Dell EMC VxFlex Ready Node 14G

Version 2.x

AMS Management Server Maintenance Guide - 14G

P/N 302-004-657
REV 02
June 2018

- Preface 2
- Overview 3
- Before you begin 4
- PowerEdge R640 server technical specifications 4
- Install a PowerEdge R640 server in a cabinet 5
- Supported hardware configurations 14
- Set up the BMC (iDRAC) IP address and BIOS 14
- Verify the status of the system hardware, storage controller, and disks - 14G servers 16
- System requirements 17
- Open the KVM console - 14G servers 18
- Map the ISO file on a VxFlex Ready Node 14G server 19
- Accept the Kylin end-user license agreement (EULA) 20
- Server node port designations 20
- Correlate the PCI slot locations and interface names 21
- Configure the ports (Linux-SLES) 24
- Validate the AMS 25
- Upgrade the AMS version 26

DELL EMC

Preface

As part of an effort to improve its product lines, Dell EMC periodically releases revisions of its software and hardware. Therefore, some functions described in this document might not be supported by all versions of the software or hardware currently in use. The product release notes provide the most up-to-date information on product features.

Contact your Dell EMC technical support professional if a product does not function properly or does not function as described in this document.

Note

This document was accurate at publication time. Go to Dell EMC Online Support (https://support.emc.com) to ensure that you are using the latest version of this document.

Previous versions of Dell EMC VxFlex OS were marketed under the name Dell EMC ScaleIO.

Similarly, previous versions of Dell EMC VxFlex Ready Node were marketed under the name Dell EMC ScaleIO Ready Node.

References to the old names in the product, documentation, or software, etc. will change over time.

Note

Software and technical aspects apply equally, regardless of the branding of the product.

Related documentation

The release notes for your version includes the latest information for your product.

The following Dell EMC publication sets provide information about your VxFlex OS or VxFlex Ready Node product:

- VxFlex OS software (downloadable as VxFlex OS Software <version> Documentation set)
- VxFlex Ready Node with AMS (downloadable as VxFlex Ready Node with AMS Documentation set)
- VxFlex Ready Node no AMS (downloadable as VxFlex Ready Node no AMS Documentation set)
- VxRack Node 100 Series (downloadable as VxRack Node 100 Series Documentation set)

You can download the release notes, the document sets, and other related documentation from Dell EMC Online Support.

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

Where to get help

Dell EMC support, product, and licensing information can be obtained as follows:

Product information

For documentation, release notes, software updates, or information about Dell EMC products, go to Dell EMC Online Support at https://support.emc.com.

Technical support

Go to Dell EMC Online Support and click Service Center. You will see several options for contacting Dell EMC Technical Support. Note that to open a service request, you must have a valid support agreement. Contact your Dell EMC sales representative for details about obtaining a valid support agreement or with questions about your account.

Your comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Send your opinions of this document to techpubcomments@emc.com.

Overview

This topic provides an overview of the steps required for the physical installation of the hardware and the deployment of the VxFlex OS Management Server Software VxFlex OS AMS 2.x image.

FRU replacement of the VxFlex OS AMS Management Server is the same as installation except with an additional recovery step at the end of the tasks for installation. Steps include:

- Gathering required information
- Installing the physical server in a cabinet.
- Deploying the operating system.

Before you begin

This topic describes prerequisites for installing the VxFlex Ready Node AMS on a physical management server.

Before beginning the procedures in this guide, ensure that you have satisfied the following prerequisites:

- Ensure that you have a PC or laptop computer with a functional network port and an available IP address (and the subnet and gateway) on the management network.
- Ensure that you have a VGA tool kit to allow console connection from a laptop computer to a server, or a computer screen and keyboard connection to the customer's rack.
- Java 1.8 latest must be installed on the laptop computer, for the purposes of using the BMC (iDRAC) HTTP client.
- In your Internet browser, ensure that pop-ups are enabled (or enabled just for the KVM).
- Download the server image file for this version. For example: `VxFlex Ready Node 2.6 AMS Server Image Download`
  You can download packages from the online support site (https://support.emc.com/products/42216).
- Ensure that you obtain the customer IP address, subnet mask, and gateway for the access to the VxFlex Ready Node AMS server SLES 12.2 operating system.
- Ensure that you know the IP address of the BMC (iDRAC)/IPMI port on the server (defined during the initial deployment process).

You can enter the gathered data in the following table:

Item	Data needed	Value
BMC (iDRAC)	IP address	
	Subnet mask	
	Gateway IP address	
	VLAN ID	
	Password	
Host	IP address	
	Subnet mask	
	Gateway IP address	
	VLAN ID	
	Root password	

PowerEdge R640 server technical specifications

Technical specifications for the PowerEdge R640 server.

This section describes the technical, power, and environmental specifications for the PowerEdge R640 server.

About your system

The Dell PowerEdge R640 is a rack server that supports up to two processors based on the Intel Xeon SP-6126 V1 processor family and up to 24 DIMMs.

Table 1 Supported configurations

PowerEdge R640 Systems	Configurations
Eight-hard-drive systems	Up to eight 2.5-inch hard drives on the front panel
Four-hard-drive systems	Up to four 3.5-inch hard drives on the front panel
Ten-hard-drive systems	Up to ten 2.5-inch hard drives on the front panel
Two-hard-drive systems	Optional support up to two 2.5-inch hard drives on the back panel

Technical specifications

For detailed technical specifications on the PowerEdge R640 system, see http://www.dell.com/support/home/us/en/04/product-support/product/poweredge-r640/manuals.

Install a PowerEdge R640 server in a cabinet

Install a PowerEdge R640 server in a cabinet.

This section describes the steps required to install the server in a cabinet.

Cabinet rail kits

There are two types of rail kits used to install the PowerEdge server in a cabinet:

- **Static Rail Kit** -- The static rail kit allows the server to be extracted a short distance before the rail locks. The travel stop position is intended for repositioning your grip for server removal. It is not intended for internal service of the server. The server must be removed from the cabinet so that the system cover can be removed to access the internal FRUs. The static rail kit consists of the following items:

  Two Dell™ ReadyRails™ static rails (1)
  Two chassis rails (2)
  Two hook and loop straps (3)

Figure 1 Static rail kit

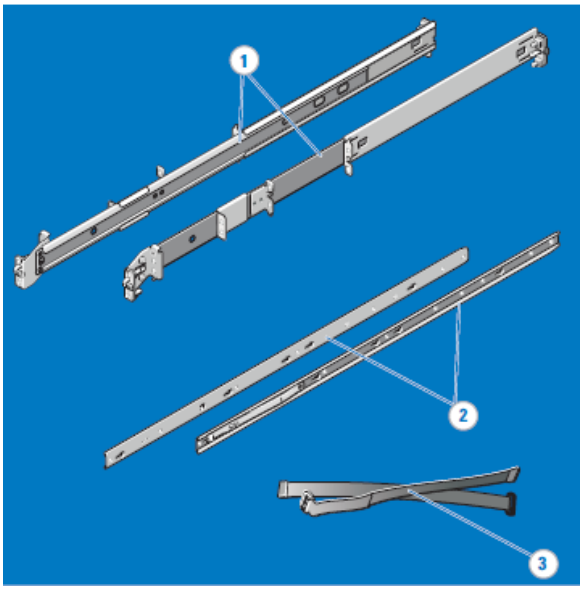

- **Sliding Rail Kit** -- The sliding rail kit allows the server to be extended from the cabinet so that the system cover can be removed to access the internal FRUs. The sliding rail kit consists of the following items:

  Two B6 Dell™ ReadyRails™ II sliding rail assemblies (1)

  Two hook and loop straps (2)

Figure 2 Sliding rail kit

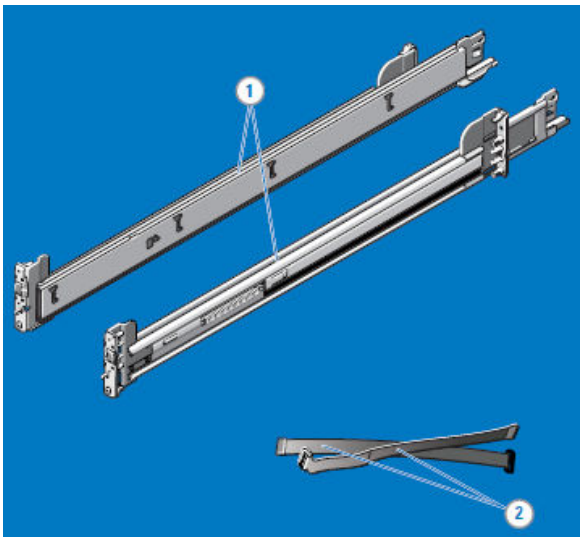

Determine rail kit that will be used

Unpack the rail kit and determine which type of rail kit is being used for installation of the server.

- **Static Rail Kit** -- Install the static rail kit and server as described in Static rail kit and server installation on page 7

- **Sliding Rail Kit** -- Install the sliding rail kit and server as described in Sliding rail kit and server installation on page 9

Static rail kit and server installation

This section describes how to install the static rail kit and the server in the cabinet.

Topics include:

- Install the static rail assemblies in the cabinet on page 7
- Install the static chassis rails on the server on page 8
- Install the server in the cabinet on page 8

Install the static rail assemblies in the cabinet

This procedures describes installation of the static rail assemblies in the cabinet.

Procedure

1. Position the left and right rail end pieces labeled FRONT facing inward and orient each end piece to seat in the holes on the front side of the vertical cabinet flanges (1).
2. Align each end piece in the bottom and top holes of the desired U spaces (2).
3. Engage the back end of the rail until it fully seats on the vertical cabinet flange and the latch clicks into place. Repeat these steps to position and seat the front end piece on the vertical cabinet flange (3).

Note

To remove the rails, pull on the latch release button on the end piece midpoint and unseat each rail (4).

Figure 3 Installing static rail assemblies

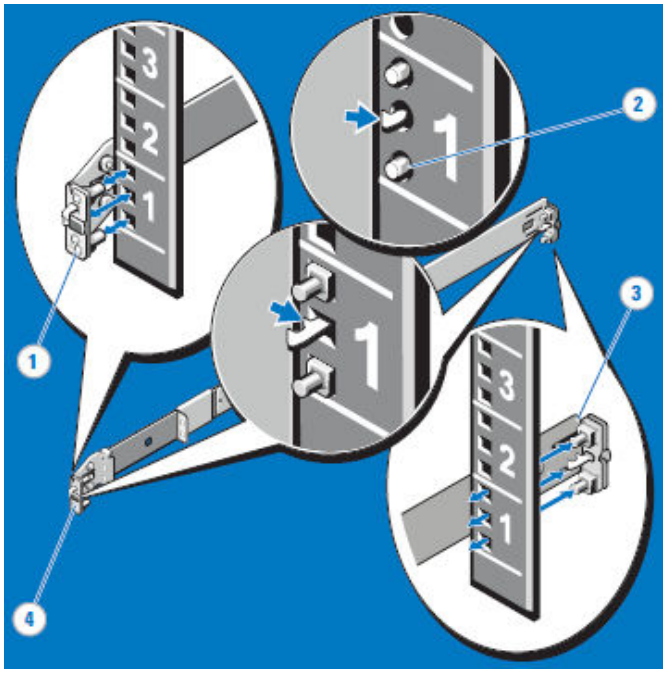

Install the static chassis rails on the server

This procedures describes installation of the static chassis rails on the server.

CAUTION

The server is heavy and should be handled by two people. To avoid personal injury and/or damage to the equipment, do not attempt to handle the server without a mechanical lift and/or help from another person.

Procedure

1. Place the server on a level surface and align the keyhole slots on the chassis rails with the pins on the server (1).
2. Slide the chassis rails toward the back of the server until they lock into place. To remove the chassis rails, lift the lock spring until it clears the pin and slide the chassis rails toward the front of the server until the pins can slip through the keyhole slots (2).

Figure 4 Installing static chassis rails

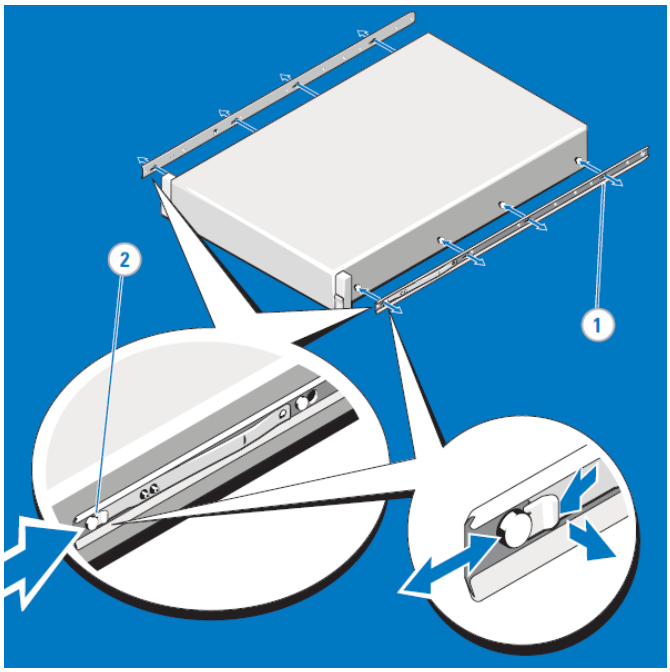

Install the server in the cabinet

This procedures describes installation of the server in the cabinet.

CAUTION

The server is heavy and should be installed in a cabinet by two people. To avoid personal injury and/or damage to the equipment, do not attempt to install the server in a cabinet without a mechanical lift and/or help from another person.

Procedure

1. Insert the ends of the chassis rails into the front of the rails mounted to the cabinet and push the server into the cabinet (1).

2. Depending on the rail configuration, the latches on the front of the server engage automatically (2). If the latches do not engage automatically, secure the server in the cabinet by tightening the screw under each latch (3).
3. Connect I/O and power cables as described in your system cabling documentation.
4. Thread the Velcro straps through the slots on the back brackets and use the Velcro straps to secure the cables (4).

Note

To remove the server from the cabinet, pull the server out until the travel stops are reached. Locate the blue tabs on the sides of the chassis rails. Push the tabs inward and pull the server until the chassis rails are clear of the rails (5). The travel stop position is intended for repositioning your grip for server removal. It is not intended for service. You must be prepared to support the weight of the server when the chassis rails separate from the cabinet rails

Figure 5 Install the server

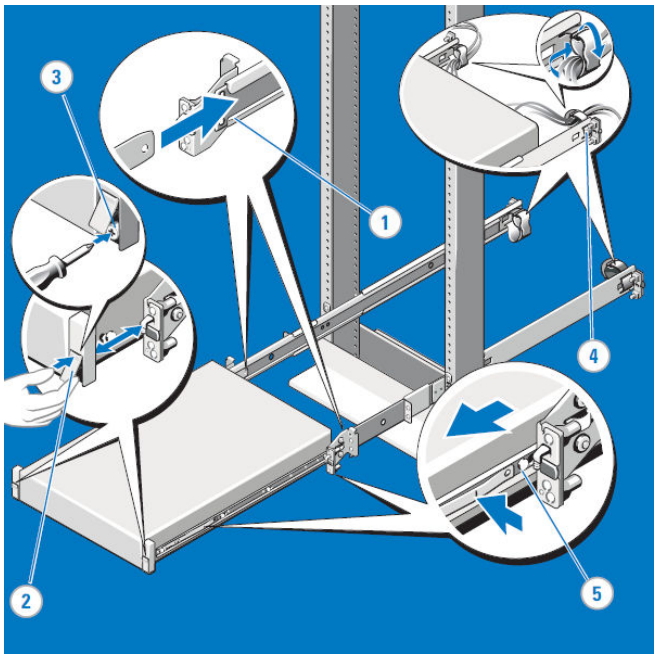

Sliding rail kit and server installation

This section describes how to install the sliding rail kit and the server in the cabinet.

Topics include:

- Install the sliding rail assemblies on page 10
- Install the server on sliding rail assemblies on page 10
- Routing the cables on page 12

Install the sliding rail assemblies

This procedures describes installation of the sliding rail assemblies in the cabinet.

Procedure

1. Position the left and right rail end pieces labeled FRONT facing inward and orient each end piece to seat in the holes on the front side of the vertical cabinet flanges (1).
2. Align each end piece in the bottom and top holes of the desired U spaces (2).
3. Engage the back end of the rail until it fully seats on the vertical cabinet flange and the latch clicks into place. Repeat these steps to position and seat the front end piece on the vertical cabinet flange (3).

Note

To remove the rails, pull on the latch release button on the end piece midpoint and unseat each rail (4).

Figure 6 Installing sliding rail assemblies

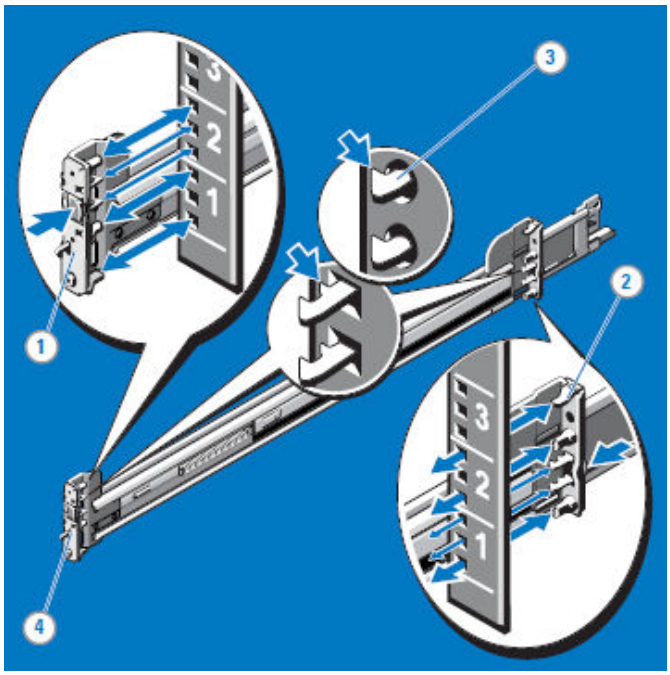

Install the server on sliding rail assemblies

This procedures describes installation of the server on the sliding rail assemblies.

⚠ CAUTION

The server is heavy and should be installed in a cabinet by two people. To avoid personal injury and/or damage to the equipment, do not attempt to install the server in a cabinet without a mechanical lift and/or help from another person.

Procedure

1. Pull the inner slide rails out of the cabinet until they lock into place (1).
2. With the help of second person, position the server above the extended rails.

3. Locate the rear rail standoff on each side of the server and lower them into the rear J-slots on the slide assemblies (2).
4. Rotate the server downward until all the rail standoffs are seated in the J-slots (3).
5. Push the server inward until the lock levers click into place.
6. Press the slide-release lock buttons on both rails (4).

Figure 7 Install server on sliding rail assemblies

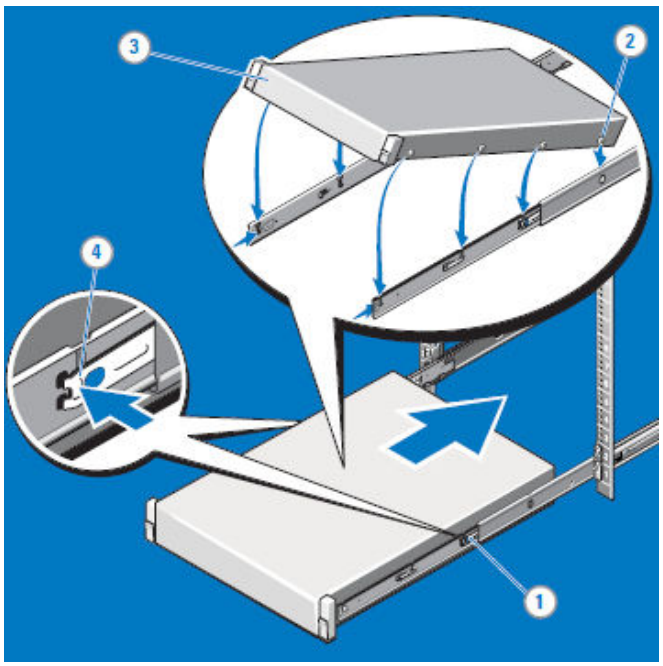

7. Push the server into the cabinet. The slam latch on either side of the server engage automatically as the server is pushed into the cabinet (1).

Note

The slam latches are released by pulling up on the latches (2). To secure the server for unstable environments, locate the hard-mount screw under each latch and tighten each screw with a #2 Phillips screwdriver (3).

Figure 8 Secure the server in the cabinet

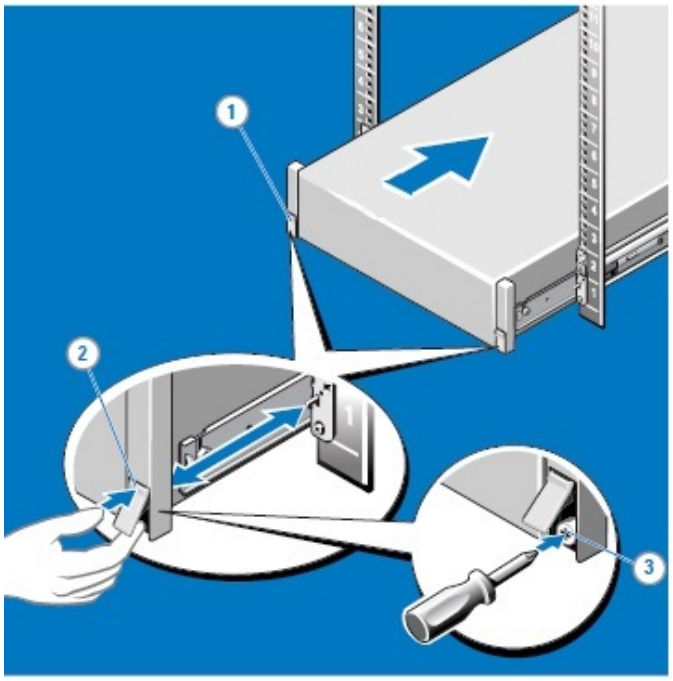

Routing the cables

This procedure describes connecting cables and use of the two hook and loop straps provided in the rail kit to route the cables at the back of the server.

Procedure

1. Connect I/O and power cables as described in your system cabling documentation.
2. Locate the outer CMA brackets on the interior sides of both cabinet flanges (1).
3. Bundle the cables gently, pulling them clear of the server connectors to the left and right sides (2).
4. Thread the hook and loop straps through the tooled slots on the outer CMA brackets on each side of the server to secure the cable bundles (3).

Figure 9 Hook and loop straps

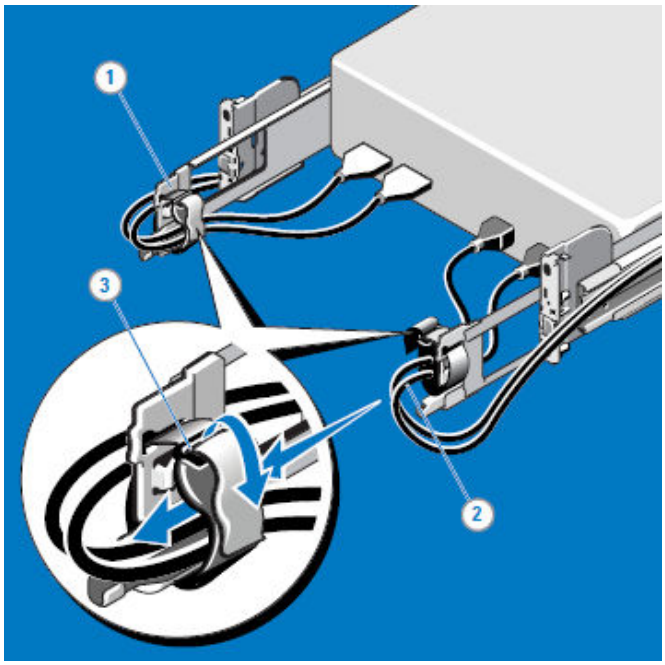

Install bezel

Bezels are application specific, and may not appear as shown. This procedure can be used to install all bezels.

Procedure

1. Position right side of bezel on the server bracket. Rotate the bezel inward so that it is flush with front of the server and the latch on the left side secures the bezel.
2. Lock the bezel with the provided key and store the key in a secure place.

Figure 10 Install the bezel

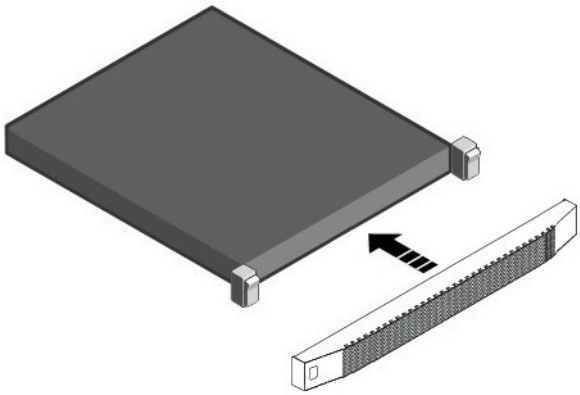

Supported hardware configurations

The 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.

The following points apply, as well:

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

  Each node has four 10/25 GbE ports. Two are used for internal network communications. The other two ports are used for client network traffic.

  The switches must have sufficient 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 GbE and one BMC (iDRAC) port per node.

Network requirements are described later in this document.

Set up the BMC (iDRAC) IP address and BIOS

Set up the BMC (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 BMC (iDRAC) IP address
- The server BMC (iDRAC) subnet mask
- The gateway IP address
- The VLAN ID of the BMC (iDRAC), if the VLAN is used

Note

The BMC (iDRAC) IP address may be an IPv4 or an IPv6 address.

Defaults:

- BMC (iDRAC) username - root

- BMC (iDRAC) password - The default password is calvin for 13G nodes and Scaleio123 for 14G nodes.
- BIOS password - emcbios

During the BMC (iDRAC) IP and BIOS setup, use the following operations:

- Use the arrow keys to navigate in the BIOS screens.
- Use the + and - keys on the keyboard to change the option selection in the BIOS screens.
- Use SPACE or ENTER keys to change the settings.
- The ENTER key opens a list to the desired values.

Procedure

1. Open 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 screen and keyboard connection to the rack.

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

   a. In the **iDRAC Settings** screen, 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 BMC (iDRAC) port:

      - **Enable IP IPv4** = **Enabled**
      - **Enable DHCP** = **Disabled**
      - **Static IP Address** = Static IP address
      - **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 BMC (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** screen.

5. From the **iDRAC Settings** screen, click **Finish**, **Yes**, and then **OK** to return to the **System Setup Main Menu** screen.

6. In the **System Setup Main Menu** screen, select the **System BIOS** menu.
7. In the **System BIOS Settings** screen, 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** screen.
8. Configure boot settings,:

   a. Select **Boot Settings**.

   b. In the **System Boot Settings--Boot Settings** pane, verify that **Boot Mode** is set to **BIOS (13G) or UEFI (14G)**.

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

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

      - In the **Boot Sequence** list, **Hard drive C:** appears as the first item.
      - In the **Hard-Disk Drive Sequence** list, the **SATADOM-ML 3SE (13G) or BOSS (14G)** device appears as the first item.

   e. Click **Back** twice to return to the **System BIOS Settings** screen.
9. Configure integrated devices:

   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** screen, click **Finish**, **Yes**, and then **OK** to return to the **System Setup Main Menu** screen.
11. Select **Finish** to exit the BIOS and apply all settings post boot.

Results

The BMC (iDRAC) IP and server BIOS address configuration is complete.

Verify the status of the system hardware, storage controller, and disks - 14G servers

Verify the status of the system hardware, storage controller, and disks in a VxFlex Ready Node 14G server.

Before you begin

Ensure that you know:

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

Procedure

1. From a browser, go to http://*<BMC/iDRAC_IP_address>*.
2. In the **DELL Console Login** window, type 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. 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 disks, and an option to **Blink/Unblink** the selected disk.
7. Verify that no disk is in failed state.

   If any of the disks is failed, refer to the relevant Disk Replace FRU.
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.

System requirements

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

Note

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

- Configure the M.2 device on the BOSS as a RAID 1 device.
- During OS installation, use Scaleio123 as the password for the user name root or administrator. (Alternately, provide a password that meets the local security criteria.)
- In a 2-Layer installation, 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 OS on the BOSS device.

Open the KVM console - 14G servers

Open the KVM console on a VxFlex Ready Node 14G server.

Before you begin

Ensure that:

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

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.

   A popup security warning screen is displayed.
4. Select **Accept** and click **Run**.

   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, then click **Run** to launch the console.
Google Chrome	a. Click **Show all downloads**.

Browser	Steps
	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.

Map the ISO file on a VxFlex Ready Node 14G server

Map the ISO file on a VxFlex Ready Node 14G server.

Before you begin

Ensure that you can access the console.

Procedure

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

   The **Initializing connection** screen is displayed with a connecting message. You may have to wait for some time for the connecting message to finish.
3. In the **DELL System Setup** screen, select **Virtual Media** > **Map CD/DVD**.
4. From the **Virtual Media - Map CD/DVD** screen, browse and set up the `ISO` file and then click **Map Device**.
5. To verify the `ISO` file selection, click **Virtual Media** and view the ISO selection in the drop down list.
6. Reset the server:

   a. In the main **DELL System Setup** console window, select **Power** > **Rest System (warm boot)** to display the list of keyboard keys with related functionality.

   b. Refer to the best practices of DELL OS installation: How to Install an Operating System on Dell PowerEdge Servers.

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

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

After you finish

Note

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

Accept the Kylin end-user license agreement (EULA)

Complete the installation by accepting the EULA.

Once the server reboots using the OS ISO file, the end-user license agreement is displayed.

Procedure

1. Scroll down to the end of the EULA and enter "**q**".
2. When prompted to accept the EULA, enter "**y**" to accept the agreement.
3. Enter **Ctrl-F1** to get a login prompt.

   This step is only necessary the first time after installing the ISO.
4. At the login prompt, log in using **root** as the username and **Scaleio123** as the password.

Results

Installation of the OS ISO is complete. Proceed with the server port designations.

Server node port designations

For single-node VxFlex Ready Node servers running Linux, connect the cables, and configure the ports as shown in the following configuration table.

Note

For the management server, the data and client network connectivity is optional for troubleshooting purposes only. The management network is mandatory.

The following additional information should be taken into account:

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

Table 2 R640 single-node server configuration

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

Note

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

Note

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

Correlate the PCI slot locations and interface names

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

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

Procedure

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

   Note

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

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

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

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

In this case, the 10GB PCI NIC is presented on slot 6, and the logical bus address of its ports is 5:0 and 5:1. (See Figure 11 on page 23, in which the correlation between PCIe slot and the logical bus address is highlighted in yellow.)

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

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

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

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

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

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

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

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

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

Figure 11 Correlating physical slot and Ethernet network names

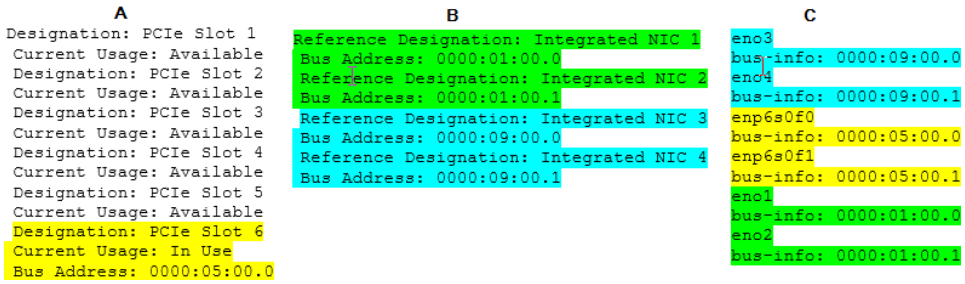

2. Create a table for your own use, similar to the example table shown below:
   - Base the table on the SLES port definitions for the relevant server type described in Server node port designations on page 20.
   - In the **Port** row, record the enXXX names that correspond with the PCI slots.
   - Map each line in turn from Step 1.a on page 22 output to the corresponding line in Step 1.b on page 22 output.

Note

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

For example, the last column in the following table was created using the information from the 4th line in step 1.a on page 22.:

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

Configure the ports (Linux-SLES)

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

Procedure

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

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

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

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

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

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

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

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

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

   Note

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

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

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

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

7. Configure the default gateway:

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

For example:

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

Note

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

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

---

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

Results

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

Validate the AMS

Validate that the AMS service is up and running.

Procedure

1. Check that the AMS service is up and running:

```
service scaleio-ams status
```

Output similar to the following is displayed:

```
* scaleio-ams.service - EMC ScaleIO AMS service
   Loaded: loaded (/etc/systemd/system/scaleio-ams.service;
enabled; vendor preset: disabled)
   Active: active (running) since Wed 2018-01-31 04:50:59 UTC;
1 weeks 1 days ago
  Process: 115380 ExecStart=/opt/emc/scaleio/ams/bin/startup.sh
> /dev/null 2> /dev/null (code=exited, status=0/SUCCESS)
 Main PID: 115395 (java)
    Tasks: 209 (limit: 512)
   CGroup: /system.slice/scaleio-ams.service
           `-115395 /usr/bin/java -
Djava.util.logging.config.file=/opt/emc/scaleio/ams/conf/
logging.properties -
Djava.util.logging.manager=org.apache.juli.ClassLoaderLogManager
 -Xms128m -Xmx1024m -Djava.endorsed.dirs=/opt/emc/scaleio/ams/
endorsed -classpath /opt/emc/scaleio/ams/bin/jul-to-
slf4j-1.7.5.jar:/opt/emc/scaleio/ams/bin/slf4j-
api-1.7.5.jar:/opt/emc/scaleio/ams/bin/logback-config:/opt/emc/
scaleio/ams/webapps/ROOT/WEB-INF/lib/logback-
classic-1.0.13.jar:/opt/emc/scaleio/ams/lib/*:/opt/emc/
scaleio/ams/webapps/ROOT/WEB-INF/lib/vasa-truststore-
manager-2.0.0-SNAPSHOT.jar:/opt/emc/scaleio/ams/bin/
bootstrap.jar:/opt/emc/scaleio/ams/bin/tomcat-juli.jar -
Dcatalina.base=/opt/emc/scaleio/ams -Dcatalina.home=/opt/emc/
scaleio/ams -Djava.io.tmpdir=/opt/emc/scaleio/ams/temp
org.apache.catalina.startup.Bootstrap start

Jan 31 04:50:59 sio-dcoe-122CJ-20 systemd[1]: Starting EMC
ScaleIO AMS service...
Jan 31 04:50:59 sio-dcoe-122CJ-20 startup.sh[115380]: Tomcat
started.
Jan 31 04:50:59 sio-dcoe-122CJ-20 systemd[1]: Started EMC
ScaleIO AMS service.
```

2. Once you have validated that the AMS service is running, proceed to the next relevant document for further instructions:
   - For an initial deployment, refer to "Deploying VxFlex Ready Node" in the *VxFlex Ready Node Deployment Guide*.
   - For an AMS recovery or Take Ownership flow, refer to the relevant troubleshooting topic in the *VxFlex Ready Node Troubleshooting Guide*.

Upgrade the AMS version

Upgrade the AMS from a previous version.

When upgrading, you should upgrade the GUI, the AMS, and the components. This procedure is described fully in the *VxFlex Ready Node Upgrade Guide*.

Procedure

1. Copy the new version of the AMS Linux package to the server.
2. Upgrade the AMS by running this command:

```
rpm -U EMC-ScaleIO-ams-<version>-<build>.x86_64.rpm
```

Results

The AMS is upgraded.

After you finish

Upgrade the other components, as described in the *VxFlex Ready Node Upgrade Guide*.

Copyright © 2018 Dell Inc. or its subsidiaries. All rights reserved.

Published June 2018

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, 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.