Dell EMC VxFlex Ready Node

Operating System Installation and Configuration Guide for ESXi

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
Supported operating systems and requirements 10
Hardware and operating systems 11
VxFlex OS packages 11
VxFlex OS component requirements 12
VxFlex OS cluster components 12
VxFlex OS Gateway server requirements—VxFlex Ready Node 12
General prerequisites 13
Java requirements 13
Supported Internet browsers 13
Connectivity requirements 13
Disk prerequisites 14
iDRAC Service Module 14
Supported hardware configurations 14
Additional equipment and network resource requirements 14

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 ESXi operating system **25**
ESXi system requirements 26
Driver support for ESXi nodes 26
Rebuild the RAID 1 boot device using the replaced M.2 module 26
Map the ESXi ISO file on a VxFlex Ready Node server 27

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

Chapter 5 Configuring network ports on ESXi hosts **35**
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 - ESXi 38
VxFlex Ready Node R740xd port designations - ESXi 39
Configure network ports on R840 servers 42
R840 server NIC and iDRAC port locations 42
VxFlex Ready Node R840 port designations - ESXi 43

Chapter 6 Installing the Drivers **47**
Install the VxFlex OS drivers on an ESXi server 48

Chapter 7 Preparing Drives 51
Verify the disk controller type 52
Enable PERCCLI on an ESXi-based system 52
Create virtual drives with PERCCLI 53
Verify the storage and disk controller status 54
Create virtual drives for an HDD using PERCCLI on ESXi systems 55
Create virtual drives for SSD with PERCCLI on ESXi systems 56
Ensure virtual drive creation with PERCCLI 56

Chapter 8 Additional configurations 59
Install OpenManage Enterprise 60
Activate the NVIDIA GPU on the server 60
Installing the iDRAC Service Module on an ESXi server 60
Install the iSM using SSH 61
Install the iSM using the vSphere CLI 61
Next steps 62

Chapter 9 Reference material 63
DTK - Hardware Update Bootable ISO 64
Dell EMC OpenManage DRAC Tools (RACADM) 64
Recommended BIOS and firmware settings 66
Troubleshooting the Hardware ISO 68
Additional resources 70

Chapter 10 Getting help 73
Contacting Dell EMC 74
Secure Remote Services 74
Recycling or End-of-Life service information 74

Figures

1 R640 and R740xd PCI slots, integrated NICs and iDRAC port locations..... 37
2 R840 PCI slots, integrated NICs, and iDRAC port locations.....42

Tables

1 ESXi....10
2 Windows, except Core systems (see note) 10
3 ESXi....11
4 Windows....11
5 VxFlex OS management IP network....30
6 VxFlex OS data IP network for Subnet #1....31
7 VxFlex OS data IP network for Subnet #2....31
8 VxFlex OS management IP network....32
9 VxFlex OS data IP network for Subnet #1....32
10 VxFlex OS data IP network for Subnet #2....32
11 R640 server configuration for 10/25GbE, 1CPU, SFP+/SFP28....38
12 R640 server configuration for 10/25GbE, 2CPU, SFP+/SFP28....38
13 R740xd server configuration for 10/25 GbE, 32xSAS, SFP+/SFP28....39
14 R740xd server configuration for 10/25GbE, 12xNVMe, SFP+/SFP28....40
15 R740xd server configuration for 25GbE, 3xGPU (M60)/ 2xGPU (M10), SFP28....40
16 R740xd server configuration for 10GbE, 3xGPU (M60)/ 2xGPU (M10), SFP+....41
17 R840 server configuration for 2x10GbE+2GbE rNDC, 2x10/25GbE PCIe slot 3, 2x10/25GbE PCIe slot 6 – SAS or NVME, no GPU....43
18 R840 server configuration for 2x25GbE rNDC, 2xGPU – SAS or NVME....44
19 Hardware ISO configuration settings....67

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
- Supported operating systems and requirements 10
- Hardware and operating systems 11
- VxFlex OS component requirements 12
- General prerequisites 13
- Supported hardware configurations 14
- Additional equipment and network resource requirements 14

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

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 ESXi

Component	Requirement
Supported OS versions	ESXi 6.7 U3 and earlier ESXi 6.5 U3 and earlier Managed by vCenter 6.7
Additional requirements	For vCenter v6.7, PowerCLI v6.5 must be used.

Connecting to a vCenter enables you to add hosts to a VcDatacenter folder.

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

Table 2 Windows, except Core systems (see note) (continued)

Component	Requirement
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," as follows.

VxFlex Ready Node R840

Table 3 ESXi

Component	Requirement
Supported OS versions	ESXi 6.7 U3 Managed by vCenter 6.7
Additional requirements	For vCenter v6.7, PowerCLI v6.5 must be used.

Connecting to a vCenter enables you to add hosts to a VcDatacenter folder.

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 12.

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.

For converged deployments on ESXi servers, add the nodes to vCenter.

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.

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 for Storage VMs (SVMs).
- 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.
- All root and administrator passwords that you plan to set or use on the SVM.

> **Note:** VxFlex Ready Node servers shipped with a factory-installed ESXi image are configured with root as the username and either the service tag number or Scaleio123 as the password.

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.

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 64.

For additional information regarding the Hardware ISO, see DTK - Hardware Update Bootable ISO on page 64.

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 64.

For more information about the Hardware ISO, see DTK - Hardware Update Bootable ISO on page 64.

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 68 .

    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 ESXi operating system

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

- ESXi system requirements 26
- Driver support for ESXi nodes 26
- Rebuild the RAID 1 boot device using the replaced M.2 module 26
- Map the ESXi ISO file on a VxFlex Ready Node server 27

ESXi system requirements

To deploy on ESXi 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.

- Ensure that the M.2 device on the BOSS device is configured as a RAID 1 device.
- During OS installation, use Scaleio123 as the password for the root or administrator username . (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.

Driver support for ESXi nodes

Ensure that correct drivers are installed in your system.

The best source for supported storage drivers is the Dell EMC customized ISO for ESXi 6.7 EP 05. If this cannot be used, consult the VMware support matrix for Dell EMC products on the VMware website.

For the Dell EMC HBA330 mini controller, both the customized and the release versions of ESXi 6.7 EP 05 have the minimum supported driver version.

Note: Before enabling DirectPath, ensure that the correct driver is installed.

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 ESXi ISO file on a VxFlex Ready Node server

Map the 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.

   Note: If the VxFlex Ready Node server was shipped with the operating system factory installed, proceed directly to Step 7.

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** > **Reset System (warm boot)** to display the list of keyboard keys with related functionality.
   b. Press F11. The **Boot Manager Main Menu** pane is displayed.
   c. Select **One-shot BIOS Boot Menu**. The **Select Legacy Boot Option** pane is displayed.
   d. Select **Virtual Optical Drive**.

   The server boots using the OS ISO file and finishes the installation.
7. If the node contains a factory-installed OS, the datastores will not be visible in the vSphere Web Client. The BOSS card is partitioned, and the OS is installed in the visible partition. To use the remaining storage on the BOSS card for the datastore, perform the following steps:

a. In the vSphere Web Client, select **Storage** > **New Datastore** > **Create New VMFS datastore**.

b. Select the empty partition for the datastore.

   Datastore 1 is created on the empty BOSS card.

CHAPTER 4

Network architecture and physical connectivity

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

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

Networking connectivity architecture and cabling best practice

The following information describes connectivity architecture, cabling best practice information, and cable connection examples from typical VxFlex OS configurations in ESXi environments in order 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*.

On ESXi servers, VLANs can be used as long as the broadcast domain is the same between the Discovery Server and the ESXi management port.

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.

When using several VLANs on the same port, the port must to be configured as Trunk, and you must assign a VLAN ID on each vSwitch on the ESXi host. This configuration can be performed on a VxFlex Ready Node server after deployment is complete by updating the VM network and relevant VMKernel port to use the VLAN ID set by the switch manager. This is similar to any other VLAN implementation.

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 ESXi 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 14 earlier in this guide.

Note: The AMS or discovery server must reside on the same Layer 2 physical network (or VLAN) as the ESXi management Port (Broadcast Domain). The Ethernet broadcast protocol is used during the initial detection and deployment process.

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. SVM_MGMT_IP = VxFlex OS storage VM management IP address 3. BMC_MGMT_IP = iDRAC IP address for each node
AMS server\Proxy IP address	The IP address used by the AMS server\proxy software to manage the VxFlex Ready Node servers and logical entities (recommended to reside on the same subnet, unless the proxy is used for Discovery operations)

Table 5 VxFlex OS management IP network (continued)

Item	Description
The formula to calculate IP address subnet pool or subnet needs is: N*Node_MGMT_IP+N*SVM_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.

Note: If the AMS is not on the same subnet as the SVM management IP addresses, the Discovery server used as a proxy must reside on the same switch network broadcast domain (Physical or VLAN), in order to allow seamless detection and MGMT IP address configuration of the SVM and ESXi IP addresses of all nodes in the system.

The VxFlex OS data IP network requires two separate subnets.

Note: Multiple IP subnets used for the VxFlex OS data network cannot be on the same subnet. For more information, see this VMware limitation in the following link: http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2010877 .

Table 6 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. SVM_DATA1_IP = VxFlex OS Storage VM Management IP addresses 3. 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+N*SVM_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. SVM_DATA2_IP = VxFlex OS Storage VM Management IP addresses 3. 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+N*SVM_DATA2_IP +MDM_Cluster_Virtual_IP_DATA2		

Note: The subnet mask and default gateway of the data networks cannot be modified after initial configuration. If you plan to have 1024 nodes in your VxFlex OS system, plan for a subnet allowing 1025 IP addresses for SVM, and

1024 IP addresses for hosts per data network. IP address pools may be modified to allow additional IP addresses to be added to new nodes.

IP addresses for R840 servers

Prepare IP addresses in your network for the VxFlex Ready Node servers in ESXi 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 14 earlier in this guide.

Table 8 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. SVM_MGMT_IP = VxFlex OS storage VM management IP address 3. 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*SVM_MGMT_IP+N*BMC_MGMT_IP	

The VxFlex OS data IP network requires two separate subnets.

Note: Multiple IP subnets used for the VxFlex OS data network cannot be on the same subnet. For more information, see this VMware limitation in the following link: http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2010877 .

Table 9 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. SVM_DATA1_IP = VxFlex OS Storage VM Management IP addresses 3. 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+N*SVM_DATA1_IP +MDM_Cluster_Virtual_IP_DATA1		

Table 10 VxFlex OS data IP network for Subnet #2

Item	Description	Comments
N	Number of nodes	

Table 10 VxFlex OS data IP network for Subnet #2 (continued)

Item	Description	Comments
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. SVM_DATA2_IP = VxFlex OS Storage VM Management IP addresses 3. 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+N*SVM_DATA2_IP +MDM_Cluster_Virtual_IP_DATA2		

Note: The subnet mask and default gateway of the data networks cannot be modified after initial configuration. Therefore it is important that you plan accordingly. For example, if you plan 1,024 nodes for the VxFlex OS system, plan for a subnet allowing 1,025 IP addresses for SVM, and 1024 IP addresses for hosts per data network. IP address pools may be modified to allow additional IP addresses to be added to new nodes.

CHAPTER 5

Configuring network ports on ESXi hosts

This section describes how to configure the ports on a VxFlex Ready Node ESXi host.

- Configure network ports on R640 or R740xd servers 36
- Configure network ports on R840 servers 42

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 iDRAC port locations

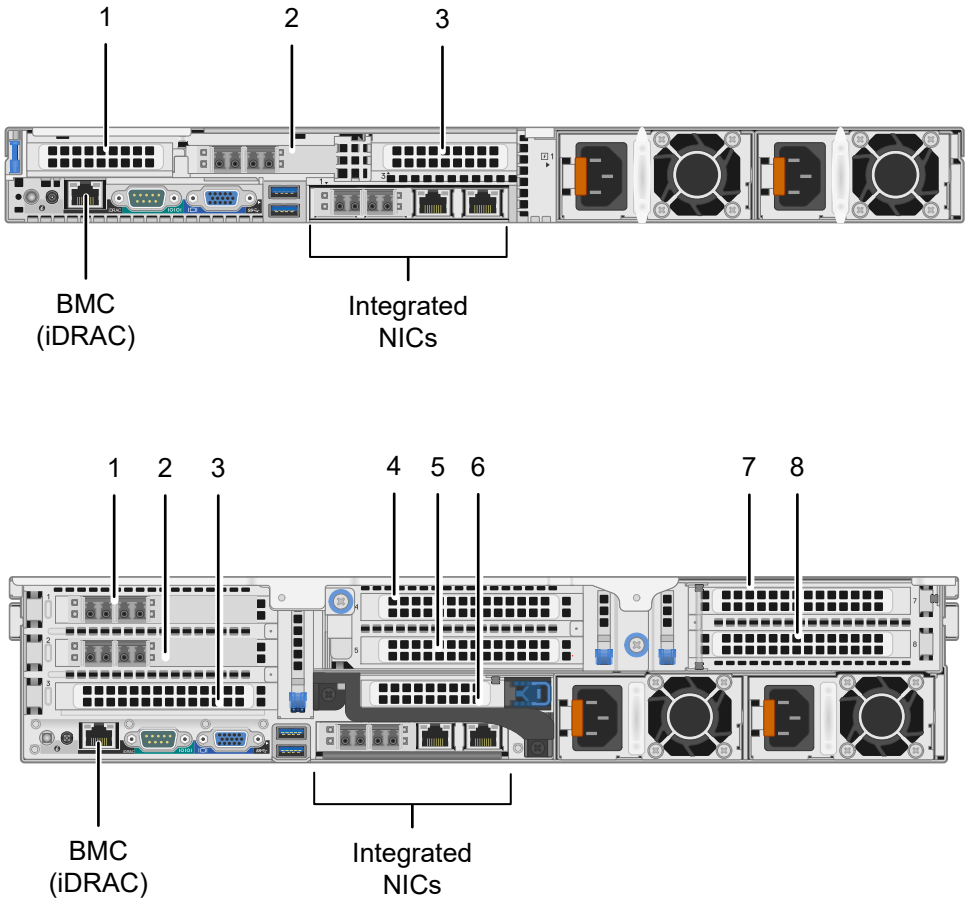

Data cables are connected to two switches for high availability, via the nodes' 10 GbE, 25 GbE, or 100 GbE ports. Management interfaces are connected to a switch on a separate management network using onboard and iDRAC ports.

In ESXi environments, the AMS/discovery server must be on the same Layer 2 network as the ESXi management ports. This allows the use of the Broadcast protocol for initial IP address configuration.

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.

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.

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.

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.

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 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 - ESXi

Connect the cables and configure the system in an ESXi-based system as shown in the following tables, depending on the server type and architecture, when relevant.

R640 servers

Note: In the following tables use SFP+ for 10 GbE, and SFP28 for 25 GbE ports.

Table 11 R640 server configuration for 10/25GbE, 1CPU, 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
Port	iDRAC	vmnic0	vmnic1	vmnic4	vmnic5	vmnic2	N/A
PCI slot	N/A	2:0	2:1	N/A	N/A	N/A	N/A
Physical Port on node	iDRAC (remote KVM)	PCI 10/25 GbE (Left)	PCI 10/25 GbE (Right)	Onboard 25 GbE (Right)	Onboard 25 GbE (Left)	Onboard 25 GbE (Left) - The same as for VxFlex OS Data2	N/A
Speed	1 GbE	10/25 GbE	10/25 GbE	25 GbE	25 GbE	25 GbE	N/A
vSwitch	N/A	sio_vSwitch1	User allocated	User allocated	vSwitch1	vSwitch0	N/A
VMkernel port		vmk2	User configurable	User configurable	vmk3	vmk0	N/A
Physical Switch	Mgmt Switch	Switch 1	Switch 1	Switch 2	Switch 2	Mgmt Switch/ Switch 2	N/A

Table 12 R640 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
Port	iDRAC	vmnic0	vmnic1	vmnic4	vmnic5	vmnic2	N/A
PCI slot	N/A	N/A	N/A	N/A	N/A	N/A	N/A

Table 12 R640 server configuration for 10/25GbE, 2CPU, SFP+/SFP28 (continued)

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
Physical Port on node	iDRAC (remote KVM)	PCI 10/25 GbE (Left)	PCI 10/25 GbE (Right)	PCI 10/25 GbE (Right)	PCI 10/25 GbE (Left)	Onboard 1 GbE (Left)	N/A
Speed	1 GbE	10/25 GbE	10/25 GbE	10/25 GbE	10/25 GbE	1 GbE	N/A
vSwitch	N/A	sio_vSwitch1	User allocated	User allocated	sio_vSwitch2	vSwitch0	N/A
VMkernel port		vmk2	User configurable	User configurable	vmk3	vmk0	N/A
Physical Switch	Mgmt Switch	Switch 1	Switch 1	Switch 2	Switch 2	Mgmt Switch	N/A

Note: This configuration supports more than one PCIe-based NIC in slots 2 and 3.

VxFlex Ready Node R740xd port designations - ESXi

Connect the cables and configure the system in an ESXi-based system as shown in the following tables, depending on the server type and architecture, when relevant.

R740xd servers

Note: In the following tables use SFP+ for 10 GbE, and SFP28 for 25 GbE ports.

Table 13 R740xd server configuration for 10/25 GbE, 32xSAS, 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
Port	iDRAC	vmnic0	vmnic1	vmnic4	vmnic5	vmnic2	vmnic3
PCI slot	N/A	1:0	1:1	2:1	2:0	N/A	N/A
Physical Port on node	iDRAC (remote KVM)	PCI 10/25 GbE (Left)	PCI 10/25 GbE (Right)	PCI 10/25 GbE (Right)	PCI 10/25 GbE (Left)	Onboard 1 GbE (Left)	N/A
Speed	1 GbE	10/25 GbE	10/25 GbE	10/25 GbE	10/25 GbE	1 GbE	N/A
vSwitch	N/A	sio_vSwitch1	User allocated	User allocated	sio_vSwitch2	vSwitch0	N/A
VMkernel port		vmk2	User configurable	User configurable	vmk3	vmk0	N/A

Table 13 R740xd server configuration for 10/25 GbE, 32xSAS, SFP+/SFP28 (continued)

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
Physical Switch	Mgmt Switch	Switch 1	Switch 1	Switch 2	Switch 2	Mgmt Switch	N/A

Note: The above configuration supports more than one PCIe slot-based NIC.

Table 14 R740xd server configuration for 10/25GbE, 12xNVMe, 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
Port	iDRAC	vmnic0	vmnic1	vmnic4	vmnic5	vmnic2	vmnic3
PCI slot	N/A	2:0	2:1	7:1	7:0	N/A	N/A
Physical Port on node	iDRAC (remote KVM)	PCI 10/25 GbE (Left)	PCI 10/25 GbE (Right)	PCI 10/25 GbE(Right)	PCI 10/25 GbE (Left)	Onboard 1 GbE (Left)	N/A
Speed	1 GbE	10/25 GbE	10/25 GbE	10/25 GbE	10/25 GbE	1 GbE	N/A
vSwitch	N/A	sio_vSwitch1	User allocated	User allocated	sio_vSwitch2	vSwitch0	N/A
VMkernel port		vmk2	User configurable	User configurable	vmk3	vmk0	N/A
Physical Switch	Mgmt Switch	Switch 1	Switch 1	Switch 2	Switch 2	Mgmt Switch	N/A

Note: The above configuration supports more than one PCIe slot-based NIC.

Table 15 R740xd server configuration for 25GbE, 3xGPU (M60)/ 2xGPU (M10), SFP28

Description	iDRAC Monitoring	VxFlex OS Data1/ VxFlex OS Management Network	VxFlex OS Data2	Not in Use
Port	iDRAC	vmnic1	vmnic0	vmnic3
PCI slot	N/A	N/A	N/A	N/A
Physical Port on node	iDRAC (remote KVM)	Onboard 25 GbE (Left)	Onboard 25 GbE (Right)	None
Speed	1 GbE	25 GbE	25 GbE	None
vSwitch	N/A	vSwitch0	sio_vSwitch1	N/A
VMkernel port		vmk2	vmk3	N/A

Table 15 R740xd server configuration for 25GbE, 3xGPU (M60)/ 2xGPU (M10), SFP28 (continued)

Description	iDRAC Monitoring	VxFlex OS Data1/ VxFlex OS Management Network	VxFlex OS Data2	Not in Use
Physical Switch	Mgmt Switch	Data Switch 1	Data Switch 2	N/A

Note: The above configuration supports only two 25 GbE onboard NICs in three-GPU configuration. In configurations with only two GPUs, another PCIe-based NIC is supported: one in slot 3 and/or one in slot 4.

Table 16 R740xd server configuration for 10GbE, 3xGPU (M60)/ 2xGPU (M10), SFP+

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
Port	iDRAC	vmnic0	vmnic1	vmnic3	vmnic2	vmnic0	N/A
PCI slot	N/A	N/A	N/A	N/A	N/A	N/A	N/A
Physical Port on node	iDRAC (remote KVM)	Onboard 10 GbE (far Left)	Onboard 10 GbE (second from Left)	Onboard 10 GbE (far Right)	Onboard 10 GbE (second from Right)	Onboard 10 GbE (far Left) - The same as for VxFlex OS Data1	None
Speed	1 GbE	10 GbE	10 GbE	10 GbE	10 GbE	10 GbE	None
vSwitch	N/A	vSwitch0	User allocated	User allocated	sio_vSwitch1	vSwitch0	N/A
VMkernel port		vmk2	User configurable	User configurable	vmk3	vmk0	N/A
Physical Switch	Mgmt Switch	Data Switch 1	Data Switch 1	Data Switch 2	Data Switch 2	Mgmt/Data1 Switch	N/A

Note: In configurations with only two GPUs, another PCIe-based NIC is supported in either slot 3 or slot 4.

Configure the Management IP address - ESXi

Configure the Management IP address using the ESXi console.

Procedure

1. Open the ESXi KVM console.
2. From the console, press F2 to open the **Customize System/View Logs** menu.

   The **Authentication Required** window is displayed.
3. Type the server password defined during installation: ScaleIO123

   Wait a few seconds for the **System Customization** window to appear; if it does not, press F2 again.
4. Select **Configure Management Network**, and press Enter.
5. In the **Configure Management Network** window, select **Network Adapters**, and press Enter.
6. Using the Up/Down arrow keys, scroll to the appropriate vmnic for your hardware (the management port vmnic differs per server type: 1U1N, 2U1N Storage-Only, 2U1N Converged, etc.)

7. Select the desired vmnic by pressing the Space bar.
8. Click **OK** to approve the selection.
9. In the **Configure Management Network** window, scroll to **IP Configuration**, and press Enter.
10. In the **IP Configuration** popup window, change the selected option to **Set static IP address and network configuration**, using the Up/Down arrow keys for navigation and the Space bar for selection.
11. Use the Up/Down arrow keys to move to the IP address fields, and type the node's Management IP Address, Subnet Mask, and Default Gateway.
12. Press Enter to apply the IP address changes.
13. In the **Configure Management Network** window, press Esc to exit the window. When the confirmation popup window appears, press Y to approve and apply the configuration.

Results

Configuration of the Management IP address is now complete. The node can now be connected to a vCenter. Once connected, validate that the IP address configuration performed in these steps is correct.

Note: Data IP addresses, vSwitches, and all Storage VM IP addresses are configured during the plug-in deployment flow.

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 2 R840 PCI slots, integrated NICs, and iDRAC port locations

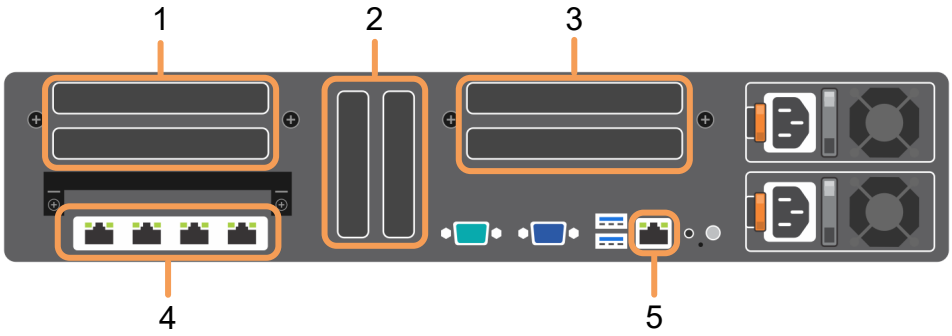

1. 2x FH PCIe slots
2. 2x LP PCIe slots
3. 2x FH PCIe slots
4. Integrated NIC/rNDC ports
5. iDRAC

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 port designations - ESXi

Connect the cables and configure the system in an ESXi-based system as shown in the following tables, depending on the server type and architecture, when relevant.

The tables below show the most common correlation between the physical NIC port location and the logical configuration in VMware. When using the same slot locations as shown below, all settings but one setting are static. The vmnic number is the dynamic setting and is determined by the VMware bus scan and boot logic. This number cannot be set persistently.

If you need to build your own table of reference, correlate the ESXi physical ports to the vmnic logical port name.

R640

R740xd

R840

Note: In the following tables use SFP+ for 10 GbE, and SFP28 for 25 GbE ports.

Table 17 R840 server configuration for 2x10GbE+2GbE rNDC, 2x10/25GbE PCIe slot 3, 2x10/25GbE PCIe slot 6 – SAS or NVME, no GPU

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
Port	iDRAC	vmnic5	vmnic4	vmnic7	vmnic6	vmnic0	vmnic1/2/3
PCI slot	N/A	4:0	4:1	6:0	6:1	N/A	N/A
Physical Port on node	iDRAC (remote KVM)	PCI 10/25 GbE (Left)	PCI 10/25 GbE (Right)	PCI 10/25 GbE (Right)	PCI 10/25 GbE (Left)	Onboard 1 GbE (Left)	N/A

Table 17 R840 server configuration for 2x10GbE+2GbE rNDC, 2x10/25GbE PCIe slot 3, 2x10/25GbE PCIe slot 6 – SAS or NVME, no GPU (continued)

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
Speed	1 GbE	10/25 GbE	10/25 GbE	10/25 GbE	10/25 GbE	1 GbE	N/A
vSwitch	N/A	sio_vSwitch	User allocated	User allocated	sio_vSwitch2	vSwitch0	N/A
VMkernel port		vmk1	User configurable	User configurable	vmk2	vmk0	N/A
Physical Switch	Mgmt Switch	Switch 1	Switch 1	Switch 2	Switch 2	Mgmt Switch	N/A

Note: The above configuration supports more than one PCIe slot-based NIC.

Table 18 R840 server configuration for 2x25GbE rNDC, 2xGPU – SAS or NVME

Description	iDRAC Monitoring	VxFlex OS Data1/ VxFlex OS Management Network	VxFlex OS Data2	Not in Use
Port	iDRAC	vmnic1	vmnic0	vmnic3
PCI slot	N/A	N/A	N/A	N/A
Physical Port on node	iDRAC (remote KVM)	Onboard 25 GbE (Left)	Onboard 25 GbE (Right)	None
Speed	1 GbE	25 GbE	25 GbE	None
vSwitch	N/A	vSwitch0	sio_vSwitch1	N/A
VMkernel port		vmk2	vmk3	N/A
Physical Switch	Mgmt Switch	Data Switch 1	Data Switch 2	N/A

Note: The above configuration supports only two 25 GbE onboard NICs in two-GPU configuration. In configurations with only two GPUs, another PCIe-based NIC is supported: one in slot 3 and/or one in slot 4.

Configure the Management IP address - ESXi

Configure the Management IP address using the ESXi Direct Console User Interface (DCUI).

Procedure

1. Open the ESXi DCUI.
2. From the console, press F2 to open the **Customize System/View Logs** menu.

   The **Authentication Required** window is displayed.
3. Type the server password defined during installation: ScaleIO123

   Wait a few seconds for the **System Customization** window to appear; if it does not, press F2 again.
4. Select **Configure Management Network**, and press Enter.
5. In the **Configure Management Network** window, select **Network Adapters**, and press Enter.

6. Using the Up/Down arrow keys, scroll to the appropriate vmnic for your hardware (the management port vmnic differs per server type: 1U1N, 2U1N Storage-Only, 2U1N Converged, etc.)
7. Select the desired vmnic by pressing the Space bar.
8. Click **OK** to approve the selection.
9. In the **Configure Management Network** window, scroll to **IP Configuration**, and press Enter.
10. In the **IP Configuration** popup window, change the selected option to **Set static IP address and network configuration**, using the Up/Down arrow keys for navigation and the Space bar for selection.
11. Use the Up/Down arrow keys to move to the IP address fields, and type the node's Management IP Address, Subnet Mask, and Default Gateway.
12. Press Enter to apply the IP address changes.
13. In the **Configure Management Network** window, press Esc to exit the window. When the confirmation popup window appears, press Y to approve and apply the configuration.

Results

Configuration of the Management IP address is now complete. The node can now be connected to a vCenter. Once connected, validate that the IP address configuration performed in these steps is correct.

Note: Data IP addresses, vSwitches, and all Storage VM IP addresses are configured during the plug-in deployment flow.

CHAPTER 6

Installing the Drivers

The following topics contain information regarding VxFlex OS drivers.

- Install the VxFlex OS drivers on an ESXi server 48

Install the VxFlex OS drivers on an ESXi server

Dell EMC recommends that you install all applicable VxFlex OS drivers on the ESXi 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 ESXi drivers that correspond to the hardware components installed on the server.
4. Using an application such as `winscp`, copy the driver files to a temporary folder on the VxFlex OS server.
5. Log in to the ESXi host.
6. Get the list of drivers currently installed on the server:

   ```
   esxcli software vib list
   ```

7. Compare the version for each `vib` in the `Zip` file to the version installed on the system, and make a note of the drivers that need to be updated.
8. If none of the drivers need to be updated, skip the rest of this procedure. Otherwise:
   - For each driver being upgraded, run:

     ```
     esxcli software vib install -d <driver_file>
     ```

   - For each driver being downgraded:
     a. Remove the current driver from the server:

     ```
     esxcli software vib remove -d <driver_file>
     ```

     Note: If the output instructs you to reboot the server at this point, ignore it.

     b. Install the correct driver version:

     ```
     esxcli software vib install -d <driver_file>
     ```

9. After all the new drivers are installed, reboot the ESXi host:

   ```
   esxcli system shutdown reboot -r 'replace driver'
   ```

10. After the system has rebooted, view the results of the previous steps:

    ```
    esxcli software profile get
    ```

11. Verify that the correct versions of drivers are running:
    a. Get the list of drivers currently installed on the server:

    ```
    esxcli software vib list
    ```

    b. Get the information for each driver that was updated:

    ```
    esxcli software vib get -n <driver_name>
    ```

CHAPTER 7

Preparing Drives

For R640 and R740xd systems, if necessary, prepare the drives so that they can be added to VxFlex OS SDS devices.

- Verify the disk controller type 52
- Enable PERCCLI on an ESXi-based system 52
- Create virtual drives with PERCCLI 53

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, proceed to the next section and perform any additional configurations required before deploying VxFlex OS.
- 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 on an ESXi-based system

Before creating virtual disks, enable the PERCCLI disk utility. You may also need to install the disk driver.

Before you begin

Ensure that you have:

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

Procedure

1. Download PERCCLI from https://downloads.dell.com/FOLDER05281643M/1/perccli_vmware_KFT4T_A07_7.529.00.zip.
2. Extract the `ZIP` file.
3. Copy the installation file, `vmware-esx-perccli-7.529.00.vib`, to `/tmp` on the server.

4. Run the installation command on the server:

```
esxcli software vib install -f -v /tmp/vmware-esx-perccli-7.529.00.vib
```

5. Before running PERCCLI commands, change your working directory to the utility's binary location:

```
cd /opt/lsi/perccli/
```

Results

PERCCLI can now be used either from the ESXi console or via SSH.

Note: PERCCLI commands are run in the following format:

```
./perccli command
```

For example, to view help, run:

```
./perccli -h
```

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:

```
cd /opt/lsi/perccli/; ./perccli 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

Node contains	Caching options
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.

Before you begin

Before using PERCCLI, ensure that your working directory is `/opt/lsi/perccli/`. If it is not, run:

```
cd /opt/lsi/perccli/
```

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/lsi/perccli/perccli /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/lsi/perccli/perccli /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/lsi/perccli/perccli /c0/vall del
```

The VDs are removed.

5. Set the controller card boot parameter:

```
/opt/lsi/perccli/perccli /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 on ESXi systems

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

Before you begin

Ensure that the PERCCLI disk utility is installed.

Before using PERCCLI, ensure that your working directory is `/opt/lsi/perccli/`. If it is not, run: `cd /opt/lsi/perccli/`

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

Example:

```
/opt/lsi/perccli/perccli /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 with PERCCLI on ESXi systems

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.

Before using PERCCLI, ensure that your working directory is `/opt/lsi/perccli/`. If it is not, run:

```
cd /opt/lsi/perccli/
```

About this task

Use this procedure both for SSD for storage and SSD to be used to provide caching acceleration. The VDs created on SSD 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/lsi/perccli/perccli /c0 add vd type=raid0 drives=<EID>:<Slt> direct wt nora
   ```

   Example:

   ```
   /opt/lsi/perccli/perccli /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.

Before you begin

Before using PERCCLI, ensure that your working directory is `/opt/lsi/perccli/`. If it is not, run:

```
cd /opt/lsi/perccli/
```

Procedure

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

```
/opt/lsi/perccli/perccli /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/lsi/perccli/perccli /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)

CHAPTER 8

Additional configurations

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

- Install OpenManage Enterprise 60
- Activate the NVIDIA GPU on the server 60
- Installing the iDRAC Service Module on an ESXi server 60
- Next steps 62

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.

Activate the NVIDIA GPU on the server

In deployments where an NVIDIA GRID GPU card is configured on the VxFlex Ready Node server, activate the NVIDIA GRID Enterprise software.

Before you begin

Verify that:

- A valid GRID Enterprise software license exists. GRID Enterprise software is licensed by the number of concurrent users and must have been purchased with the GPU-enabled VxFlex Ready Node servers. The following editions are available for NVIDIA GRID Software license:
  - Virtual Applications (vAPP)
  - Virtual PC (vPC)
  - Virtual Workstation (vWS)

Procedure

1. Locate the email with the PAK and registration link to the NVIDIA licensing portal.
2. Log in to the licensing portal to access the Product Activation Key and driver for your NVIDIA card.
3. Follow the instructions for enabling and setting up the GPU card.

Installing the iDRAC Service Module on an ESXi server

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.

The iSM is available as a Zip archive for installing on systems running the VMware ESXi operating system. You can install the iSM via SSH or using the vSphere CLI.

The Driverpack file for the approved iSM package is available from https://support.emc.com/products/42216.

Install the iSM using SSH

Install the iDRAC Service Module (iSM) software on the ESXi host via SSH.

Procedure

1. Download the Driverpack file for the iSM package from https://support.emc.com/products/42216.
2. Unzip the file to a local directory on the system.
3. Upload the VIB to the ESXi host's datastore.
4. Log in to the vCenter via the vSphere Web Client.
5. For each guest operating system, select the guest OS, and from the **Actions** > **Power** menu, select **Shut Down Guest OS**.
6. When all the guest operating systems are powered off, right-click the ESXi host and select **Enter Maintenance Mode**.
7. Establish an SSH connection to the ESXi host and log in using the root username and password.
8. Install the iSM software:

   ```
   esxcli software vib install - d /vmfs/volumes/datastore/<iDRAC Service Module file>
   ```

   The command output tells you whether the installation succeeded.
9. Right-click the ESXi host and select **Power** > **Reboot**.
10. From the iDRAC web interface, open the **iDRAC settings** > **Overview** page and ensure that the iDRAC Service Module status is **Running**

Install the iSM using the vSphere CLI

Install the iDRAC Service Module (iSM) software on the ESXi host using the vSphere CLI.

Procedure

1. Download the Driverpack file for the iSM package from https://support.emc.com/products/42216.
2. Unzip the file to a local directory on the system.
3. Log in to the vCenter via the vSphere Web Client.
4. For each guest operating system, select the guest OS, and from the **Actions** > **Power** menu, select **Shut Down Guest OS**.
5. When all the guest operating systems are powered off, right-click the ESXi host and select **Enter Maintenance Mode**.
6. If you are using the vSphere CLI from a Windows server, navigate to the directory where you have installed the vSphere CLI utilities.

   Note: If you are using the vSphere CLI on a Linux server, perform the command from any directory.
7. Install the iSM software:

   ```
   esxcli --server <host_IP> software vib install - d /var/log/vmware/<iDRAC Service Module
   file>
   ```
8. At the prompt, type the root username and password of the ESXi host.

   The command output tells you whether the installation succeeded.

9. Right-click the ESXi host and select **Power** > **Reboot**.
10. Optionally view the information about the installed VIBs:

```
esxcli --server <IP> software vib get
```

or

```
software vib list
```

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

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.

NVDIMM configuration on ESXi nodes is performed after deploying VxFlex OS. Complete instructions are included in the *Dell EMC Configure and Customize VxFlex OS Guide*.

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 9

Reference material

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

- DTK - Hardware Update Bootable ISO 64

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 65 - 9 on page 65 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 68 .

    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 65) 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 19 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 68.

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 68.

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 10

Getting help

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

- Contacting Dell EMC 74
- Secure Remote Services 74
- Recycling or End-of-Life service information 74

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.