysconfig/network/ifcfg-emX
echo 1 >/sys/bus/pci/rescan
sleep 2
ifup emX
```

7. Configure the default gateway:

```
echo "default <XXX.XXX.XXX.XXX> - em3 " >/etc/sysconfig/network/ifroute-emX
```

For example:

```
echo "default - em3" > /etc/sysconfig/network/ifroute-em3
```

Note: Only the one default gateway is required. Most likely this will be the management network interface. Create this file by using the following command:

```
touch /etc/sysconfig/network/ifroute-emX
```

8. Perform reboot to the host.
9. After the host is up, log in to the management IP address using SSH, and ping all data IP addresses of another node to make sure that you have the correct connectivity on each node.

Results

The network configuration on the SLES node is complete. Next perform disk configuration procedures.

CHAPTER 6

Installing the Drivers

The following topics contain information regarding VxFlex OS drivers.

- Install the VxFlex OS drivers on a Linux server 54

Install the VxFlex OS drivers on a Linux server

Dell EMC recommends that you install all applicable VxFlex OS drivers on each Linux server in preparation for installing the VxFlex OS software.

Procedure

1. Download the VxFlex OS driver Zip file from https://support.emc.com/products/42216 to a Windows-based server.
2. Using a Windows OS-based system, unzip the file's contents to a temporary location on Drive C, such as `C:\temp`.

   The unzipped file displays a list of folders that correspond to the hardware components that may be installed on your server.
3. Browse through the folders and locate the Linux drivers that correspond to the hardware components installed on the server.
4. Using an application such as `winscp`, copy the driver files to the VxFlex OS server.
5. Open a console connection to the VxFlex OS server and install each driver file using the following command:

   ```
   rpm -ivh <rpm name>
   ```

   Note: In some cases a dependent rpm is required and should be in the same folder.
6. After all the VxFlex OS drivers are installed, reboot the server.
7. To verify that the drivers are loaded, run the `lsmod` command to list the loaded drivers, and then use `modinfo` to view the driver information for each of the drivers in the list.

CHAPTER 7

Preparing Disks

The following topics describe how to prepare disks so that they can be added to VxFlex OS SDS devices.

- Verify the disk controller type 56
- Enable PERCCLI 56
- Create virtual drives with PERCCLI 57
- Retrieve device paths on the server 61

Verify the disk controller type

Verify the controller type installed on a VxFlex Ready Node R640 or R740xd server using the integrated Dell Remote Access Controller (iDRAC) web utility.

Before you begin

Ensure that you have access to:

- The IP address of the iDRAC port
- The username and password for accessing the iDRAC

About this task

The server contains one of two types of controller cards:

- HBA330 controller card: This server has only SSD.
- H730p/H740p controller card: This server may have either only HDD, or a mix of HDD and SSD.

Procedure

1. From your Internet browser, go to http://*<iDRAC_IP_address>*.
2. In the **DELL Console Login** window, type the user name and password, and click **Login**.

   The **Integrated Remote Access Controller** screen appears.
3. In the navigation pane, select **Storage** > **Controllers**.

   Controller-related information is displayed in the **Health and Properties** table, with the controller type shown in the **Name** column.

   For example:

   ```
   PERC H740p Mini (Embedded)
   ```

After you finish

- If the system uses an HBA330 controller card, retrieve the device paths, as described in Retrieve device paths on the server on page 61.
- If the system uses an H730p or H740p controller card, create a virtual drive on the drive, starting by enabling the PERCCLI disk utility on the server.

Enable PERCCLI

Enable the PERCCLI disk utility on a VxFlex Ready Node server.

Before you begin

Ensure that you have:

- Root access to the server on which PERCCLI is to be installed
- Internet access to enable you to download the PERCCLI utility

Procedure

1. Download PERCCLI from https://downloads.dell.com/FOLDER05235308M/1/perccli_linux_NF8G9_A07_7.529.00.tar.gz%20.
2. Copy the installation file, `perccli_linux_NF8G9_A07_7.529.00.tar.gz`, to the `/tmp` folder on the node.
3. Establish a connection to the SDS node:

- For ESXi-based systems using DirectPath, connect to the SVM where the SDS is installed.
- For Linux-based systems, connect to the SDS.

4. Install the PERCCLI package:

```
rpm -Uvh /tmp/perccli_linux_NF8G9_A07_7.529.00.tar.gz
```

Results

PERCCLI is ready for use.

Create virtual drives with PERCCLI

Create virtual drives (VDs) on drives using the PERCCLI utility. You are required to perform this step before adding drives to a VxFlex OS SDS server in the VxFlex Ready Node system.

Note: Skip this step if the server contains the HBA330 controller or if you are configuring an VxFlex Ready Node R840 system.

When creating virtual drives, the following rules apply:

- Create one virtual drive for each kind of drive, HDD or SSD.
- The VDs for HDD are configured with Direct IO, write-back, and read-ahead enabled (abbreviated in PERCCLI as RWBD).
- The VDs for SSD are configured with Direct IO and write-through, but read-ahead is not enabled. (Abbreviated in PERCCLI as NRWTD.)

In the following PERCCLI command, the controller ID is assumed to be 0, designated by `/c0`. In some cases, the controller ID may be different, which will cause a `Controller 0 not found` error to be displayed. If this occurs, determine your controller ID by running the following command:

```
/opt/MegaRAID/perccli/perccli64 show
```

The controller ID is displayed in the `System Overview` section of the output, in the `Ctl` column.

The following table lists the different drive configurations on a VxFlex Ready Node R640 or R740xd server:

Node contains	Caching options
HDD for storage	NA
Mix of HDD and SSD for storage	Caching not used
Mix of HDD and SSD, where some HDDs or SSDs are for storage and others provide acceleration	RFcache

Before creating VDs, you must gather some information about the drives and devices configured on each node.

Verify the storage and disk controller status

Before creating virtual drives (VDs), verify the storage and disk controller status.

About this task

In this task, use the PERCCLI utility to perform the following tasks:

- Verifying that all the disks are in UGood (unconfigured, but good) status

- Removing all the existing VDs from the node
- Setting up the controller card boot parameter

Procedure

1. Log in to the node.
2. Display the disk information on the node:

```
/opt/MegaRAID/perccli/perccli64 /c0/eall/sall show
```

Output similar to the following is displayed:

```
Drive Information :
==================

---------------------------------------------------------------------------
EID:Slt DID State DG       Size Intf Med SED PI SeSz Model            Sp
---------------------------------------------------------------------------
32:0      0 UGood -  744.625 GB SAS  SSD N   N  512B LT0800MO         U
32:1      1 UGood -  744.625 GB SAS  SSD N   N  512B LT0800MO         U
32:2      2 UGood -  744.625 GB SAS  SSD N   N  512B LT0800MO         U
32:3      3 UGood -  744.625 GB SAS  SSD N   N  512B LT0800MO         U
32:4      4 UGood -  744.625 GB SAS  SSD N   N  512B LT0800MO         U
32:5      5 UGood -  744.625 GB SAS  SSD N   N  512B LT0800MO         U
32:6      6 UGood -    1.091 TB SAS  HDD N   N  512B ST1200MM0088     U
32:7      7 UGood -    1.091 TB SAS  HDD N   N  512B ST1200MM0088     U
32:8      8 UGood -    1.091 TB SAS  HDD N   N  512B ST1200MM0088     U
32:9      9 UGood -    1.091 TB SAS  HDD N   N  512B ST1200MM0088     U
32:10    10 UGood -    1.091 TB SAS  HDD N   N  512B ST1200MM0088     U
32:11    11 UGood -    1.091 TB SAS  HDD N   N  512B ST1200MM0088     U
```

The output shows the following:

- Enclosure ID (EID): Used in a later step when creating VDs
- Slot ID (SLT) of each drive: Used in a later step when creating VDs
- Drive type: Displays SSD or HDD as the device type. (SSD drives are smaller—744 GB in this case.)
- Drive state: Should be UGood (unused and good)

3. Verify that the VDs exist on the node:

```
/opt/MegaRAID/perccli/perccli64 /c0/vall show
```

Output similar to the following is displayed:

```
Controller = 0
Status = Success
Description = No VDs have been configured
```

If the output confirms that VDs exist on the node, perform the next step to remove them. Otherwise, skip the next step.

4. Remove all existing VDs from the node:

```
/opt/MegaRAID/perccli/perccli64 /c0/vall del
```

The VDs are removed.

5. Set the controller card boot parameter:

```
/opt/MegaRAID/perccli/perccli64 /c0 set bootwithpinnedcache=on
```

The boot parameter of the controller card is defined.

Results

You have verified that the disks are in the UGood (unconfigured, but good) state and that the controller card boot parameter is set to **on**. You can now create virtual drives.

Create virtual drives for an HDD using PERCCLI

Use PERCCLI to create virtual drives (VDs) for HDDs on a VxFlex Ready Node server. The devices can be used as part of a Storage Pool or to provide caching.

Before you begin

Ensure that the PERCCLI disk utility is installed.

About this task

You need to create one VD for each HDD. The VDs for HDDs are configured with Direct IO, write-back, and read-ahead enabled.

Procedure

1. Log in to the server.
2. Create a VD that can be accelerated using a utility such as RFcache, or that will not be accelerated at all:

   ```
   /opt/MegaRAID/perccli/perccli64 /c0 add vd type=raid0 drives=<EID>:<Slt> direct wb ra
   pdcache=off
   ```

   Example:

   ```
   /opt/MegaRAID/perccli/perccli64 /c0 add vd type=raid0 drives=32:6 direct wb ra
   pdcache=off
   ```

   where *<EID>* and *<Slt>* are the Enclosure ID and Slot ID values, identified in a previous step.

   Output similar to the following is displayed:

   ```
   Controller = 0
   Status = Success
   Description = Add VD Succeeded
   ```

Results

The VD was created successfully.

Create virtual drives for SSD using PERCCLI

Use PERCCLI to create virtual drives (VDs) for SSDs on a VxFlex Ready Node server.

Before you begin

Ensure that the PERCCLI disk utility is installed.

About this task

Use this procedure both for SSD for storage and SSD to be used to provide caching acceleration. The VDs created on SSDs are configured with Direct IO and write-through, but read-ahead is not enabled.

Procedure

1. Log in to the server.

2. Create the VD:

```
/opt/MegaRAID/perccli/perccli64 /c0 add vd type=raid0 drives=<EID>:<Slt> direct wt nora
```

Example:

```
opt/MegaRAID/perccli/perccli64 /c0 add vd type=raid0 drives=32:0 direct wt nora
```

where *EID* and *Slt* are the Enclosure ID and Slot ID values, which in this example are 32:0.

Results

The VD was created successfully.

After you finish

Continue with the procedure that referred you to this topic.

Ensure virtual drive creation with PERCCLI

Use the PERCCLI utility to ensure that a virtual drive (VD) was created properly on each drive.

Procedure

1. Log in to the server.
2. Display the physical device (PD) information:

```
/opt/MegaRAID/perccli/perccli64 /c0/eall/sall show
```

Output similar to the following is displayed:

```
Drive Information :
==================

---------------------------------------------------------------------------
EID:Slt DID State DG       Size Intf Med SED PI SeSz Model            Sp
---------------------------------------------------------------------------
32:0      0 Onln   1 744.625 GB SAS  SSD N   N  512B LT0800MO         U
32:1      1 Onln   2 744.625 GB SAS  SSD N   N  512B LT0800MO         U
32:2      2 Onln   3 744.625 GB SAS  SSD N   N  512B LT0800MO         U
32:3      3 Onln   4 744.625 GB SAS  SSD N   N  512B LT0800MO         U
32:4      4 Onln   5 744.625 GB SAS  SSD N   N  512B LT0800MO         U
32:5      5 Onln   6 744.625 GB SAS  SSD N   N  512B LT0800MO         U
32:6      6 Onln   7   1.090 TB SAS  HDD N   N  512B ST1200MM0088     U
32:7      7 Onln   8   1.090 TB SAS  HDD N   N  512B ST1200MM0088     U
32:8      8 Onln   9   1.090 TB SAS  HDD N   N  512B ST1200MM0088     U
32:9      9 Onln  10   1.090 TB SAS  HDD N   N  512B ST1200MM0088     U
32:10    10 Onln  11   1.090 TB SAS  HDD N   N  512B ST1200MM0088     U
32:11    11 Onln  12   1.090 TB SAS  HDD N   N  512B ST1200MM0088     U
```

3. Verify that all the devices have a virtual drive created on them.

   In the **Drive Information** output, the device state displayed as "Onln" in the **State** column shows that the device has a virtual drive created on it.

4. Display device group (DG) and virtual drive (VD) information:

```
/opt/MegaRAID/perccli/perccli64 /c0/vall show
```

Output similar to the following is displayed:

```
Virtual Drives :
==============

--------------------------------------------------------------
DG/VD TYPE  State Access Consist Cache sCC       Size Name
--------------------------------------------------------------
0/0   RAID0 Optl  RW     Yes     NRWTD -   744.625 GB
1/1   RAID0 Optl  RW     Yes     NRWTD -   744.625 GB
2/2   RAID0 Optl  RW     Yes     NRWTD -   744.625 GB
3/3   RAID0 Optl  RW     Yes     NRWTD -   744.625 GB
4/4   RAID0 Optl  RW     Yes     NRWTD -   744.625 GB
5/5   RAID0 Optl  RW     Yes     NRWTD -   744.625 GB
6/6   RAID0 Optl  RW     Yes     RWBD  -     1.090 TB
7/7   RAID0 Optl  RW     Yes     RWBD  -     1.090 TB
8/8   RAID0 Optl  RW     Yes     RWBD  -     1.090 TB
9/9   RAID0 Optl  RW     Yes     RWBD  -     1.090 TB
10/10 RAID0 Optl  RW     Yes     RWBD  -     1.090 TB
11/11 RAID0 Optl  RW     Yes     RWBD  -     1.090 TB
```

5. Verify that all the VDs are configured correctly.

   The following values should appear in the display:

   - Cache = NRWTD (for SSDs)
   - Cache = RWBD (for HDDs)

After you finish

Continue by retrieving the device path on the server.

Retrieve device paths on the server

The manner of retrieving device paths in a Linux-based VxFlex Ready Node server differs, depending on the type of controller card in the node.

Proceed to the section that matches your configuration:

- Retrieve device paths in a Linux server with an HBA330 controller on page 61
- Retrieving device paths in a Linux server with an H730p/H740p controller on page 62
- Retrieving device paths in a Linux server with NVMe devices configured on page 63

Retrieve device paths in a Linux server with an HBA330 controller

Retrieve the OS device paths in a Linux-based VxFlex Ready Node server with an HBA330 controller.

About this task

In Linux-based systems, the OS device path is in the following format: `/dev/sdX`, where *X* is a sequential letter of the alphabet in lowercase, in single or multiple format. For example: a, b, c...z, followed by aa, ab, ac...zz.

Procedure

1. Log in to the server.
2. Display the device list:

```
ls -l /dev/disk/by-path/
```

Output similar to the following appears:

```
[root@I22CJ-6 ~]# ls -l /dev/disk/by-path/
total 0
lrwxrwxrwx 1 root root  9 Jun  9 01:22 pci-0000:00:1f.2-scsi-4:0:0:0 -> ../../sdk
lrwxrwxrwx 1 root root 10 Jun  9 01:22 pci-0000:00:1f.2-scsi-4:0:0:0-part1 -> ../../sdk1
lrwxrwxrwx 1 root root 10 Jun  9 01:22 pci-0000:00:1f.2-scsi-4:0:0:0-part2 -> ../../sdk2
lrwxrwxrwx 1 root root 10 Jun  9 01:22 pci-0000:00:1f.2-scsi-4:0:0:0-part3 -> ../../sdk3
lrwxrwxrwx 1 root root  9 Jun  9 01:22 pci-0000:02:00.0-sas-0x50000396dc89a826-lun-0 -> ../../sdj
lrwxrwxrwx 1 root root  9 Jun  9 01:22 pci-0000:02:00.0-sas-0x50000396dc89a832-lun-0 -> ../../sda
lrwxrwxrwx 1 root root  9 Jun  9 01:22 pci-0000:02:00.0-sas-0x50000396dc89a836-lun-0 -> ../../sdh
lrwxrwxrwx 1 root root  9 Jun  9 01:22 pci-0000:02:00.0-sas-0x50000396dc89a83a-lun-0 -> ../../sdd
lrwxrwxrwx 1 root root  9 Jun  9 01:22 pci-0000:02:00.0-sas-0x50000396dc89a83e-lun-0 -> ../../sdc
lrwxrwxrwx 1 root root  9 Jun  9 01:22 pci-0000:02:00.0-sas-0x50000396dc89a842-lun-0 -> ../../sde
lrwxrwxrwx 1 root root  9 Jun  9 01:22 pci-0000:02:00.0-sas-0x50000396dc89a846-lun-0 -> ../../sdi
lrwxrwxrwx 1 root root  9 Jun  9 01:22 pci-0000:02:00.0-sas-0x50000396dc89a86a-lun-0 -> ../../sdf
lrwxrwxrwx 1 root root  9 Jun  9 01:22 pci-0000:02:00.0-sas-0x50000396dc89a86e-lun-0 -> ../../sdg
lrwxrwxrwx 1 root root  9 Jun  9 01:22 pci-0000:02:00.0-sas-0x50000396dc89a876-lun-0 -> ../../sdb
```

3. In the output, search for the lines starting with:

```
pci-0000:0#:00.0-sas-
```

where *#* is a number. (For example: 2 or 3.)

4. In the lines you just located, search for `sdX` at the end of the lines.

   The devices paths are `/dev/sdX`. Use these paths when adding devices to a VxFlex OS SDS.

Retrieving device paths in a Linux server with an H730p/H740p controller

Retrieve the OS device path for a specific virtual drive (VD) in a Linux-based VxFlex Ready Node server with an H730p/H740p controller.

Before you begin

Ensure that you have root access to the server.

About this task

In Linux servers, the OS device path is in the format: `/dev/sdX`, where X is a sequential letter of the alphabet in lowercase, in single or multiple format. For example: a, b, c...z, followed by aa, ab, ac...zz.

Procedure

1. Log in to the server.
2. Display the VD information:

```
/opt/MegaRAID/perccli/perccli64 /c0/vall show
```

Output similar to the following appears:

```
Virtual Drives :
===============

--------------------------------------------------------------
DG/VD TYPE  State Access Consist Cache sCC       Size Name
--------------------------------------------------------------
0/0   RAID0 Optl  RW     Yes     NRWTD -   744.625 GB
1/1   RAID0 Optl  RW     Yes     NRWTD -   744.625 GB
2/2   RAID0 Optl  RW     Yes     NRWTD -   744.625 GB
3/3   RAID0 Optl  RW     Yes     NRWTD -   744.625 GB
4/4   RAID0 Optl  RW     Yes     NRWTD -   744.625 GB
5/5   RAID0 Optl  RW     Yes     NRWTD -   744.625 GB
6/6   RAID0 Optl  RW     Yes     RWBD  -     1.090 TB
7/7   RAID0 Optl  RW     Yes     RWBD  -     1.090 TB
8/8   RAID0 Optl  RW     Yes     RWBD  -     1.090 TB
9/9   RAID0 Optl  RW     Yes     RWBD  -     1.090 TB
10/10 RAID0 Optl  RW     Yes     RWBD  -     1.090 TB
11/11 RAID0 Optl  RW     Yes     RWBD  -     1.090 TB
```

The output displays the device groups (DG) and their associated VD.

3. Match a VD to a device path:

   a. Run the following command:

   ```
   ls -l /dev/disk/by-path/
   ```

   Output similar to the following appears:

   ```
   sio-dcoe-122CJ-11:~ # ls -l /dev/disk/by-path/
   total 0
   lrwxrwxrwx 1 root root  9 Jun  9 04:18 pci-0000:00:1f.2-scsi-9:0:0:0 -> ../../sda
   lrwxrwxrwx 1 root root 10 Jun  9 04:18 pci-0000:00:1f.2-scsi-9:0:0:0-part1 -> ../../sda1
   lrwxrwxrwx 1 root root 10 Jun  9 04:18 pci-0000:00:1f.2-scsi-9:0:0:0-part2 -> ../../sda2
   lrwxrwxrwx 1 root root  9 Jun 13 08:01 pci-0000:02:00.0-scsi-0:2:0:0 -> ../../sdb
   lrwxrwxrwx 1 root root  9 Jun 13 08:02 pci-0000:02:00.0-scsi-0:2:1:0 -> ../../sdc
   lrwxrwxrwx 1 root root  9 Jun 13 08:02 pci-0000:02:00.0-scsi-0:2:2:0 -> ../../sdd
   lrwxrwxrwx 1 root root  9 Jun 13 08:02 pci-0000:02:00.0-scsi-0:2:3:0 -> ../../sde
   lrwxrwxrwx 1 root root  9 Jun 13 08:02 pci-0000:02:00.0-scsi-0:2:4:0 -> ../../sdf
   lrwxrwxrwx 1 root root  9 Jun 13 08:02 pci-0000:02:00.0-scsi-0:2:5:0 -> ../../sdg
   lrwxrwxrwx 1 root root  9 Jun 13 08:02 pci-0000:02:00.0-scsi-0:2:6:0 -> ../../sdh
   lrwxrwxrwx 1 root root  9 Jun 13 08:02 pci-0000:02:00.0-scsi-0:2:7:0 -> ../../sdi
   lrwxrwxrwx 1 root root  9 Jun 13 08:02 pci-0000:02:00.0-scsi-0:2:8:0 -> ../../sdj
   lrwxrwxrwx 1 root root  9 Jun 13 08:02 pci-0000:02:00.0-scsi-0:2:9:0 -> ../../sdk
   ```

   b. In the command output, search for the line:

   `pci-0000:02:00.0-scsi-0:2:X:0`

   where *X* is the number assigned to the VD.

   c. In the line you just located, search for the `sdX` at the end of the line.

   The device path for the VD is `/dev/sdX`.
   Use this path when adding devices to a VxFlex OS SDS.

Retrieving device paths in a Linux server with NVMe devices configured

Retrieve the OS device paths in a Linux-based VxFlex Ready Node server with NVMe devices configured.

Before you begin

Ensure that you have root access to the server

Procedure

1. Log in to the server.

2. Display the NVMe devices:

```
ls -l /dev/nvm*n1
```

Output similar to the following appears:

```
[root@sio-dcoe-96q-12 ~]# ls -l /dev/nvm*n1
brw-rw----. 1 root disk 259, 4 Jul 10 19:44 /dev/nvme0n1
brw-rw----. 1 root disk 259, 2 Jul 10 20:04 /dev/nvme1n1
brw-rw----. 1 root disk 259, 5 Jul 10 20:19 /dev/nvme2n1
brw-rw----. 1 root disk 259, 3 Jul 10 19:14 /dev/nvme3n1
brw-rw----. 1 root disk 259, 0 Jul 10 19:44 /dev/nvme4n1
brw-rw----. 1 root disk 259, 1 Jul 10 19:29 /dev/nvme5n1
```

Use these paths when adding devices to a VxFlex OS SDS.

CHAPTER 8

Additional configurations

This section describes additional configurations required for the VxFlex Ready Node server installation process.

- Install OpenManage Enterprise 66
- Disable Smartmontool error messages 66
- iDRAC Service Module 67
- Prepare the DAX devices 67
- iDRAC Service Module 71

Install OpenManage Enterprise

Dell EMC OpenManage Enterprise is a web-based console that simplifies hardware monitoring and firmware maintenance on VxFlex Ready Node servers. It is recommended that you install OpenManage Enterprise.

About this task

Note: An additional SupportAssist option is available for hardware home call capabilities.

To learn about using OpenManage Enterprise with VxFlex Ready Node servers, see *Managing VxFlex Ready Nodes using OpenManage Enterprise*. Both this publication and the *VxFlex OS OpenManage Enterprise Catalog* are available from https://support.emc.com/products/42216.

Note: This link requires yourDell EMC account credentials.

Contact your VxFlex OS account team for information on how to order OpenManage Enterprise and SupportAssist.

Disable Smartmontool error messages

Smartmontools may throw repeated drive errors due to an outdated list of supported drives. You can disable these error messages if they become problematic.

About this task

Note: When the Smartmon tool is disabled, device alerts will not be posted and the system may run with components that are not functioning properly.

VxFlex OS utilizes Smartmontools as part of hardware awareness. Smartmontools depends on a local database of drive definitions to work properly. These drive definitions are updated periodically and are maintained by the drive manufacturers.

On occasion, drives become available for use before the drive definition database is updated. Because these drives are not identified properly, error messages are posted to `/var/log/messages`.

For example:

```
smartd[2007]: Device: /dev/sdo [SAT], SMART Usage Attribute: 245 Unknown_Attribute changed from
95 to 94
```

Where `/dev/sdo` is the location of the device triggering the error message.

For additional information regarding Smartmontools and supported devices, see https://www.smartmontools.org/

Procedure

1. Disable the error message on the node:
   - On servers with an HBA330 controller card, at root run:

     ```
     smartctl -s off <device_path>
     ```

     where *<device_path>* is the location of the device that Smartmontools does not recognize.

     For example:

     ```
     smartctl -s off /dev/sdo
     ```

     The command is persistent and will run each time the node reboots.

- On servers with an H730p/H740p controller card, create a startup script containing the following command so that it runs each time the node reboots:

  ```
  smartctl -s off <device_path>
  ```

  where *<device_path>* is the location of the device that Smartmontools does not recognize.

  For example:

  ```
  smartctl -s off /dev/sdo
  ```

2. Repeat the above step on every node on which the errors occur.

iDRAC Service Module

The iDRAC Service module (iSM) is a small OS-resident process that expands iDRAC management into supported host operating systems. Services that the iSM adds include OS information, automatic system recovery, and remote server power cycle.

It also enables NVMe device removal without shutting down or rebooting the system.

Install the iDRAC Service Module on a Linux server

The iDRAC Service Module (iSM) is available as an BIN file for installing on systems running Linux.

Before you begin

Ensure that the openIPMI package is installed on the server.

Procedure

1. Locate the `Systems-Management_Application_49J6N_LN64_3.x.x.x_xxx.bin` file in the drive ZIP file you downloaded from https://support.emc.com/products/42216.
2. Execute the DUP package in your local directory.

   The DUP execution runs on the server.

   Note: You might need to run `chmod a+x` *<file path+name>* prior to executing the file, in order to run the file cd to the relevant folder and run the binary `Management_Application_49J6N_LN64_3.x.x.x_xxx.bin`.

3. From the iDRAC web interface, open the **iDRAC settings** > **Overview** page and ensure that the iDRAC Service Module status is Running.

Prepare the DAX devices

If the deployment will be using NVDIMM devices for acceleration, perform the following tasks to prepare the NVDIMMs as DAX acceleration devices.

NVDIMM information table

Before you prepare the NVDIMM as a DAX device, create or print the following NVDIMM information table. Use it to record the information that you discover while performing the following task. You will need the information to add the acceleration device to VxFlex OS.

Note: If the failed and replacement NVDIMM share a column, the value for both entries must be the same.

Table 16 NVDIMM information table

Item	Replacement NVDIMM
Serial Number	
Device name (NMEM)	
Namespace	
DAX device name (chardev)	
Acceleration device path	

Prepare an NVDIMM as a DAX device

Prepare a new or replacement NVDIMM as a DAX device before adding it to the SDS. This step is optional when replacing an NVDIMM battery.

About this task

Note: If you replaced the system board, perform the following steps for each NVDIMM mounted on the server.

Procedure

1. Using SSH, log in to the Linux server.
2. Vew information for the new/replacement NVDIMM:

```
    dmidecode --type memory | grep "Non-" -B 3 -A 3 | grep -E 'Locator|Serial' | grep -v
Bank
```

Output similar to the following appears:

```
Locator: A7
        Serial Number: 17496594
        Locator: B7
        Serial Number: 174965AC
```

3. Find the serial number in the output and record it in the NVDIMM information table.
4. Display the correlation between the ID and NMEM device name of each NVDIMM mounted on the server:

```
ndctl list -Dvvv | jq '.[0] .dimms'
```

Output similar to the following appears:

```
[
  {
    "dev": "nmem1",
    "id": "802c-0f-1722-174965ac",
    "handle": 4097,
    "phys_id": 4370,
    "health": {
      "health_state": "ok",
      "temperature_celsius": 255,
      "life_used_percentage": 29
    }
  },
  {
    "dev": "nmem0",
```

```
    "id": "802c-0f-1722-17496594",
    "handle": 1,
    "phys_id": 4358,
    "health": {
      "health_state": "ok",
      "temperature_celsius": 255,
      "life_used_percentage": 29
    }
  }
]
```

5. In the output from the previous step, find the device (dev) with the `id` that partially correlates with the serial number you discovered previously for the failed device.

   For example:

   - The NVDIMM output displays serial number 19BA9C2D for the NVDIMM device.
   - In the previous step, the output displays the ID of device nmem0 as 802c-0f-1746-19ba9c2d.

Serial number	Device ID

`Type Detail: Synchronous Non-Volatile Registered (Buffered)`
`Speed: 2666 MHz`
`Manufacturer: 002C00B3002C`
`Serial Number: 19BA9C2D`
`Asset Tag: 0F174651`
	`{`
	`"dev":"nmem0",`
	`"id":"802c-0f-1746-19ba9c2d",`
	`"handle":1,`
	`"phys_id":4358`
	`}`

   In the above example, the NMEM name is nmem0.

6. Record the NMEM name in the `Device name` row of the NVDIMM information table.
7. Correlate between the NMEM DIMM and the namespace/DAX device:

```
ndctl list -d all -v -H
```

   Output similar to the following appears:

```
{
  "dimms":[
    {
      "dev":"nmem1",
      "id":"802c-0f-1746-19ba9c47",
      "handle":4097,
      "phys_id":4370,
      "health":{
        "health_state":"ok",
        "temperature_celsius":36.000000,
        "life_used_percentage":5
      }
    },
    {
      "dev":"nmem0",
      "id":"802c-0f-1746-19ba9c2d",
      "handle":1,
      "phys_id":4358,
      "health":{
        "health_state":"ok",
        "temperature_celsius":32.000000,
        "life_used_percentage":5
      }
    }
  ],
  "namespaces":[
```

```
    {
      "dev":"namespace1.0",
      "mode":"devdax",
      "map":"dev",
      "size":16909336576,
      "uuid":"7d905ce0-49ed-42ba-8ad3-3981eb434f4d",
      "numa_node":1
    },
    {
      "dev":"namespace0.0",
      "mode":"devdax",
      "map":"dev",
      "size":16909336576,
      "uuid":"47165786-f91f-4d33-86bd-6d80aa3141f1",
      "numa_node":0
    }
  ]
```

8. In the output displayed in the previous step, locate the namespace that correlates with the NMEM name and DIMM serial number, and record it in the NVDIMM information table.

   In the above example, nmem0's namespace is namespace0.0.

9. Convert the namespace device to the acceleration device name of type /dev/daxX.X:

```
ndctl create-namespace -f -e namespaceX.X --mode=devdax -align=4K
```

   For ndctl build 62 or later, run:

```
ndctl create-namespace -f -e namespaceX.X --mode=devdax -align=4K --no-autolabel
```

   where namespace*X.X* is the actual namespace of the NVDIMM.

   For example, to convert namespace0.0 to /dev/dax0.0, where /dev/daxX.X is the acceleration device name, depending on the ndctl build run:

   Note: `--no-autolabel` is only relevant for RHEL 7.6.

```
ndctl create-namespace -f -e namespace0.0 -mode=dax -align=4K
```

   or

```
ndctl create-namespace -f -e namespace0.0 -mode=dax -align=4K --no-autolabel
```

10. Record the acceleration device name in the NVDIMM information table.
11. Run the `namespace-to-dax-device` correlation command to find the DAX device name of the replacement NVDIMM:

```
ndctl list --regions --verbose
```

   Output similar to the following appears:

```
{
  "regions":[
    {
```

```
      "dev":"region0",
      "size":17179869184,
      "available_size":0,
      "max_available_extent":0,
      "type":"pmem",
      "numa_node":0,
      "persistence_domain":"unknown",
      "namespaces":[
        {
          "dev":"namespace0.0",
          "mode":"devdax",
          "map":"dev",
          "size":16909336576,
          "uuid":"47165786-f91f-4d33-86bd-6d80aa3141f1",
          "daxregion":{
            "id":0,
            "size":16909336576,
            "align":2097152,
            "devices":[
              {
                "chardev":"dax0.0",
                "size":16909336576
              }
            ]
          },
          "numa_node":0
        }
      ]
    }
  ]
}
```

The DAX device name appears in the output as the chardev value.

In the example output above, the DAX device name is dax0.0.

12. Record the DAX device name in the NVDIMM information table.
13. Find the full acceleration device path:

```
ls -ltr /dev/ | grep dax
```

The following output should appear:

```
/dev/daxX.X
```

For example:

```
/dev/dax0.0
```

14. Record the acceleration device path in the NVDIMM information table.

Results

You are now ready to add the DAX device to the NVDIMM Acceleration Pool.

iDRAC Service Module

The iDRAC Service module (iSM) is a small OS-resident process that expands iDRAC management into supported host operating systems. Services that the iSM adds include OS information, automatic system recovery, and remote server power cycle.

It also enables NVMe device removal without shutting down or rebooting the system.

iDRAC Service Module

The iDRAC Service module (iSM) is a small OS-resident process that expands iDRAC management into supported host operating systems. Services that the iSM adds include OS information, automatic system recovery, and remote server power cycle.

It also enables NVMe device removal without shutting down or rebooting the system.

Install the iDRAC Service Module on a Linux server

The iDRAC Service Module (iSM) is available as an BIN file for installing on systems running Linux.

Before you begin

Ensure that the openIPMI package is installed on the server.

Procedure

1. Locate the `Systems-Management_Application_49J6N_LN64_3.x.x.x_xxx.bin` file in the drive ZIP file you downloaded from https://support.emc.com/products/42216.
2. Execute the DUP package in your local directory.

   The DUP execution runs on the server.

   Note: You might need to run `chmod a+x` *<file path+name>* prior to executing the file, in order to run the file cd to the relevant folder and run the binary `Management_Application_49J6N_LN64_3.x.x.x_xxx.bin`.

3. From the iDRAC web interface, open the **iDRAC settings** > **Overview** page and ensure that the iDRAC Service Module status is Running.

CHAPTER 9

Next steps

Next steps and operating system-specific guidelines for deploying VxFlex OS on the servers that you have prepared

- Next steps 74

Next steps

The server is now ready for VxFlex OS deployment. Take note of the following required steps for deploying your system:

The *Dell EMC Deploy VxFlex OS Guide* explains how to deploy VxFlex OS on VxFlex Ready Node servers. Follow the preparation guidelines and deployment procedures relevant to your environment.

If you are new to VxFlex OS and VxFlex Ready Node products, Dell EMC recommends that you refer to: How to Find Information for VxFlex OS.

Deployment limitations and requirements:

Pay attention to the following when using the *Deploy VxFlex OS Guide* to deploy VxFlex OS on R840 systems:

- VxFlex Ready Node R840 with NVDIMM acceleration is not supported at this time.
- For VxFlex Ready Node R840 functioning as a Compute Node (SDC and LIA only) as part of a 2-layer architecture based system, refer to the "Linux" section, 2-layer system topics in the *Deploy VxFlex OS Guide*.
- For VxFlex Ready Node R840 with HBA330 SAS controller on ESXi, refer to the "VMware ESXi" section, DirectPath configuration flow in the *Deploy VxFlex OS Guide*. All other VMware-based options are not supported for this controller.
- For VxFlex Ready Node R840 with NVMe disks on ESXi, refer to the "VMware ESXi" section, RDM configuration flow in the *Deploy VxFlex OS Guide*.

CHAPTER 10

Reference material

This section contains additional information that may be required for the procedures described in this document.

- DTK - Hardware Update Bootable ISO 76

DTK - Hardware Update Bootable ISO

The Dell EMC OpenManage Deployment Toolkit (DTK) includes a set of utilities, sample scripts, and sample configuration files that you can use to deploy and configure the Dell EMC systems. This section provides additional information on using the VxFlex Ready Node Hardware Update Bootable ISO ("Hardware ISO") to update drivers, BIOS, and firmware on VxFlex Ready Node servers.

Dell EMC OpenManage DRAC Tools (RACADM)

An integrated Dell Remote Access Controller (iDRAC) with Lifecycle Controller is embedded in every Dell EMC VxFlex Ready Node server. The RACADM command-line utility provides a scriptable interface that enables you to configure the iDRAC either locally or remotely. The utility runs on the management station and the managed system.

The RACADM utility supports the following interfaces:

- Local - Supports running RACADM commands from the managed server's operating system. To run local RACADM commands, install the OpenManage software on the managed server. Only one instance of Local RACADM can be executed on a system at a time.
- SSH or Telnet (also known as Firmware RACADM) - Firmware RACADM is accessible by logging in to iDRAC using SSH or Telnet.
- Remote - Supports running RACADM commands from a remote management station such as a laptop or desktop running Windows or Linux. To run Remote RACADM commands, install the DRAC Tools utility from the OpenManage software on the remote computer.

Note: For information regarding RACADM commands, see the RACADM CLI Guide.

Update the hardware using remote RACADM

You can install and execute the Dell EMC RACADM tool from any management system with access to the iDRAC network. The remote RACADM command set is useful in this situation to mount and execute the Hardware ISO to a large number of VxFlex Ready Node servers. This prevents the need for multiple web browser windows and manual keyboard and mouse clicks for each server.

Procedure

1. Configure a laptop or server with access to the iDRAC network.
2. Install the DRAC tools, including RACADM:
   a. In any Internet browser, go to the RACADM Command Line Interface for DRAC page in the Dell EMC Systems Management wiki.
   b. Follow the instructions for downloading and installing Remote RACADM for your operating system.
3. Download the VxFlex Ready Node Hardware Update Bootable ISO from the https://support.emc.com/products/42216 and make it accessible on the network share folder.
4. Create a CIFS or NFS network share that includes the Hardware ISO. The share will need to be accessible to the iDRAC network.

   Refer to OS vendor guidelines for instructions on how to create CIFS (windows) or NFS (Linux or XenServer) shares
5. Open a terminal/command prompt with superuser/administrator accessibility. This prompt should allow execution of the RACADM utility.

6. Mount the Hardware ISO to the iDRAC from the remote share, where the following command is all on one line:

```
racadm -r <dracIP> -u root -p <password> remoteimage -c -u <myuser> -p <mypass> -l //
<myip>/SIORN/VxFlex-Ready-Node-Hardware-Update-for-Dell_14G_2018_May_A00.iso
```

Where:

- *<dracIP>* is the iDRAC IP address
- *<password>* is the password for the server
- *<myuser>* is the NFS/CIFS share username
- *<mypass>* is the NFS/CIFS share password
- *<myip>* is the NFS/CIFS server IP

Note: The default password is Scaleio123.

7. Enable the iDRAC VirtualCD to boot once using these Remote RACADM commands:

```
racadm -r <dracIP> -u root -p <password> set iDRAC.ServerBoot.BootOnce Enabled
racadm -r <dracIP> -u root -p <password> set iDRAC.ServerBoot.FirstBootDevice VCD-DVD
```

8. Verify that the server is in a state in which it can be power-cycled. (The OS should be in maintenance mode, as needed).
9. Power-cycle the VxFlex Ready Node server:

```
racadm -r <dracIP> -u root -p <password> serveraction graceshutdown
```

10. Repeat steps 6 on page 77 - 9 on page 77 for each VxFlex Ready Node server needing the BIOS and firmware updates and configuration. All servers can be updated in parallel.
11. Wait for the configuration and firmware updates to complete. The server console screen will indicate when the script is complete.

WARNING Do not reboot the VxFlex Ready Node server while the update process is being performed!

The iDRAC will be reset several times during the update process. This causes the iDRAC virtual console viewer to close, virtual media to disconnect, and the iDRAC browser window to be unavailable for several minutes during each reset. The hardware update scripts will continue to run from RAM on the server.

The update script will generate a log indicating whether each configuration and firmware flash is successful.

12. (Optional) Check each VxFlex Ready Node server's log for successful competition:
    a. Connect to the iDRAC KVM console screen.
    b. After the update script completes, press Alt+F2 to access a user console, and then Enter to log in.
    c. Check the log contents for errors:

```
less /bundleapplicationlogs/apply_components.log
```

For more information, see Troubleshooting the Hardware ISO on page 80 .

    d. Press q to exit the log viewer, and then Alt+F1 to access the original console screen.

13. Reboot the servers and allow the update and configuration jobs to complete. You can do this remotely using RACADM (see Step 9 on page 77) or locally from the server console screen.
14. For each server, check the iDRAC job queue to verify that the iDRAC jobs created by the bootable ISO have completed successfully:

    ```
    racadm -r <dracIP> -u root -p <password> jobqueue view
    ```

    Any job failures may require running the bootable ISO again, or further troubleshooting.

15. For each server, after the hardware updates are finalized, clean up the iDRAC job queue:

    ```
    racadm -r <dracIP> -u root -p <password> jobqueue delete --all
    ```

16. Continue with OS installation and configuration, VxFlex OS deployment, and other required tasks.

    At the end of the process, the BIOS password changes to emcbios.

Recommended BIOS and firmware settings

This section describes the BIOS, firmware, and configuration settings included in the Hardware ISO.

VxFlex OS ID module

The Hardware ISO runs a script that automatically flashes the ID module, as needed.

This action rebrands server naming, and is intended to make it compatible with the VxFlex Ready Node offering.

The VxFlex Ready Node ID module should be installed only on VxFlex Ready Node servers. In the event of a system board failure, the Hardware ISO can assist with the reinstallation of the ID module on the replacement system board.

BIOS and firmware

The Hardware ISO runs a script that automatically forces the server to install the necessary firmware updates.

This firmware is consistent with the qualified Dell EMC VxFlex Ready Node Driver and Firmware Matrix, located at https://support.emc.com/products/42216.

Some of the firmware listed in the table is dependent on the VxFlex Ready Node hardware configuration. The Hardware ISO attempts to apply all firmware updates, but only those updates that are compatible will be installed.

Applying settings using RACADM

The individual firmware files are also available on the Dell EMC Online Support site, and can easily be installed using the following remote RACADM command:

```
racadm -r <dracIP> -u root -p <password> update -f <filename.exe>
```

Where:

- *<dracIP>* is the iDRAC IP address
- *<password>* is the password for the server
- *<filename.exe>* is the name of the Dell EMC Windows update packages
- *<dracIP>* is the iDRAC IP address
- *<password>* is the password for the server

Note: The default password is Scaleio123.

Configuration settings

The Hardware ISO runs a script that automatically configures the BIOS and iDRAC settings listed in the table below. Some settings are dependent on the server model.

Table 17 Hardware ISO configuration settings

Description	Setting	Value
Server BIOS Boot Sequence	BIOS.BiosBootSettings.BootSeq	HardDisk.List.1-1
Hard Disk Boot Order	BIOS.BiosBootSettings.HddSeq	Disk.SATAEmbedded.J-1,RAID.Integrated.1-1 or Disk.SATAEmbedded.D-1,RAID.Integrated.1-1
Server Boot Mode	BIOS.BiosBootSettings.BootMode	Bios
SRIOV Global Enablement	BIOS.IntegratedDevices.SriovGlobalEnable	Enabled
Memory Performance Tuning	BIOS.MemSettings.SnoopMode	EarlySnoop
System Change Tracking	BIOS.MiscSettings.InSystemCharacterization	Disabled
CPU Virtualization Features	BIOS.ProcSettings.ProcVirtualization	Enabled
CPU Cores	BIOS.ProcSettings.ProcCores	All
CPU X2APIC Mode	BIOS.ProcSettings.ProcX2Apic	Enabled
CPU Turbo Engagement	BIOS.ProcSettings.ControlledTurbo	Disabled
System Power Profile	BIOS.SysProfileSettings.SysProfile	PerfOptimized
OS to iDRAC Pass-through Mode	iDRAC.OS-BMC.PTMode	usb-p2p
OS to iDRAC Enablement	iDRAC.OS-BMC.AdminState	Enabled
iDRAC DHCP Enablement	iDRAC.IPv4.DHCPEnable	Disabled
iDRAC Default Credentials	iDRAC.Tuning.DefaultCredentialWarning	Disabled
iDRAC IPMI Enablement	iDRAC.IPMILan.Enable	Enabled
iDRAC Alert Enablement	iDRAC.IPMILan.AlertEnable	Enabled
iDRAC IPv6 Enablement	iDRAC.IPv6.Enable	Enabled

Applying settings using RACADM

The individual settings can also be applied using the remote RACADM command:

```
racadm -r <dracIP> -u root -p <password> set <setting> <value>
```

Where:

- *<dracIP>* is the iDRAC IP address

- *<password>* is the password for the server
- *<setting>* is the BIOS/iDRAC setting name
- *<value>* is the BIOS/iDRAC setting value

Note:
When setting the BIOS configuration, include this command:

```
racadm -r <dracIP> -u root -p <password> jobqueue create BIOS.Setup.1-1 <value>
```

Troubleshooting the Hardware ISO

This section describes troubleshooting procedures for problems you may encounter while using the Hardware ISO.

Troubleshoot general iDRAC failures

When problems occur with iDRAC or Lifecycle Controller jobs, you can delete all jobs with a single iDRAC command. All of the completed jobs, plus any orphaned pending jobs, are deleted, and the data manager service on the iDRAC is restarted.

Procedure

1. Clear the iDRAC job queue:

   ```
   racadm -r <dracIP> -u root -p <password> jobqueue delete -i JID_CLEARALL_FORCE
   ```

   Where:

   - *<dracIP>* is the iDRAC IP address
   - *<password>* is the password for the server

   Note: The default password is Scaleio123.

2. Wait 120 seconds.

   The iDRAC is unable to process any other jobs during this time.

3. Reset the iDRAC

   ```
   racadm -r <dracIP> -u root -p <password> racreset
   ```

   The iDRAC becomes accessible on the network 3 to 5 minutes after the reset.

iDRAC virtual console issues

Within the iDRAC virtual console window, if the keyboard or some of the keys are not responding, perform the following checks:

- Ensure on your keyboard that the Scroll Lock button is off.
- In the iDRAC virtual console, ensure that the **Keyboard/Mouse Attach State** is set to **Auto-attached**.

For persistent iDRAC virtual console issues, see Troubleshoot general iDRAC failures on page 80.

iDRAC virtual media issues

The following are solutions to problems that may occur when using iDRAC virtual media:

1. If the iDRAC is stuck, one of these actions may assist with recovery:
   - In the server's **Attached Media** screen, disconnect the iDRAC Remote File Share from the iDRAC browser.
   - Disconnect the iDRAC remote image:

     ```
     racadm -r <dracIP> -u root -p <password> remoteimage -d
     ```

     Where:

     - *<dracIP>* is the iDRAC IP address
     - *<password>* is the password for the server

   Note:
   The default password is Scaleio123.
2. Issues when booting to iDRAC virtual media may cause CPU machine check errors at POST. When this occurs, clear the iDRAC job queue and reset the iDRAC, as described in Troubleshoot general iDRAC failures on page 80.

Rebranding operation fails

When updating the Personality Module on servers using the DTK, the operation may fail and the following error message may appear:

```
UEFI0052 Unable to complete the rebranding operation.
```

To resolve the problem, upload the PM file manually:

1. Contact Customer Support for the VxFlex OS PM file.
2. Follow the instructions in the *Upgrade VxFlex OSGuide* for "Single server firmware and BIOS upgrade to a specific version" to upload the file.

Check the logs for error messages

You can view VxFlex Ready Node Hardware Update Bootable ISO logs after the update script completes.

Procedure

1. Press Alt-F2 to access a user console, and then press Enter.
2. Open the log to check the contents for errors:

   ```
   less /bundleapplicationlogs/apply_components.log
   ```

3. You can also view the script for the Hardware ISO, which is useful in helping to identify and troubleshoot log entries:

   ```
   less /opt/dell/toolkit/systems/drm_files/apply_bundles.sh
   ```

Results

The script attempts to configure several boot order commands, regardless of the hardware configuration of the VxFlex Ready Node server. This allows the script to support multiple hardware platform configurations.
Therefore, it is normal to see these error messages within the logs:

```
Apply R730xd bootorder
[Key=BIOS.Setup.1-1#BiosBootSettings]
RAC1017: Successfully modified the object value and the change is in
       pending state.
       To apply modified value, create a configuration job and reboot
       the system. To create the commit and reboot jobs, use "jobqueue"
       command. For more information about the "jobqueue" command, see RACADM
       help.

ERROR: BOOT016: Input source argument value for the boot device is incorrect or
        not found among the boot devices on the system.

ERROR: BOOT016: Input source argument value for the boot device is incorrect or
        not found among the boot devices on the system.

ERROR: BOOT016: Input source argument value for the boot device is incorrect or
        not found among the boot devices on the system.

[Key=BIOS.Setup.1-1#BiosBootSettings]
RAC1017: Successfully modified the object value and the change is in
       pending state.
:
```

Firmware updates may also display and log the following message:

```
This update is not compatible with your system configuration.
```

```
iDRAC Settings Complete
Proceeding with server BIOS and firmware flash.
Collecting inventory...
Running validation...

This Update Package is not compatible with your system configuration.

Collecting inventory...
Running validation...
```

These firmware are included on the ISO in order to support various VxFlex Ready Node platforms. The log messages do not necessarily indicate a failure.

Additional resources

This section contains information regarding additional resources that may be helpful useful when using the Hardware ISO.

VxFlex OS resources

VxFlex Ready Node deployments have specific guidelines regarding server installation, rack and stack procedures, and power and networking requirements.

The *VxFlex Ready Node Hardware Installation Guide* describes how to install the physical components of a VxFlex Ready Node system. For additional information regarding the VxFlex Ready Node product, documentation, advisories, downloads, and white papers, visit https://support.emc.com/products/42216.

Dell Lifecycle Controller (LC)

At the heart of the VxFlex Ready Node servers' embedded management is the iDRAC with Lifecycle Controller (LC) technology.

This technology allows users to perform useful tasks such as configuring BIOS and hardware settings, deploying operating systems, updating drivers, changing RAID settings, and saving hardware profiles. Together, they provide a robust set of management functions that can be leveraged throughout the entire server lifecycle.

For more information, visit the Lifecycle Controller wiki homepage.

Dell EMC OpenManage Deployment Toolkit (DTK)

The Dell EMC OpenManage Deployment Toolkit (DTK) includes a set of utilities, sample scripts, and sample configuration files that you can use to deploy and configure Dell systems.

You can use the DTK to build script-based and RPM-based installation for deploying large number of systems on a pre-operating system environment in a reliable way, without changing their current deployment processes. Using DTK you can install operating systems on Dell EMC systems in Unified Extensible Firmware Interface (UEFI) mode.

For more information, visit the Dell OpenManagement Deployment Toolkit wiki homepage.

CHAPTER 11

Getting help

This section explains the different resources available for getting help for your system.

- Contacting Dell EMC 86
- Secure Remote Services 86
- Recycling or End-of-Life service information 86

Contacting Dell EMC

Use this topic to learn how to contact Dell EMC.

About this task

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

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

Procedure

1. Go to **https://dell.com/support/incidents-online/contactus**.
2. For customized support:
   a. Type your system **Service Tag** in the **Service Tag** box.
   b. Click **Submit**.
3. For general support:
   a. Select your product category.
   b. Select your product segment.
   c. Select your product.

   The support page that lists the various support categories is displayed.

Secure Remote Services

Secure Remote Services (formerly ESRS) is a secure, two-way connection between Dell EMC products and Dell EMC Customer Support that helps customers avoid and resolve issues up to 73 percent faster. It is completely virtual and offers flexibility for enterprise environments of any size.

Key benefits

- Proactive wellness monitoring and issue prevention
- Automated issue detection, notification and case creation for quicker uptime
- Predictive analytics-based recommendations through MyService360 and product consoles

The Secure Remote Services lifeline is a heartbeat that pulses outbound from the Secure Remote Services gateway to Dell EMC Customer Service in 30-second intervals, providing Dell EMC with connectivity status as well as the status of each product. The heartbeat ensures continuous monitoring, notification, and if necessary, proactive remote troubleshooting to ensure high availability of your products. As a result, you will experience faster resolution and greater uptime.

Recycling or End-of-Life service information

Take back and recycling services are offered for this product in certain countries.

If you want to dispose of system components, visit Dell.com/recyclingworldwide and select the relevant country.