

Dell EMC PowerVault ME4 Series Storage System

Deployment Guide

Notes, cautions, and warnings

 **NOTE:** A NOTE indicates important information that helps you make better use of your product.

 **CAUTION:** A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

 **WARNING:** A WARNING indicates a potential for property damage, personal injury, or death.

Chapter 1: Before you begin.....	6
Unpack the enclosure.....	6
Safety guidelines.....	7
Safe handling.....	7
Safe operation.....	8
Electrical safety.....	8
Rack system safety precautions.....	9
Installation checklist.....	9
Planning for installation.....	10
Preparing for installation.....	10
Preparing the site and host server.....	11
Required tools.....	11
Requirements for rackmount installation.....	11
Disk drive module.....	11
Drive carrier module in 2U chassis.....	11
Drive status indicators.....	12
Blank drive carrier modules.....	13
DDIC in a 5U enclosure.....	13
Populating drawers with DDICs.....	14
Chapter 2: Mount the enclosures in the rack.....	15
Rackmount rail kit.....	15
Install the 2U enclosure.....	15
Install the 2U enclosure front bezel.....	16
Install the 5U84 enclosure.....	16
Connect optional expansion enclosures.....	18
Cable requirements for expansion enclosures.....	18
Chapter 3: Connect to the management network.....	21
Chapter 4: Cable host servers to the storage system.....	22
Cabling considerations.....	22
Connecting the enclosure to hosts.....	22
CNC technology.....	22
Fibre Channel protocol.....	23
iSCSI protocol.....	23
SAS protocol.....	25
Host connection.....	25
16 Gb Fibre Channel host connection.....	25
10 GbE iSCSI host connection.....	25
10Gbase-T host connection.....	25
12 Gb HD mini-SAS host connection.....	26
Connecting direct attach configurations.....	26
Single-controller module configurations.....	26

Dual-controller module configurations.....	26
Chapter 5: Connect power cables and power on the storage system.....	30
Power cable connection.....	30
Chapter 6: Perform system and storage setup.....	33
Record storage system information.....	33
Using guided setup.....	33
Web browser requirements and setup.....	33
Access the PowerVault Manager.....	33
Update firmware.....	34
Use guided setup in the PowerVault Manager Welcome panel.....	34
Chapter 7: Perform host setup.....	43
Host system requirements.....	43
About multipath configuration.....	43
Windows hosts.....	43
Configuring a Windows host with FC HBAs.....	43
Configuring a Windows host with iSCSI network adapters.....	45
Configuring a Windows host with SAS HBAs.....	48
Linux hosts.....	50
Configuring a Linux host with FC HBAs	50
Configure a Linux host with iSCSI network adapters.....	52
SAS host server configuration for Linux.....	55
VMware ESXi hosts.....	57
Fibre Channel host server configuration for VMware ESXi.....	57
iSCSI host server configuration for VMware ESXi.....	58
SAS host server configuration for VMware ESXi.....	62
Citrix XenServer hosts.....	64
Fibre Channel host server configuration for Citrix XenServer.....	64
iSCSI host server configuration for Citrix XenServer.....	66
SAS host server configuration for Citrix XenServer.....	69
Chapter 8: Troubleshooting and problem solving.....	71
Locate the service tag.....	71
Operators (Ops) panel LEDs.....	71
2U enclosure Ops panel.....	71
5U enclosure Ops panel.....	72
Initial start-up problems.....	74
2U enclosure LEDs.....	75
5U enclosure LEDs.....	77
Module LEDs.....	80
Troubleshooting 2U enclosures.....	81
Troubleshooting 5U enclosures.....	82
Fault isolation methodology.....	83
Options available for performing basic steps.....	83
Performing basic steps.....	84
If the enclosure does not initialize.....	85
Correcting enclosure IDs.....	85

Host I/O.....	86
Dealing with hardware faults.....	86
Appendix A: Cabling for replication.....	90
Connecting two storage systems to replicate volumes.....	90
Host ports and replication.....	91
Example cabling for replication.....	91
Single-controller module configuration for replication.....	91
Dual-controller module configuration for replication.....	92
Isolating replication faults.....	94
Diagnostic steps for replication setup.....	95
Appendix B: SFP+ transceiver for FC/iSCSI ports.....	98
Appendix C: System Information Worksheet.....	100
Appendix D: Setting network port IP addresses using the CLI port and serial cable.....	103
Mini-USB Device Connection.....	106
Microsoft Windows drivers.....	106
Linux drivers.....	107

Before you begin

This document describes the initial hardware setup for Dell EMC PowerVault ME4 Series storage systems.

This document might contain third-party content that is not under the control of Dell EMC. The language in the third-party content might be inconsistent with the current guidelines for Dell EMC content. Dell EMC reserves the right to update this document after the content is updated by the relevant third parties.

Topics:

- [Unpack the enclosure](#)
- [Safety guidelines](#)
- [Installation checklist](#)
- [Planning for installation](#)
- [Preparing for installation](#)
- [Disk drive module](#)
- [Populating drawers with DDICs](#)

Unpack the enclosure

Examine the packaging for crushes, cuts, water damage, or any other evidence of mishandling during transit. If you suspect that damage has happened, photograph the package before opening, for possible future reference. Retain the original packaging materials for use with returns.

- Unpack the 2U storage system and identify the items in your shipment.

NOTE: The cables that are used with the enclosure are not shown in [Unpacking the 2U12 and 2U24 enclosures](#) on page 6. The rail kit and accessories box is located below the 2U enclosure shipping box lid.




Figure 1. Unpacking the 2U12 and 2U24 enclosures

- | | |
|---------------------------------------|-----------------------------|
| 1. Storage system enclosure | 2. Rackmount left rail (2U) |
| 3. Rackmount right rail (2U) | 4. Documentation |
| 5. Enclosure front-panel bezel option | 6. Rack mount ears |
- 2U enclosures are shipped with the controller modules or input/output modules (IOMs) installed. Blank drive carrier modules must be installed in the unused drive slots.
 - For enclosures configured with CNC controller modules, locate the SFP+ transceivers included with the shipment. See [SFP+ transceiver for FC/iSCSI ports](#) on page 98.
 - Unpack the 5U84 storage system and identify the items in your shipment.

NOTE: The cables that are used with the enclosure are not shown in [Unpacking the 5U84 enclosure](#) on page 7. The rail kit and accessories box is located below the 5U84 enclosure shipping box lid.




Figure 2. Unpacking the 5U84 enclosure

- | | |
|--------------------------------|-----------------------------------|
| 1. Storage system enclosure | 2. DDICs (Disk Drive in Carriers) |
| 3. Documentation | 4. Rackmount left rail (5U84) |
| 5. Rackmount right rail (5U84) | 6. Drawers |
- DDICs ship in a separate container and must be installed into the enclosure drawers during product installation. For rackmount installations, DDICs are installed after the enclosure is mounted in the rack. See [Populating drawers with DDICs](#) on page 14.
 - For enclosures configured with CNC controller modules, locate the SFP+ transceivers included with the shipment. See [SFP+ transceiver for FC/iSCSI ports](#) on page 98.
- CAUTION:** A 5U enclosure does not ship with DDICs installed, but the rear panel controller modules or IOMs are installed. This partially populated enclosure weighs approximately 64 kg (142 lb). You need a minimum of two people to remove the enclosure from the box.

Safety guidelines

Always follow these safety guidelines to avoid injury and damage to ME4 Series components.

If you use this equipment in a manner that is not specified by Dell EMC, the protection that is provided by the equipment could be impaired. For your safety and precaution, observe the rules that are described in the following sections:

NOTE: See the *Dell EMC PowerVault ME4 Series Storage System Getting Started Guide* for product safety and regulatory information. Warranty information is included as a separate document.

Safe handling

Dell EMC recommends that only individuals with rack-mounting experience install an enclosure into a rack.

CAUTION: Use this equipment in a manner specified by Dell EMC. Failure to do so may cancel the protection that is provided by the equipment.

- Unplug the enclosure before you move it or if you think that it has become damaged in any way.
- A safe lifting height is 20U.
- Always remove the power cooling modules (PCMs) to minimize weight before you move the enclosure.
- Do not lift the enclosures by the handles on the PCMs—they are not designed to take the weight.

CAUTION: Do not try to lift the enclosure by yourself:

- Fully configured 2U12 enclosures can weigh up to 32 kg (71 lb)
- Fully configured 2U24 enclosures can weigh up to 30 kg (66 lb)

- Fully configured 5U84 enclosures can weigh up to 135 kg (298 lb). An unpopulated enclosure weighs 46 kg (101 lb).
- Use a minimum of two people to lift the 5U84 enclosure from the shipping box and install it in the rack.

Before lifting the enclosure:

- Avoid lifting the enclosure using the handles on any of the CRUs because they are not designed to take the weight.
- Do not lift the enclosure higher than 20U. Use mechanical assistance to lift above this height.
- Observe the lifting hazard label affixed to the storage enclosure.

Safe operation

Operation of the enclosure with modules missing disrupts the airflow and prevents the enclosure from receiving sufficient cooling.

i NOTE: For a 2U enclosure, all IOM and PCM slots must be populated. In addition, empty drive slots (bays) in 2U enclosures must hold blank drive carrier modules. For a 5U enclosure, all controller module, IOM, FCM, and PSU slots must be populated.

- Follow the instructions in the module bay caution label affixed to the module being replaced.
- Replace a defective PCM with a fully operational PCM within 24 hours. Do not remove a defective PCM unless you have a replacement model of the correct type ready for insertion.
- Before removal/replacement of a PCM or PSU, disconnect supply power from the module to be replaced. See the *Dell EMC PowerVault ME4 Series Storage System Owner's Manual*.
- Follow the instructions in the hazardous voltage warning label affixed to power cooling modules.

△ CAUTION: 5U84 enclosures only

- To prevent a rack from tipping over, drawer interlocks stop users from opening both drawers simultaneously. Do not attempt to force open a drawer when the other drawer in the enclosure is already open. In a rack containing more than one 5U84 enclosure, do not open more than one drawer per rack at a time.
- Observe the hot surface label that is affixed to the drawer. Operating temperatures inside enclosure drawers can reach 60°C (140°F). Take care when opening drawers and removing DDICs.
- Due to product acoustics, ear protection should be worn during prolonged exposure to the product in operation.
- Observe the drawer caution label. Do not use open drawers to support any other objects or equipment.

Electrical safety

- The 2U enclosure must be operated from a power supply input voltage range of 100–240 VAC, 50/60Hz.
- The 5U enclosure must be operated from a power supply input voltage range of 200–240 VAC, 50/60Hz.
- Provide a power source with electrical overload protection to meet the requirements in the technical specification.
- The power cord must have a safe electrical grounding connection. Check the grounding connection of the enclosure before you switch on the power supply.

i NOTE: The enclosure must be grounded before applying power.

- The plug on the power supply cord is used as the main disconnect device. Ensure that the socket outlets are located near the equipment and are accessible.
- 2U enclosures are intended to operate with two PCMs.
- 5U84 enclosures are intended to operate with two PSUs.
- Follow the instructions that are shown on the power-supply disconnection caution label that is affixed to power cooling modules.

△ CAUTION: Do not remove the covers from the enclosure or any of the modules as there is a danger of electric shock inside.

Rack system safety precautions

The following safety requirements must be considered when the enclosure is mounted in a rack:

- The rack construction must support the total weight of the installed enclosures. The design should incorporate stabilizing features to prevent the rack from tipping or being pushed over during installation or in normal use.
- When loading a rack with enclosures, fill the rack from the bottom up; and empty the rack from the top down.
- Always remove all power supply modules to minimize weight, before loading the enclosure into the rack.
- Do not try to lift the enclosure by yourself.

CAUTION: To prevent of the rack falling over, never move more than one enclosure out of the cabinet at any one time.

- The system must be operated with low-pressure rear exhaust installation. The back pressure that is created by rack doors and obstacles must not exceed 5 pascals (0.5 mm water gauge).
- The rack design should take into consideration the maximum operating ambient temperature for the enclosure. The maximum operating temperature is 35°C (95°F) for controllers and 40°C (104°F) for expansion enclosures.
- The rack should have a safe electrical distribution system. It must provide overcurrent protection for the enclosure. Make sure that the rack is not overloaded by the total number of enclosures that are installed in the rack. Consideration should be given to the electrical power consumption rating shown on the nameplate.
- The electrical distribution system must provide a reliable connection for each enclosure in the rack.
- Each PSU or PCM in each enclosure has a grounding leakage current of 1.0 mA. The design of the electrical distribution system must take into consideration the total grounding leakage current from all the PSUs/PCMs in all the enclosures. The rack requires labeling with “High Leakage Current. Grounding connection essential before connecting supply.”

Installation checklist

This section shows how to plan for and successfully install your enclosure system into an industry standard 19-inch rack cabinet.

CAUTION: Use only the power cables supplied when installing the storage system.

The following table outlines the steps that are required to install the enclosures, and initially configure and provision the storage system:

NOTE: To ensure successful installation, perform the tasks in the order presented.

Table 1. Installation checklist

Step	Task	Where to find procedure
1	Unpack the enclosure.	See Unpack the enclosure on page 6.
2	Install the controller enclosure and optional expansion enclosures in the rack. ¹	See Required tools on page 11. See Requirements for rackmount installation on page 11. See Install the 2U enclosure on page 15. See Install the 5U84 enclosure on page 16.
3	Populate drawers with disks (DDICs) in 5U84 enclosure; 2U enclosures ship with disks installed.	See Populating drawers with DDICs on page 14.
4	Cable the optional expansion enclosures.	See Connect optional expansion enclosures on page 18.
5	Connect the management ports.	See Connect to the management network on page 21.
6	Cable the controller host ports. ²	See Connecting the enclosure to hosts on page 22.
7	Connect the power cords and power on the system.	See Power cable connection on page 30.
8	Perform system and storage setup.	See Using guided setup on page 33.

Table 1. Installation checklist (continued)

Step	Task	Where to find procedure
9	Perform host setup: <ul style="list-style-type: none"> • Attach the host servers. • Install the required host software. 	See Host system requirements on page 43. See Windows hosts on page 43. See Linux hosts on page 50. See VMware ESXi hosts on page 57. See Citrix XenServer hosts on page 64.
10	Perform the initial configuration tasks. ³	See Using guided setup on page 33.

¹ The environment in which the enclosure operates must be dust-free to ensure adequate airflow.

² For more information about hosts, see the *About hosts* topic in the *Dell EMC PowerVault ME4 Series Storage System Administrator's Guide*.

³ The PowerVault Manager is introduced in [Using guided setup](#) on page 33. See the *Dell EMC PowerVault ME4 Series Storage System Administrator's Guide* or online help for additional information.

Planning for installation

Before beginning the enclosure installation, familiarize yourself with the system configuration requirements.

Table 2. System configuration

Module type	Location	Description
Drive carrier modules	2U front panel	All drive slots must hold either a drive carrier or blank drive carrier module. Empty slots are not allowed. At least one disk must be installed.
DDIC	5U front panel drawers	Maximum 84 disks are installed (42 disks per drawer). Minimum 28 disks are required. Follow the drawer population rules in Populating drawers with DDICs on page 14.
Power cooling modules	2U rear panel	Two PCMs provide full power redundancy, allowing the system to continue to operate while a faulty PCM is replaced.
Power supply unit modules	5U rear panel	Two PSUs provide full power redundancy, allowing the system to continue to operate while a faulty PSU is replaced.
Fan cooling modules	5U rear panel	Five FCMs provide airflow circulation, maintaining all system components below the maximum temperature allowed.
Controller modules and IOMs	Rear panel	<ul style="list-style-type: none"> • One or two controller modules may be installed in 2U12 and 2U24 enclosures. • Two controller modules must be installed in 5U84 enclosures. • Two IOMs must be installed in 2U12, 2U24, and 5U84 enclosures.

Preparing for installation

- NOTE:** Enclosure configurations:
- 2U enclosures are delivered with CRUs and all drive carrier modules installed.
 - 5U84 enclosures are delivered with CRUs installed; however, DDICs must be installed during system setup.
 - 5U84 enclosures require 200–240VAC for operation. See the *Environmental requirements* topic in the *Dell EMC PowerVault ME4 Series Storage System Owner's Manual* for detailed information

- CAUTION:** Lifting enclosures:
- **A 2U enclosure, including all its component parts, is too heavy for one person to lift and install into the rack cabinet. Two people are required to safely move a 2U enclosure.**

- **A 5U enclosure, which is delivered without DDICs installed, requires two people to lift it from the box. A mechanical lift is required to hoist the enclosure for positioning in the rack.**

Make sure that you wear an effective antistatic wrist or ankle strap and follow conventional ESD precautions when touching modules and components. Do not touch the midplane, motherboard, or module connectors. See [Safety guidelines](#) on page 7 for important preparation requirements and handling procedures to use during product installation.

Preparing the site and host server

Before beginning the enclosure installation, verify that the site where you plan to install your storage system has the following:

- Each redundant power supply module requires power from an independent source or a rack power distribution unit with Uninterruptible Power Supply (UPS). 2U enclosures use standard AC power and the 5U84 enclosure requires high-line (high-voltage) AC power.
- A host computer configured with the appropriate software, BIOS, and drives. Contact your supplier for the correct software configurations.

Before installing the enclosure, verify the existence of the following:

- Depending upon the controller module: SAS, Fibre Channel (FC), or iSCSI HBA and appropriate switches (if used)
- Qualified cable options for host connection
- One power cord per PCM or PSU
- Rail kit (for rack installation)

Contact your supplier for a list of qualified accessories for use with the enclosure. The accessories box contains the power cords and other accessories.

Required tools

The following tools are required to install an ME4 Series enclosure:

- Phillips screwdriver
- Torx T20 bit for locks and select CRU replacement

Requirements for rackmount installation

You can install the enclosure in an industry standard 19-inch cabinet capable of holding 2U form factors.

NOTE: See the *Dell EMC PowerVault ME4 Series Owner's Manual* for front and rear panel product views.

- Minimum depth: 707 mm (27.83") from rack posts to maximum extremity of enclosure (includes rear panel cabling and cable bend radii).
- Weight:
 - Up to 32 kg (71 lb), dependent upon configuration, per 2U enclosure.
 - Up to 128 kg (282 lb), dependent upon configuration, per 5U enclosure.
- The rack should cause a maximum back pressure of 5 pascals (0.5 mm water gauge).
- Before you begin, ensure that you have adequate clearance in front of the rack for installing the rails.

Disk drive module

The ME4 Series Storage System supports different disk drive modules for use in 2U and 5U84 enclosures.

- The disk drive modules that are used in 2U enclosures are referred to as drive carrier modules.
- The disk drive modules that are used in 5U84 enclosures are referred to as Disk Drive in Carrier (DDIC) modules.

Drive carrier module in 2U chassis

The drive carrier module consists of a disk drive that is installed in a carrier module.

- Each 2U12 drive slot holds a single low profile 1.0 in. high, 3.5 in. form factor disk drive in its carrier. The disk drives are horizontal. A 2.5" to 3.5" carrier adapter is available to accommodate 2.5" disk drives.

- Each 2U24 drive slot holds a single low profile 5/8 inch high, 2.5 in. form factor disk drive in its carrier. The disk drives are vertical.

The carriers have mounting locations for:

- Direct dock SAS drives.

A sheet steel carrier holds each drive, which provides thermal conduction, radio frequency, and electro-magnetic induction protection, and physically protects the drive.

The front cap also has an ergonomic handle which gives the following functions:

- Secure location of the carrier into and out of drive slots.
- Positive spring-loading of the drive/midplane connector.

The carrier can use this interface:

- Dual path direct dock Serial Attached SCSI.

The following figures display the supported drive carrier modules:




Figure 3. Dual path LFF 3.5" drive carrier module




Figure 4. Dual path SFF 2.5" drive carrier module




Figure 5. 2.5" to 3.5" hybrid drive carrier adapter

Drive status indicators

Green and amber LEDs on the front of each drive carrier module indicate disk drive status.

Blank drive carrier modules

Blank drive carrier modules, also known as drive blanks, are provided in 3.5" (2U12) and 2.5" (2U24) form factors. They must be installed in empty disk slots to create a balanced air flow.




Figure 6. Blank drive carrier modules: 3.5" drive slot (left); 2.5" drive slot (right)

DDIC in a 5U enclosure

Each disk drive is installed in a DDIC that enables secure insertion of the disk drive into the drawer with the appropriate SAS carrier transition card.

The DDIC features a slide latch button with directional arrow. The slide latch enables you to install and secure the DDIC into the disk slot within the drawer. The slide latch also enables you to disengage the DDIC from its slot, and remove it from the drawer. The DDIC has a single Drive Fault LED, which illuminates amber when the disk drive has a fault.

The following figure shows a DDIC with a 3.5" disk drive:




Figure 7. 3.5" disk drive in a DDIC

The following figure shows a DDIC with a hybrid drive carrier adapter and a 2.5" disk drive:




Figure 8. 2.5" drive in a 3.5" DDIC with a hybrid drive carrier adapter

Populating drawers with DDICs

The 5U84 enclosure does not ship with DDICs installed. Before populating drawers with DDICs, ensure that you adhere to the following guidelines:

- The minimum number of disks that are supported by the enclosure is 28, 14 in each drawer.
- DDICs must be added to disk slots in complete rows (14 disks at a time).
- Beginning at the front of each drawer, install DDICs consecutively by number, and alternately between the top drawer and the bottom drawer. For example, install first at slots 0–13 in the top drawer, and then 42–55 in the bottom drawer. After that, install slots 14–27, and so on.
- The number of populated rows must not differ by more than one row between the top and bottom drawers.
- Hard disk drives (HDD) and solid-state drives (SDD) can be mixed in the same drawer.
- HDDs installed in the same row should have the same rotational speed.
- DDICs holding 3.5" disks can be intermixed with DDICs holding 2.5" disks in the enclosure. However, each row should be populated with disks of the same form factor (all 3.5" disks or 2.5" disks).

The following figure shows a drawer that is fully populated with DDICs:

- See [3.5" disk drive in a DDIC](#) on page 13 for the DDIC holding the 3.5" disk
- See [2.5" drive in a 3.5" DDIC with a hybrid drive carrier adapter](#) on page 14 for the DDIC holding the 2.5" disk with 3.5" adapter




Figure 9. 5U84 enclosure drawer fully populated with DDICs

Mount the enclosures in the rack

This section describes how to unpack the ME4 Series Storage System equipment, prepare for installation, and safely mount the enclosures into the rack.

Topics:

- [Rackmount rail kit](#)
- [Install the 2U enclosure](#)
- [Install the 5U84 enclosure](#)
- [Connect optional expansion enclosures](#)

Rackmount rail kit

Rack mounting rails are available for use in 19-inch rack cabinets.

The rails have been designed and tested for the maximum enclosure weight. Multiple enclosures may be installed without loss of space in the rack. Use of other mounting hardware may cause some loss of rack space. Contact Dell EMC to ensure that suitable mounting rails are available for the rack you plan to use.

Install the 2U enclosure

The 2U enclosure is delivered with the disks installed.

1. Remove the rail kit from the accessories box, and examine for damage.
2. Ensure that the preassembled rails are the correct length for the rack.
3. Use the following procedure to install the rail in the rack:
 - a. Loosen the position locking screws on the rail.
 - b. Identify the rack holes for installing the rails in the rack and insert the rail pins into the rear rack post.
 - c. Extend the rail to fit between the front and rear rack posts and insert the rail pins into the front rack post.

NOTE: Ensure that the rail pins are fully inserted in the rack holes in the front and rear rack posts.
 - d. Use the clamping screws to secure the rail to the rack posts and tighten the position locking screws on the rail.




Figure 10. Install the rail in the rack (left hand rail shown for 2U enclosure)

Table 3. Install the rail in the rack

Item	Description	Item	Description
1	Front rack post (square hole)	6	Clamping screw
2	Rail pins (two per rail)	7	Enclosure fastening screw
3	Left rail	8	2U Ops panel installation detail (exploded view)
4	Rear rack post (square hole)	9	Position locking screw
5	Clamping screw	10	Enclosure fastening screw

- e. Repeat the previous steps to install the other rail in the rack.
- 4. Install the enclosure into the rack:
 - a. Lift the enclosure and align it with the installed rack rails.
 - NOTE:** Ensure that the enclosure remains level while installing it in the rack.
 - b. Carefully insert the slides on each side of the enclosure into the rack rails.
 - c. Push the enclosure fully into the rack.
 - d. Secure the enclosure to the rack using the enclosure fastening screws.

Install the 2U enclosure front bezel

Install the bezel if it was included with the enclosure.

While holding the bezel in your hands, face the front panel of the 2U12 or 2U24 enclosure.

1. Hook the right end of the bezel onto the right ear cover of the storage system.




Figure 11. Attach the bezel to the front of the 2U enclosure

2. Insert the left end of the bezel into the securing slot until the release latch snaps into place.
3. Secure the bezel with the keylock as shown in [Attach the bezel to the front of the 2U enclosure](#) on page 16.

NOTE: To remove the bezel from the 2U enclosure front panel, reverse the order of the previous steps.

Install the 5U84 enclosure

The 5U84 enclosure is delivered without the disks installed.

NOTE: Due to the weight of the enclosure, install it into the rack without DDICs installed, and remove the rear panel CRUs to decrease the enclosure weight.

The adjustment range of the rail kit from the front post to the rear post is 660 mm–840 mm. This range suits a one-meter deep rack within Rack Specification IEC 60297.

1. Remove the rail kit from the accessories box, and examine for damage.
2. Ensure that the preassembled rails are the correct length for the rack.
3. Use the following procedure to install the rail in the rack:
 - a. Loosen the position locking screws on the rail.
 - b. Identify the rack holes for installing the rails in the rack and insert the rail pins into the rear rack post.

- c. Extend the rail to fit between the front and rear rack posts and insert the rail pins into the front rack post.
 - NOTE:** Ensure that the rail pins are fully inserted in the rack holes in the front and rear rack posts.
- d. Use the clamping screws to secure the rail to the rack posts and tighten the position locking screws on the rail.
- e. Ensure the four rear spacer clips (not shown) are fitted to the edge of the rack post.




Figure 12. Install the rail in the rack (left hand rail shown for 5U enclosure)

Table 4. Install the rail in the rack

Item	Description	Item	Description
1	Enclosure fastening screws (A)	7	5U84 chassis section shown for reference
2	Left rail	8	Front rack post (square hole)
3	Rear rack post (square hole)	9	Position locking screws
4	Clamping screw (B)	10	5U84 chassis section shown for reference
5	Clamping screw (B)	11	Enclosure fastening screw (C)
6	Rail pins (quantity 4 per rail)	12	Rail kit fasteners used in rackmount installation (A= fastening; B= clamping; C= fastening)

- f. Repeat the previous steps to install the other rail in the rack.
4. Install the enclosure into the rack:
 - a. Lift the enclosure and align it with the installed rack rails.
 - CAUTION:** A mechanical lift is required to safely lift the enclosure for positioning in the rack.
 - b. Slide the enclosure onto the rails until it is fully seated.
 - c. Fasten the front of the enclosure to the rack using the fastening screws.
 - d. Fix the rear of the enclosure to the sliding bracket with the rear enclosure fastening screws.

Reinsert the rear panel modules and install the DDICs into the drawers. See the instructions in the *Dell EMC PowerVault ME4 Series Storage System Owner's Manual*.

- Installing a controller module
- Installing an IOM
- Installing a fan cooling module
- Installing a PSU
- Installing a DDIC

Connect optional expansion enclosures

ME4 Series controller enclosures support 2U12, 2U24, and 5U84 expansion enclosures. 2U12 and 2U24 expansion enclosures can be intermixed, however 2U expansion enclosures cannot be intermixed with 5U84 expansion enclosures in the same storage system.

NOTE: To add expansion enclosures to an existing storage system, power down the controller enclosure before connecting the expansion enclosures.

- ME4 Series 2U controller enclosures support up to ten 2U enclosures (including the controller enclosure), or a maximum of 240 disk drives.
- ME4 Series 5U controller enclosures support up to four 5U enclosures (including the controller enclosure), or a maximum of 336 disk drives.
- ME4 Series expansion enclosures are equipped with dual-IOMs. These expansion enclosures cannot be cabled to a controller enclosure equipped with a single IOM.
- The enclosures support reverse SAS cabling for adding expansion enclosures. Reverse cabling enables any drive enclosure to fail—or be removed—while maintaining access to other enclosures. Fault tolerance and performance requirements determine whether to optimize the configuration for high availability or high performance when cabling.

Cable requirements for expansion enclosures

ME4 Series supports 2U12, 2U24, and 5U84 form factors, each of which can be configured as a controller enclosure or an expansion enclosure. Key enclosure characteristics include:

NOTE: To add expansion enclosures to an existing storage system, power down the controller enclosure before connecting the expansion enclosures.

- When connecting SAS cables to IOMs, use only supported HD mini-SAS x4 cables.
- Qualified HD mini-SAS to HD mini-SAS 0.5 m (1.64 ft.) cables are used to connect cascaded enclosures in the rack.
- The maximum enclosure cable length that is allowed in any configuration is 2 m (6.56 ft.).
- When adding more than two expansion enclosures, you may need to purchase additional cables, depending upon the number of enclosures and cabling method used.
- You may need to order additional or longer cables when reverse-cabling a fault-tolerant configuration.

Per common convention in cabling diagrams, the controller enclosure is shown atop the stack of connected expansion enclosures. In reality, you can invert the order of the stack for optimal weight and placement stability in the rack. The schematic representation of cabling remains unchanged. See [Mount the enclosures in the rack](#) on page 15 for more detail.

When connecting multiple expansion enclosures to an expansion enclosure, use reverse cabling to ensure the highest level of fault tolerance.

The ME4 Series identifies controller modules and IOMs by enclosure ID and IOM ID. In the following figure, the controller modules are identified as 0A and 0B, the IOMs in the first expansion enclosure are identified as 1A and 1B, and so on. Controller module 0A is connected to IOM 1A, with a chain of connections cascading down (blue). Controller module 0B is connected to the lower IOM (9B), of the last expansion enclosure, with connections moving in the opposite direction (green). Reverse cabling enables any expansion enclosure to fail—or be removed—while maintaining access to other enclosures.

NOTE: The cabling diagrams show only relevant details such as module face plate outlines and expansion ports.

[Cabling connections between a 2U controller enclosure and 2U expansion enclosures](#) on page 19 shows the maximum cabling configuration for a 2U controller enclosure with 2U expansion enclosures.

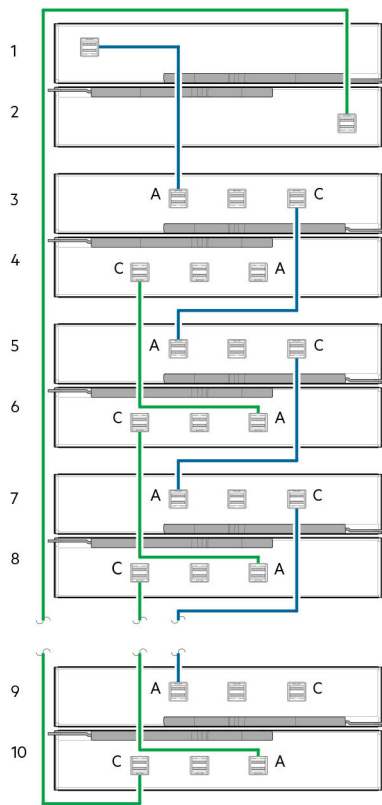


Figure 13. Cabling connections between a 2U controller enclosure and 2U expansion enclosures

- | | |
|-----------------------------|-----------------------------|
| 1. Controller module A (0A) | 2. Controller module B (0B) |
| 3. IOM (1A) | 4. IOM (1B) |
| 5. IOM (2A) | 6. IOM (2B) |
| 7. IOM (3A) | 8. IOM (3B) |
| 9. IOM (9A) | 10. IOM (9B) |

Cabling connections between a 5U controller enclosure and 5U expansion enclosures on page 19 shows the maximum cabling configuration for a 5U84 controller enclosure with 5U84 expansion enclosures (four enclosures including the controller enclosure).




Figure 14. Cabling connections between a 5U controller enclosure and 5U expansion enclosures

- | | |
|-----------------------------|-----------------------------|
| 1. Controller module A (0A) | 2. Controller module B (0B) |
| 3. IOM (1A) | 4. IOM (1B) |
| 5. IOM (2A) | 6. IOM (2B) |
| 7. IOM (3A) | 8. IOM (3B) |

Cabling connections between a 2U controller enclosure and 5U84 expansion enclosures on page 20 shows the maximum cabling configuration for a 2U controller enclosure with 5U84 expansion enclosures (four enclosures including the controller enclosure).

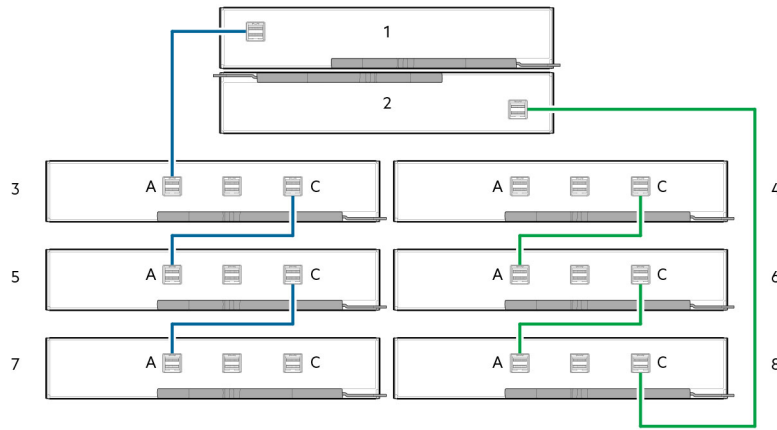


Figure 15. Cabling connections between a 2U controller enclosure and 5U84 expansion enclosures

- | | |
|-----------------------------|-----------------------------|
| 1. Controller module A (0A) | 2. Controller module B (0B) |
| 3. IOM (1A) | 4. IOM (1B) |
| 5. IOM (2A) | 6. IOM (2B) |
| 7. IOM (3A) | 8. IOM (3B) |

Label the back-end cables

Make sure to label the back-end SAS cables that connect the controller enclosure and the expansion enclosures.

Connect to the management network

Perform the following steps to connect a controller enclosure to the management network:

1. Connect an Ethernet cable to the network port on each controller module.
2. Connect the other end of each Ethernet cable to a network that your management host can access, preferably on the same subnet.

NOTE: If you connect the iSCSI and management ports to the same physical switches, Dell EMC recommends using separate VLANs.




Figure 16. Connect a 2U controller enclosure to the management network

- | | |
|--------------------------------|--------------------------------|
| 1. Controller module in slot A | 2. Controller module in slot B |
| 3. Switch | 4. SAN |




Figure 17. Connect a 5U controller enclosure to the management network

- | | |
|--------------------------------|--------------------------------|
| 1. Controller module in slot A | 2. Controller module in slot B |
| 3. Switch | 4. SAN |

NOTE: See also the topic about configuring network ports on controller modules in the *Dell EMC PowerVault ME4 Series Storage System Administrator's Guide*.

Cable host servers to the storage system

This section describes the different ways that host servers can be connected to a storage system.

Topics:

- [Cabling considerations](#)
- [Connecting the enclosure to hosts](#)
- [Host connection](#)

Cabling considerations

Host interface ports on ME4 Series controller enclosures can connect to respective hosts using direct-attach or switch-attach methods.

Another important cabling consideration is cabling controller enclosures to enable the replication feature. The FC and iSCSI product models support replication, but SAS product models do not support replication. See [Cabling for replication](#) on page 90.

Use only Dell EMC cables for host connections:

- Qualified 16 Gb FC SFP+ transceivers and cable options
- Qualified 10 GbE iSCSI SFP+ transceivers and cable options
- Qualified 10Gbase-T cable options
- Qualified 12 Gb mini-SAS HD cable options

Connecting the enclosure to hosts

A *host* identifies an external port to which the storage system is attached. The external port may be a port in an I/O adapter (such as an FC HBA) in a server. Cable connections vary depending on configuration. This section describes host interface protocols supported by ME4 Series controller enclosures, while showing a few common cabling configurations. ME4 Series controllers use Unified LUN Presentation (ULP), which enables a host to access mapped volumes through any controller host port.

ULP can show all LUNs through all host ports on both controllers, and the interconnect information is managed by the controller firmware. ULP appears to the host as an active-active storage system, allowing the host to select any available path to access the LUN, regardless of disk group ownership.

CNC technology

The ME4 Series FC/iSCSI models use Converged Network Controller (CNC) technology.

The CNC technology enables you to select the host interface protocols to use on the storage system. The small form-factor pluggable (SFP+) connectors that are used in CNC ports are further described in the following subsections:

NOTE:

- Controller modules are not always shipped with preinstalled SFP+ transceivers. You might need to install SFP transceivers into the controller modules. Within your product kit, locate the qualified SFP+ transceivers and install them into the CNC ports. See [SFP+ transceiver for FC/iSCSI ports](#) on page 98.
- Use the PowerVault Manager to set the host interface protocol for CNC ports using qualified SFP+ transceivers. ME4 Series models ship with CNC ports configured for FC. When connecting CNC ports to iSCSI hosts, you must configure these ports for iSCSI.

CNC ports used for host connection

ME4 Series SFP+ based controllers ship with CNC ports that are configured for FC.

If you must change the CNC port mode, you can do so using the PowerVault Manager.

Alternatively, the ME4 Series enables you to set the CNC ports to use FC and iSCSI protocols in combination. When configuring a combination of host interface protocols, host ports 0 and 1 must be configured for FC, and host ports 2 and 3 must be configured for iSCSI. The CNC ports must use qualified SFP+ connectors and cables for the selected host interface protocol. For more information, see [SFP+ transceiver for FC/iSCSI ports](#) on page 98.

Fibre Channel protocol

ME4 Series controller enclosures support controller modules with CNC host interface ports.

Using qualified FC SFP+ transceiver/cable options, these CNC ports can be configured to support Fibre Channel protocol in either four or two CNC ports. Supported data rates are 8 Gb/sec or 16 Gb/s.

The controllers support Fibre Channel Arbitrated Loop (public or private) or point-to-point topologies. Loop protocol can be used in a physical loop or for direct connection between two devices. Point-to-point protocol is used to connect to a fabric switch. Point-to-point protocol can also be used for direct connection, and it is the only option supporting direct connection at 16 Gb/s.

The Fibre Channel ports are used for:

- Connecting to FC hosts directly, or through a switch used for the FC traffic.
- Connecting two storage systems through a switch for replication. See [Cabling for replication](#) on page 90.

The first option requires that the host computer must support FC and optionally, multipath I/O.

Use the PowerVault Manager to set FC port speed and options. See the topic about configuring host ports in the *Dell EMC PowerVault ME4 Series Storage System Administrator's Guide*. You can also use CLI commands to perform these actions:

- Use the `set host-parameters` CLI command to set FC port options.
- Use the `show ports` CLI command to view information about host ports.

iSCSI protocol

ME4 Series controller enclosures support controller modules with CNC host interface ports.

CNC ports can be configured to support iSCSI protocol in either four or two CNC ports. The CNC ports support 10 GbE but do not support 1 GbE.

The 10 GbE iSCSI ports are used for:

- Connecting to 10 GbE iSCSI hosts directly, or through a switch used for the 10 GbE iSCSI traffic.
- Connecting two storage systems through a switch for replication.

The first option requires that the host computer supports Ethernet, iSCSI, and optionally, multipath I/O.

See the topic about configuring CHAP in the *Dell EMC PowerVault ME4 Series Storage System Administrator's Guide*.

Use the PowerVault Manager to set iSCSI port options. See the topic about configuring host ports in the *Dell EMC PowerVault ME4 Series Storage System Administrator's Guide*. You can also use CLI commands to perform these actions:

- Use the `set host-parameters` CLI command to set iSCSI port options.
- Use the `show ports` CLI command to view information about host ports.

iSCSI settings

The host should be cabled to two different Ethernet switches for redundancy.

If you are using switches with mixed traffic (LAN/iSCSI), then a VLAN should be created to isolate iSCSI traffic from the rest of the switch traffic.

Example iSCSI port address assignments

The following figure and the supporting tables provide example iSCSI port address assignments featuring two redundant switches and two IPv4 subnets:

NOTE: For each callout number, read across the table row for the addresses in the data path.




Figure 18. Two subnet switch example (IPv4)

Table 5. Two subnet switch example

No.	Device	IP Address	Subnet
1	A0	192.68.10.200	10
2	A1	192.68.11.210	11
3	A2	192.68.10.220	10
4	A3	192.68.11.230	11
5	B0	192.68.10.205	10
6	B1	192.68.11.215	11
7	B2	192.68.10.225	10
8	B3	192.68.11.235	11
9	Switch A	N/A	N/A
10	Switch B	N/A	N/A
11	Host server 1, Port 0	192.68.10.20	10
12	Host server 1, Port 1	192.68.11.20	11
13	Host server 2, Port 0	192.68.10.21	10
14	Host server 2, Port 1	192.68.11.21	11

To enable CHAP, see the topic about configuring CHAP in the *Dell EMC PowerVault ME4 Series Storage System Administrator's Guide*.

SAS protocol

ME4 Series SAS models use 12 Gb/s host interface protocol and qualified cable options for host connection.

12Gb HD mini-SAS host ports

ME4 Series 12 Gb SAS controller enclosures support two controller modules. The 12 Gb/s SAS controller module provides four SFF-8644 HD mini-SAS host ports. These host ports support data rates up to 12 Gb/s. HD mini-SAS host ports are used for attachment to SAS hosts directly. The host computer must support SAS and optionally, multipath I/O. Use a qualified cable option when connecting to a host.

Host connection

ME4 Series controller enclosures support up to eight direct-connect server connections, four per controller module.

Connect appropriate cables from the server HBAs to the controller module host ports as described in the following sections.

16 Gb Fibre Channel host connection

To connect controller modules supporting FC host interface ports to a server HBA or switch, using the controller CNC ports, select a qualified FC SFP+ transceiver. For information about configuring HBAs, see the Fibre Channel topics under [Perform host setup](#) on page 43.

Use the cabling diagrams to connect the host servers to the switches. See the *Dell EMC Storage Support Matrix* for supported Fibre Channel HBAs.

- Install and connect each FC HBA to a switch that is connected to the host ports on the two controllers that are shown in [Connecting hosts: ME4 Series 2U switch-attached – two servers, two switches](#) on page 29 and [Connecting hosts: ME4 Series 5U switch-attached – two servers, two switches](#) on page 29.
- In hybrid examples, one server and switch manage FC traffic, and the other server and switch manage iSCSI traffic.
- For FC, each initiator must be zoned with a single host port or multiple host ports only (single initiator, multi-target of the same kind).

Connecting host servers directly to the storage system is also supported.

Qualified options support cable lengths of 1 m (3.28'), 2 m (6.56'), 5 m (16.40'), 15 m (49.21'), 30 m (98.43'), and 50 m (164.04') for OM4 multimode optical cables and OM3 multimode FC cables. A 0.5 m (1.64') cable length is also supported for OM3. In addition to providing host connection, these cables are used for connecting two storage systems through a switch, to facilitate use of the optional replication feature.

10 GbE iSCSI host connection

To connect controller modules supporting 10 GbE iSCSI host interface ports to a server HBA or switch, using the controller CNC ports, select a qualified 10 GbE SFP+ transceiver. For information about configuring iSCSI initiators/HBAs, see the iSCSI topics under [Perform host setup](#) on page 43.

Use the cabling diagrams to connect the host servers to the switches.

- Install and connect each Ethernet NIC to a switch that is connected to the host ports on the two controllers that are shown in [Connecting hosts: ME4 Series 2U switch-attached – two servers, two switches](#) on page 29 and [Connecting hosts: ME4 Series 5U switch-attached – two servers, two switches](#) on page 29.
- In hybrid examples, one server and switch manage iSCSI traffic, and the other server and switch manage FC traffic.

Connecting host servers directly to the storage system is also supported.

10Gbase-T host connection

To connect controller modules with 10Gbase-T iSCSI host interface ports to a server HBA or switch, select a qualified 10Gbase-T cable option.

For information about configuring network adapters and iSCSI HBAs, see the iSCSI topics under [Perform host setup](#) on page 43. See also, the cabling instructions in [10 GbE iSCSI host connection](#) on page 25.

12 Gb HD mini-SAS host connection

To connect controller modules supporting HD mini-SAS host interface ports to a server HBA, using the SFF-8644 dual HD mini-SAS host ports on a controller, select a qualified HD mini-SAS cable option. For information about configuring SAS HBAs, see the SAS topics under [Perform host setup](#) on page 43.

A qualified SFF-8644 to SFF-8644 cable option is used for connecting to a 12Gb/s enabled host. Qualified SFF-8644 to SFF-8644 options support cable lengths of 0.5 m (1.64'), 1 m (3.28'), 2 m (6.56'), and 4 m (13.12').

Connecting direct attach configurations

A dual-controller configuration improves application availability. If a controller failure occurs, the affected controller fails over to the healthy partner controller with little interruption to data flow.

A failed controller can be replaced without the need to shut down the storage system.

NOTE: In the following examples, a single diagram represents CNC, SAS, and 10Gbase-T host connections for ME4 Series controller enclosures. The location and sizes of the host ports are similar. Blue cables show controller A paths and green cables show controller B paths for host connection.

Single-controller module configurations

A single controller module configuration does not provide redundancy if a controller module fails.

This configuration is intended only for environments where high availability is not required. If the controller module fails, the host loses access to the storage data until failure recovery actions are completed.

NOTE: Expansion enclosures are not supported in a single controller module configuration.




Figure 19. Connecting hosts: ME4 Series 2U direct attach – one server, one HBA, single path

1. Server
2. Controller module in slot A
3. Controller module blank in slot B

NOTE: If the ME4 Series 2U controller enclosure is configured with a single controller module, the controller module must be installed in the upper slot. A controller module blank must be installed in the lower slot. This configuration is required to enable sufficient air flow through the enclosure during operation.

Dual-controller module configurations

A dual-controller module configuration improves application availability.

If a controller module failure occurs, the affected controller module fails over to the partner controller module with little interruption to data flow. A failed controller module can be replaced without the need to shut down the storage system.

In a dual-controller module system, hosts use LUN-identifying information from both controller modules to determine the data paths are available to a volume. Assuming MPIO software is installed, a host can use any available data path to access a volume that is owned by either controller module. The path providing the best performance is through the host ports on the controller module that owns the volume. Both controller modules share one set of 1,024 LUNs (0-1,023) for use in mapping volumes to hosts.

Dual-controller module configurations – directly attached

In the following figures, blue cables show controller module A paths, and green cables show controller module B paths for host connection:




Figure 20. Connecting hosts: ME4 Series 2U direct attach – one server, one HBA, dual path

1. Server
2. Controller module in slot A
3. Controller module in slot B




Figure 21. Connecting hosts: ME4 Series 5U direct attach – one server, one HBA, dual path

1. Server
2. Controller module in slot A
3. Controller module in slot B




Figure 22. Connecting hosts: ME4 Series 2U direct attach – two servers, one HBA per server, dual path

1. Server 1
2. Server 2
3. Controller module in slot A
4. Controller module in slot B




Figure 23. Connecting hosts: ME4 Series 5U direct attach – two servers, one HBA per server, dual path

1. Server 1
2. Server 2
3. Controller module in slot A
4. Controller module in slot B




Figure 24. Connecting hosts: ME4 Series 2U direct attach— four servers, one HBA per server, dual path

- | | |
|------------------------|------------------------|
| 1. Server 1 | 2. Server 2 |
| 3. Server 3 | 4. Server 4 |
| 5. Controller module A | 6. Controller module B |




Figure 25. Connecting hosts: ME4 Series 5U direct attach – four servers, one HBA per server, dual path

- | | |
|------------------------|------------------------|
| 1. Server 1 | 2. Server 2 |
| 3. Server 3 | 4. Server 4 |
| 5. Controller module A | 6. Controller module B |

Dual-controller module configurations – switch-attached

A switch-attached solution—or SAN—places a switch between the servers and the controller enclosures within the storage system. Using switches, a SAN shares a storage system among multiple servers, reducing the number of storage systems required for a particular environment. Using switches increases the number of servers that can be connected to the storage system.

- i NOTE:** About switch-attached configurations:
- See the recommended switch-attached examples for host connection in the *Setting Up Your Dell EMC PowerVault ME4 Series Storage System* document that is provided with your controller enclosure.
 - See [Two subnet switch example \(IPv4\)](#) on page 24 for an example showing host port and controller port addressing on an IPv4 network.




Figure 26. Connecting hosts: ME4 Series 2U switch-attached – two servers, two switches

- | | |
|------------------------|------------------------|
| 1. Server 1 | 2. Server 2 |
| 3. Switch A | 4. Switch B |
| 5. Controller module A | 6. Controller module B |




Figure 27. Connecting hosts: ME4 Series 5U switch-attached – two servers, two switches

- | | |
|------------------------|------------------------|
| 1. Server 1 | 2. Server 2 |
| 3. Switch A | 4. Switch B |
| 5. Controller module A | 6. Controller module B |

Label the front-end cables

Make sure to label the front-end cables to identify the controller module and host interface port to which each cable connects.

Connect power cables and power on the storage system

Before powering on the enclosure system, ensure that all modules are firmly seated in their correct slots.

Verify that you have successfully completed the [Installation checklist](#) on page 9 instructions. Once you have completed steps 1–7, you can access the management interfaces using your web-browser to complete the system setup.

Topics:

- [Power cable connection](#)

Power cable connection

Connect a power cable from each PCM or PSU on the enclosure rear panel to the PDU (power distribution unit) as shown in the following figures:




Figure 28. Typical AC power cable connection from PDU to PCM (2U)

1. Controller enclosure with redundant PCMs
2. Redundant PCM to PDU (AC UPS shown) connection




Figure 29. Typical AC power cable connection from PDU to PSU (5U)

1. Controller enclosure with redundant PSUs
2. Redundant PSU to PDU (AC UPS shown) connection

NOTE: The power cables must be connected to at least two separate and independent power supplies to ensure redundancy. When the storage system is ready for operation, ensure that each PCM or PSU power switch is set to the On position. See also [Powering on](#) on page 31).

CAUTION: Always remove the power connections before you remove the PCM (2U) or PSU (5U84) from the enclosure.

Testing enclosure connections

See [Powering on](#) on page 31. Once the power-on sequence succeeds, the storage system is ready to be connected as described in [Connecting the enclosure to hosts](#) on page 22.

Grounding checks

The enclosure system must be connected to a power source that has a safety electrical grounding connection.

CAUTION: If more than one enclosure goes in a rack, the importance of the grounding connection to the rack increases because the rack has a larger Grounding Leakage Current (Touch Current). Examine the grounding connection to the rack before power on. An electrical engineer who is qualified to the appropriate local and national standards must do the examination.

Powering on

CAUTION: Do not operate the enclosure system until the ambient temperature is within the specified operating range that is described in the system specifications section of the *Dell EMC PowerVault ME4 Series Storage System Owner's Manual*. If the drive modules have been recently installed, ensure that they have had time to adjust to the environmental conditions before they are used with production data for I/O.

- With 2U enclosures, power on the storage system by connecting the power cables from the PCMs to the PDU, and moving the power switch on each PCM to the On position. See [Typical AC power cable connection from PDU to PCM \(2U\)](#) on page 30.

The System Power LED on the 2U Ops panel should be lit green when the enclosure power is activated.

- With 5U84 enclosures, power on the storage system by connecting the power cables from the PSUs to the PDU, and moving the power switch on each PSU to the On position. See [Typical AC power cable connection from PDU to PSU \(5U\)](#) on page 30.

The Power on/Standby LED on the 5U84 Ops panel should be lit green when the enclosure power is activated.

- When powering up, ensure to power up the enclosures and associated data host in the following order:
 - Drive enclosures first – Ensures that the disks in the drive enclosure have enough time to completely spin up before being scanned by the controller modules within the controller enclosure. The LEDs blink while the enclosures power up. After the LEDs stop blinking – if the LEDs on the front and back of the enclosure are not amber – the power-on sequence is complete, and no faults have been detected.
 - Controller enclosure next – Depending upon the number and type of disks in the system, it may take several minutes for the system to become ready.
 - Data host last (if powered off for maintenance purposes).

When powering off, reverse the order of steps that are used for powering on.

NOTE: If main power is lost for any reason, the system automatically restarts when power is restored.

Enclosure Ops panels

- See [2U enclosure Ops panel](#) on page 71 for details pertaining to 2U Ops panel LEDs and related fault conditions.
- See [5U enclosure Ops panel](#) on page 72 for details pertaining to 5U84 Ops panel LEDs and related fault conditions.

Guidelines for powering enclosures on and off

- Remove the AC cord before inserting or removing a PCM (2U) or PSU (5U84).
- Move the PCM or PSU switch to the Off position before connecting or disconnecting the AC power cable.
- Allow 15 seconds between powering off and powering on the PCM or PSU.
- Allow 15 seconds before powering on one PSU or PCM in the system, and powering off another PCM or PSU.
- Never power off a PCM or PSU while any amber LED is lit on the partner PCM or PSU.

- A 5U84 enclosure must be left in a power on state for 30 seconds following resumption from standby before the enclosure can be placed into standby again.
- Although the enclosure supports standby, the expansion module shuts off completely during standby and cannot receive a user command to power back on. An AC power cycle is the only method to return the 5U84 to full power from standby.

Perform system and storage setup

The following sections describe how to setup a Dell EMC PowerVault ME4 Series storage system:

Topics:

- [Record storage system information](#)
- [Using guided setup](#)

Record storage system information

Use the [System Information Worksheet](#) on page 100 to record the information that you need to install the ME4 Series storage system.


Using guided setup

Upon completing the hardware installation, use PowerVault Manager to configure, provision, monitor, and manage the storage system.

When first accessing the PowerVault Manager, perform a firmware update before configuring your system. After the firmware update is complete, use the guided setup to verify the web browser requirements and then access the PowerVault Manager.

Web browser requirements and setup

The PowerVault Manager web interface requires Mozilla Firefox 57 or later, Google Chrome 57 or later, Microsoft Internet Explorer 10 or 11, or Apple Safari 10.1 or later.

 **NOTE:** You cannot view PowerVault Manager help content if you are using the Microsoft Edge browser that ships with Windows 10.

- To see the help window, you must enable pop-up windows.
- To optimize the display, use a color monitor and set its color quality to the highest setting.
- Do not use the Back, Forward, Reload, or Refresh buttons in the browser. The PowerVault Manager has a single page for which content changes as you perform tasks and automatically updates to show current data.
- To navigate past the Sign In page (with a valid user account):
 - Verify that cookies are allowed for the IP address of each controller network port.
 - For Internet Explorer, set the local-intranet security option on the browser to medium or medium-low.
 - For Internet Explorer, add each network IP address for each controller as a trusted site.
 - For HTTPS, ensure that Internet Explorer is set to use TLS 1.2.

Access the PowerVault Manager

Do not turn on more than one unconfigured controller enclosure at a time to avoid IP conflicts.

1. Temporarily set the management host NIC to a 10.0.0.x address or to the same IPv6 subnet to enable communication with the storage system.
2. In a supported web browser:
 - Type `https://10.0.0.2` to access controller module A on an IPv4 network.
 - Type `https://fd6e:23ce:fed3:19d1::1` to access controller module A on an IPv6 network.
3. If the storage system is running G275 firmware:
 - a. Sign in to the PowerVault Manager using the following user name and password:
 - User name: **manage**

- Password: **!manage**

- Read the Commercial Terms of Sale and End User License Agreement, and click **Accept**.

The storage system displays the Welcome panel. The Welcome panel provides options for setting up and provisioning your storage system.

- If the storage system is running G280 firmware:
 - Click **Get Started**.
 - Read the Commercial Terms of Sale and End User License Agreement, and click **Accept**.
 - Type a new user name for the storage system in the **Username** field.

The user name requirements are described in the *Dell EMC PowerVault ME4 Series Storage System Administrator's Guide*.

- Type password for the new username in the **Password** and **Confirm Password** fields.

The password requirements are described in the *Dell EMC PowerVault ME4 Series Storage System Administrator's Guide*.

- Click **Apply and Continue**.

The storage system creates the user and displays the Welcome panel. The Welcome panel provides options for setting up and provisioning your storage system.

NOTE: If you are unable to use the 10.0.0.x network to configure the storage system, see [Setting network port IP addresses using the CLI port and serial cable](#) on page 103.

Update firmware

After powering on the storage system for the first time, verify that the controller modules, expansion modules, and disk drives are using the current firmware release.

NOTE: Expansion module firmware is updated automatically with controller module updates.

- Using the PowerVault Manager, select **Action > Update Firmware** in the System topic.

The **Update Firmware** panel opens. The Update Controller Modules tab shows versions of firmware components that are installed in each controller module.

- Locate firmware updates at www.dell.com/support. If newer versions of the firmware are available, download the bundle file or relevant firmware component files.
- Click **Browse**, select the firmware bundle file or component file to install, and then click **OK**.

When the update is complete, the system restarts.

Use guided setup in the PowerVault Manager Welcome panel

The Welcome panel provides options for you to quickly set up your system by guiding you through the configuration and provisioning process.

With guided setup, you must first configure your system settings by accessing the System Settings panel and completing all required options. After these options are complete, you can provision your system by accessing the Storage Setup panel and the Host Setup panel and completing the wizards.


The Welcome panel also displays the health of the system. If the health of the system is degraded or faulty, you can click **System Information** to access the System topic. In the System topic, you can view information about each enclosure, including its physical components, in front, rear, and tabular views.

If the system detects that it has only one controller, its health shows as degraded. If you are operating the system with a single controller, acknowledge this message in the panel.

If you installed two controllers, click **System Information** to diagnose the problem. If the system health is degraded, you can still to configure and provision the system. However, if the health of the system is bad, you cannot configure and provision the system until you resolve the problem affecting system health.

To use guided setup:

- From the Welcome panel, click **System Settings**.
- Choose options to configure your system.

 **NOTE:** Tabs with a red asterisk next to them contain required settings.

3. Save your settings and exit System Settings to return to the Welcome panel.
4. Click **Storage Setup** to access the Storage Setup wizard and follow the prompts to begin provisioning your system by creating disk groups and pools. For more information about using the Storage Setup wizard, see [Configuring storage setup](#) on page 40.
5. Save your settings and exit Storage Setup to return to the Welcome panel.
6. Click **Host Setup** to access the Host Setup wizard and follow the prompts to continue provisioning your system by attaching hosts.
For more information, see [Host system requirements](#) on page 43.

Configuring system settings

The System Settings panel provides options for you to quickly configure your system.


Navigate the options by clicking the tabs on the left side of the panel. Tabs with a red asterisk next to them are required. To apply and save changes, click **Apply**. To apply changes and close the panel, click **Apply and Close**.

At a minimum, Dell EMC recommends that you perform the following actions:

- [Configuring controller network ports](#) on page 35
- [Setting up system notifications](#) on page 37
- [Setting up SupportAssist and CloudIQ](#) on page 37
- [Changing host port settings](#) on page 38

Configuring controller network ports

You can manually set static IP address parameters for network ports or you can specify that IP addresses be set automatically. IP addresses can be set automatically using DHCP for IPv4 or Auto for IPv6, which uses DHCPv6 and/or SLAAC.

 **NOTE:** If you used the default 10.0.0.2/10.0.0.3 addresses to access the guided setup, consider changing those IPv4 addresses to avoid an IP conflict if you have more than one ME4 Series array on your network.

When setting IP values, you can choose either IPv4 or IPv6 formatting for each controller. You can also set the addressing mode and IP version differently for each controller and use them concurrently. For example, you could set IPv4 on controller A to Manual to enable static IP address, and IPv6 on controller B to Auto to enable automatic IP address.

When using DHCP mode, the system obtains values for the network port IP address, subnet mask, and gateway from a DHCP server if one is available. If a DHCP server is unavailable, current address is unchanged. You must have some means of determining what addresses have been assigned, such as the list of bindings on the DHCP server. When using Auto mode, addresses are retrieved from both DHCP and Stateless address auto-configuration (SLAAC). DNS settings are also automatically retrieved from the network.

Each controller has the following factory-default IP settings:

- IP address source: Manual
- Controller A IP address: 10.0.0.2
- Controller B IP address: 10.0.0.3
- IP subnet mask: 255.255.255.0
- Gateway IP address: 10.0.0.1

When DHCP is enabled in the storage system, the following initial values are set and remain set until the system can contact a DHCP server for new addresses:

- Controller IP addresses: 169.254.x.x (where the value of x.x is the lowest 16 bits of the controller serial number)
- IP subnet mask: 255.255.0.0
- Gateway IP address: 10.0.0.0

169.254.x.x addresses (including gateway 169.254.0.1) are on a private subnet that is reserved for unconfigured systems and the addresses are not routable. This prevents the DHCP server from reassigning the addresses and possibly causing a conflict where two controllers have the same IP address. As soon as possible, change these IP values to proper values for your network.

For IPv6, when Manual mode is enabled you can enter up to four static IP addresses for each controller. When Auto is enabled, the following initial values are set and remain set until the system can contact a DHCPv6 and/or SLAAC server for new addresses:


- Controller A IP address: fd6e:23ce:fed3:19d1::1
- Controller B IP address: fd6e:23ce:fed3:19d1::2
- Gateway IP address: fd6e:23ce:fed3:19d1::3

 **CAUTION:** Changing IP settings can cause management hosts to lose access to the storage system after the changes are applied in the confirmation step.

Set IPv4 addresses for network ports

Perform the following steps to set IPv4 addresses for the network ports:

1. In the **Welcome** panel, select **System Settings**, and then click the **Network** tab.
2. Select the **IPv4** tab.
IPv4 uses 32-bit addresses.
3. Select the type of IP address settings to use for each controller from the **Source** drop-down menu:
 - Select **Manual** to specify static IP addresses.
 - Select **DHCP** to allow the system to automatically obtain IP addresses from a DHCP server.
4. If you selected **Manual**, perform the following steps: , and then
 - a. Type the IP address, IP mask, and Gateway addresses for each controller.
 - b. Record the IP addresses.


 **NOTE:** The following IP addresses are reserved for internal use by the storage system: 169.254.255.1, 169.254.255.2, 169.254.255.3, 169.254.255.4, and 127.0.0.1. Because these addresses are routable, do not use them anywhere in your network.

5. If you selected **DHCP**, complete the remaining steps to allow the controller to obtain IP addresses from a DHCP server.
6. Click **Apply**.
A confirmation panel appears.
7. Click **OK**.
If you selected **DHCP** and the controllers successfully obtained IP addresses from the DHCP server, the new IP addresses are displayed.
8. Sign out to use the new IP address to access PowerVault Manager.

Set IPv6 values for network ports

Perform the following steps to set IPv6 addresses for the network ports:

1. In the **Welcome** panel, select **System Settings**, and then click the **Network** tab.
2. Select the **IPv6** tab.
IPv6 uses 128-bit addresses.
3. Select the type of IP address settings to use for each controller from the **Source** drop-down menu:
 - Select **Manual** to specify up to four static IP addresses for each controller.
 - Select **Auto** to allow the system to automatically obtain IP addresses.
4. If you chose **Manual**, perform the following steps for each controller:
 - a. Click **Add Address**.
 - b. Type the IPv6 addresses in the **IP Address** field.
 - c. Type a label for the IP address in the **Address Label** field.
 - d. Click **Add**.
 - e. Record the IPv6 address.

 **NOTE:** The following IP addresses are reserved for internal use by the storage system: 169.254.255.1, 169.254.255.2, 169.254.255.3, 169.254.255.4, and 127.0.0.1. Because these addresses are routable, do not use them anywhere in your network.

5. If you selected **Auto**, complete the remaining steps to allow the controllers to obtain IP addresses.
6. Click **Apply**.
A confirmation panel appears.
7. Click **OK**.

8. Sign out and use the new IP address to access PowerVault Manager.

Setting up system notifications

Dell EMC recommends enabling at least one notification service to monitor the system.


Enable email notifications

Perform the following steps to enable email notifications:

1. In the Welcome panel, select **System Settings**, and then click the **Notifications** tab.
2. Select the Email tab and ensure that the SMTP Server and SMTP Domain options are set.
3. Set the email notification:
 - To enable email notifications, select the **Enable Email Notifications** check box.
 - To disable email notifications, clear the **Enable Email Notifications** check box.
4. If email notification is enabled, select the minimum severity for which the system should send email notifications: **Critical** (only); **Error** (and Critical); **Warning** (and Error and Critical); **Resolved** (and Error, Critical, and Warning); Informational (all).
5. If email notification is enabled, in one or more of the Email Address fields enter an email address to which the system should send notifications. Each email address must use the format `user-name@domain-name`. Each email address can have a maximum of 320 bytes. For example: **Admin@mydomain.com** or **IT-team@mydomain.com**.
6. Perform one of the following:
 - To save your settings and continue configuring your system, click **Apply**.
 - To save your settings and close the panel, click **Apply and Close**.A confirmation panel is displayed.
7. Click **OK** to save your changes. Otherwise, click **Cancel**.

Test notification settings

Perform the following steps to test notifications:

1. Configure your system to receive trap and email notifications.
2. Click **Send Test Event**. A test notification is sent to each configured trap host and email address.
3. Verify that the test notification reached each configured email address.
 **NOTE:** If there was an error in sending a test notification, event 611 is displayed in the confirmation.


Setting up SupportAssist and CloudIQ

SupportAssist provides an enhanced support experience for ME4 Series storage systems by sending configuration and diagnostic information to technical support at regular intervals. CloudIQ provides storage monitoring and proactive service, giving you information that is tailored to your needs, access to near real-time analytics, and the ability to monitor storage systems from anywhere at any time.

Perform the following steps to set up SupportAssist and enable CloudIQ:

1. In the Welcome panel, select **System Settings**, and then click the **SupportAssist** tab.
2. Select the **SupportAssist** checkbox to enable SupportAssist for the storage system.
The SupportAssist agreement is displayed.
3. Read through the agreement, then acknowledge it by clicking **Accept**.

The system attempts to establish connectivity with the remote support server. Once connectivity is established, the system collects an initial full debug log dump and sends it to the SupportAssist server.


 **NOTE:** If the system cannot contact the remote support server, an error message is displayed that contains details about the connection failure and provides recommended actions.

4. In the **Contact Information** tab, type the primary contact information and select the preferred contact settings.
To receive email messages when a storage system issue occurs, select the **Yes, I would like to receive emails from SupportAssist when issues arise, including hardware failure notifications** checkbox.

5. If the storage array does not have direct access to the Internet, you can use a web proxy server to send SupportAssist data to technical support.

To use a web proxy, click the **Web Proxy** tab, select the **Web Proxy** checkbox, and type the web proxy server settings in the appropriate fields.

6. To enable CloudIQ, click the **CloudIQ Settings** tab and select the **Enable CloudIQ** checkbox.

 **NOTE:** For more information about CloudIQ, contact technical support or go to the [CloudIQ](#) product page.

7. Click **Apply** or **Apply and Close**, and click **OK** on the confirmation panel.

Changing host port settings


You can configure controller host-interface settings for ports except for systems with a 4-port SAS controller module or 10Gbbase-T iSCSI controller module.

To enable the system to communicate with hosts, you must configure the host-interface options on the system.

For a system with a 4-port SAS controller module or 10Gbbase-T iSCSI controller module, there are no host-interface options.

For a system with 4-port SFP+ controller modules (CNC), all host ports ship from the factory in Fibre Channel (FC) mode. However, the ports can be configured as a combination of FC or iSCSI ports. FC ports support use of qualified 16 Gb/s SFP transceivers. You can set FC ports to auto-negotiate the link speed or to use a specific link speed. iSCSI ports support use of qualified 10 Gb/s SFP transceivers.

For information about setting host parameters such as FC port topology, and the host-port mode, see the *Dell EMC PowerVault ME4 Series Storage System CLI Reference Guide*.

 **NOTE:** If the current settings are correct, port configuration is optional.

Configure FC ports

Perform the following steps to configure FC ports:

1. In the Welcome panel, select **System Settings**, and then click the **Ports** tab.
2. On the **Port Settings** tab, set the port-specific options:
 - Set the **Speed** option to the proper value to communicate with the host, or to auto, which auto-negotiates the proper link speed. Because a speed mismatch prevents communication between the port and host, set a speed only if you need to force the port to use a known speed.
 - Set the **Connection Mode** to either point-to-point or auto:
 - **point-to-point** – Fibre Channel point-to-point.
 - **auto** – Automatically sets the mode based on the detected connection type.
3. Perform one of the following:
 - To save your settings and continue configuring your system, click **Apply**.
 - To save your settings and close the panel, click **Apply and Close**.A confirmation panel appears.
4. Click **OK**.

Configure iSCSI ports

Perform the following steps to configure iSCSI ports:

1. In the Welcome panel, select **System Settings**, and then click the **Ports** tab.
2. On the **Port Settings** tab, set the port-specific options:
 - IP Address. For IPv4 or IPv6, the port IP address. For corresponding ports in each controller, assign one port to one subnet and the other port to a second subnet. Ensure that each iSCSI host port in the storage system is assigned a different IP address. For example, in a system using IPv4:
 - Controller A port 0: 10.10.10.100
 - Controller A port 1: 10.11.10.120
 - Controller B port 0: 10.10.10.110
 - Controller B port 1: 10.11.10.130
 - Controller A port 2: 10.10.10.200

- Controller A port 3: 10.11.10.220
 - Controller B port 2: 10.10.10.210
 - Controller B port 3: 10.11.10.230
 - Netmask: For IPv4, subnet mask for assigned port IP address.
 - Gateway: For IPv4, gateway IP address for assigned port IP address.
 - Default Router: For IPv6, default router for assigned port IP address.
3. In the Advanced Settings section of the panel, set the options that apply to all iSCSI ports:

Table 6. Options for iSCSI ports

Enable Authentication (CHAP)	Enables or disables use of Challenge Handshake Authentication Protocol. Enabling or disabling CHAP in this panel updates the setting in the Configure CHAP panel (available in the Hosts topic by selecting Action > Configure CHAP). CHAP is disabled by default.
Link Speed	<ul style="list-style-type: none"> ● auto—Auto-negotiates the proper speed. ● 1 Gb/s—Forces the speed to 1 Gbit/sec, overriding a downshift that can occur during auto-negotiation with 1 Gb/sec HBAs. This setting does not apply to 10 Gb/sec SFPs.
Enable Jumbo Frames	Enables or disables support for jumbo frames. Allowing for 100 bytes of overhead, a normal frame can contain a 1400-byte payload whereas a jumbo frame can contain a maximum 8900-byte payload for larger data transfers. i NOTE: Use of jumbo frames can succeed only if jumbo-frame support is enabled on all network components in the data path.
iSCSI IP Version	Specifies whether IP values use Internet Protocol version 4 (IPv4) or version 6 (IPv6) format. IPv4 uses 32-bit addresses. IPv6 uses 128-bit addresses.
Enable iSNS	Enables or disables registration with a specified Internet Storage Name Service server, which provides name-to-IP-address mapping.
iSNS Address	Specifies the IP address of an iSNS server.
Alternate iSNS Address	Specifies the IP address of an alternate iSNS server, which can be on a different subnet. ⚠ CAUTION: Changing IP settings can cause data hosts to lose access to the storage system.

4. Perform one of the following:
- To save your settings and continue configuring your system, click **Apply**.
 - To save your settings and close the panel, click **Apply and Close**.
- A confirmation panel is displayed.
5. Click **Yes** to save your changes. Otherwise, click **No**.

Configure two ports as FC and two ports as iSCSI per controller

Perform the following steps on each controller to configure two ports as FC and two ports as iSCSI:

1. In the Welcome panel, select **System Settings**, and then click the **Ports** tab.
2. From the Host Post Mode list, select **FC-and-iSCSI**.
i **NOTE:** Ports 0 and 1 are FC ports. Ports 2 and 3 are iSCSI ports.
3. Set the FC port-specific options:
 - Set the **Speed** option to the proper value to communicate with the host, or to auto, which auto-negotiates the proper link speed. A speed mismatch prevents communication between the port and host. Set a speed only if you want to force the port to use a known speed.
 - Set the **FC Connection Mode** to either point-to-point or auto:
 - **point-to-point:** Fibre Channel point-to-point.
 - **auto:** Automatically sets the mode that is based on the detected connection type.
4. Set the iSCSI port-specific options:

Table 7. iSCSI port-specific options

IP Address	For IPv4 or IPv6, the port IP address. For corresponding ports in each controller, assign one port to one subnet and the other port to a second subnet. Ensure that each iSCSI host port in the storage system is assigned a different IP address. For example, in a system using IPv4: <ul style="list-style-type: none"> • Controller A port 2: 10.10.10.100 • Controller A port 3: 10.11.10.120 • Controller B port 2: 10.10.10.110 • Controller B port 3: 10.11.10.130
Netmask	For IPv4, subnet mask for assigned port IP address.
Gateway	For IPv4, gateway IP address for assigned port IP address.
Default Router	For IPv6, default router for assigned port IP address.

- In the Advanced Settings section of the panel, set the options that apply to all iSCSI ports:
 - Enable Authentication (CHAP). Enables or disables use of Challenge Handshake Authentication Protocol. Enabling or disabling CHAP in this panel updates the setting in the Configure CHAP panel (available in the Hosts topic by selecting **Action > Configure CHAP**). CHAP is disabled by default.
 - Link Speed.
 - auto—Auto-negotiates the proper speed.
 - 1 Gb/s— This setting does not apply to 10 Gb/sec HBAs.
 - Enable Jumbo Frames: Enables or disables support for jumbo frames. Allowing for 100 bytes of overhead, a normal frame can contain a 1400-byte payload whereas a jumbo frame can contain a maximum 8900-byte payload for larger data transfers.
 - NOTE:** Use of jumbo frames can succeed only if jumbo-frame support is enabled on all network components in the data path.
 - iSCSI IP Version: Specifies whether IP values use Internet Protocol version 4 (IPv4) or version 6 (IPv6) format. IPv4 uses 32-bit addresses. IPv6 uses 128-bit addresses.
 - Enable iSNS: Enables or disables registration with a specified Internet Storage Name Service server, which provides name-to-IP-address mapping.
 - iSNS Address: Specifies the IP address of an iSNS server.
 - Alternate iSNS Address: Specifies the IP address of an alternate iSNS server, which can be on a different subnet.
 - CAUTION:** Changing IP settings can cause data hosts to lose access to the storage system.
- Perform one of the following:
 - To save your settings and continue configuring your system, click **Apply**.
 - To save your settings and close the panel, click **Apply and Close**.
A confirmation panel is displayed.
- Click **OK** to save your changes. Otherwise, click **Yes**.

Configuring storage setup

The Storage Setup wizard guides you through each step of creating disk groups and pools in preparation for attaching hosts and volumes.

NOTE: You can cancel the wizard at any time, but the changes that are made in completed steps are saved.

Access the Storage Setup wizard from the Welcome panel or by choosing **Action > Storage Setup**. When you access the wizard, you must select the storage type for your environment. After selecting a storage type, you are guided through the steps to create disk groups and pools. The panels that are displayed and the options within them are dependent upon:

- Whether you select a virtual or linear storage type
- Whether the system is brand new (all disks are empty and available and no pools have been created)
- Whether the system has any pools
- Whether you are experienced with storage provisioning and want to set up your disk groups in a certain way


On-screen directions guide you through the provisioning process.


Select the storage type

When you first access the wizard, you are prompted to select the type of storage to use for your environment.

Read through the options and make your selection, and then click **Next** to proceed.

- Virtual storage supports the following features:
 - Tiering
 - Snapshots
 - Replication
 - Thin provisioning
 - One pool per installed RAID controller and up to 16 disk groups per pool
 - Maximum 1 PB usable capacity per pool with large pools feature enabled
 - RAID levels 1, 5, 6, 10, and ADAPT
 - Adding individual disks to increase RAID capacity is only supported for ADAPT disk groups
 - Capacity can be increased by adding additional RAID disk groups
 - Page size is static (4 MB)
 - SSD read cache
 - Global and/or dynamic hot spares
- Linear storage supports the following features:
 - Up to 32 pools per installed RAID controller and one disk group per pool
 - RAID levels 0, 1, 3, 5, 6, 10, 50, ADAPT, and NRAID
 - Adding individual disks to increase RAID capacity is supported for RAID 0, 3, 5, 6, 10, 50, and ADAPT disk groups
 - Configurable chunk size per disk group
 - Global, dedicated, and/or dynamic hot spares

 **NOTE:** Dell EMC recommends using virtual storage.

 **NOTE:** After you create a disk group using one storage type, the system will use that storage type for additional disk groups. To switch to the other storage type, you must first remove all disk groups.

Creating disk groups and pools

The panel that is displayed when creating disk groups and pools is dependent upon whether you are operating in a virtual storage environment or a linear storage environment.

Virtual storage environments

If you are operating in a virtual storage environment, the system scans all available disks, recommends one optimal storage configuration, and displays the suggested disk group layout within the panel.

In a virtual storage environment, the storage system automatically groups disk groups by pool and tier. The disk groups also include a description of the total size and number of disks to be provisioned, including the configuration of spares and unused disks.

If the system is unable to determine a valid storage configuration, the wizard lists the reasons why and provides directions on how to achieve a proper configuration. If the system is unhealthy, an error is displayed along with a description of how to fix it. Follow the recommendations in the wizard to correct the errors, then click **Rescan** to view the optimized configuration.

For a system with no pools provisioned, if you are satisfied with the recommended configuration, click **Create Pools** to provision the system as displayed in the panel and move on to attaching hosts. For a system that contains a pool, if you are satisfied with the recommended configuration, click **Expand Pools** to provision the system as displayed in the panel.

If your environment requires a unique setup, click **Go To Advanced Configuration** to access the Create Advanced Pools panel. Select Add Disk Group and follow the instructions to manually create disk groups one disk at a time. Select Manage Spares and follow the instructions to manually select global spares.

Linear storage environments

If you are operating in a linear storage environment, the Create Advanced Pools panel opens.

Select **Add Disk Groups** and follow the instructions to manually create disk groups one at a time. Select **Manage Spares** and follow the instructions to manually select global spares. Click the icon for more information about options presented.

Open the guided disk group and pool creation wizard

Perform the following steps to open the disk group and pool creation wizard:

1. Access Storage Setup by performing one of the following actions:
 - From the Welcome panel, click **Storage Setup**.
 - From the Home topic, click **Action > Storage Setup**.
2. Follow the on-screen directions to provision your system.

Perform host setup

This section describes how to perform host setup for Dell EMC PowerVault ME4 Series storage systems.

Dell EMC recommends performing host setup on only one host at a time.

For a list of supported HBAs or iSCSI network adapters, see the *Dell EMC PowerVault ME4 Series Storage System Support Matrix*.

For more information, see the topics about initiators, hosts, and host groups, and attaching hosts and volumes in the *Dell EMC PowerVault ME4 Series Storage System Administrator's Guide*.

Topics:

- [Host system requirements](#)
- [Windows hosts](#)
- [Linux hosts](#)
- [VMware ESXi hosts](#)
- [Citrix XenServer hosts](#)

Host system requirements

Hosts connected to ME4 Series controller enclosures must meet the following requirements:

Depending on your system configuration, host operating systems may require that multipathing is supported.

If fault tolerance is required, then multipathing software may be required. Host-based multipath software should be used in any configuration where two logical paths between the host and any storage volume may exist simultaneously. This includes most configurations where there are multiple connections to the host or multiple connections between a switch and the storage.

About multipath configuration

ME4 Series storage systems comply with the SCSI-3 standard for Asymmetrical Logical Unit Access (ALUA).

ALUA-compliant storage systems provide optimal and non-optimal path information to the host during device discovery. To implement ALUA, you must configure your servers to use multipath I/O (MPIO).

Windows hosts

Ensure that the HBAs or network adapters are installed, the drivers are installed, and the latest supported BIOS and firmware are installed.

Configuring a Windows host with FC HBAs

The following sections describe how to configure a Windows host with Fibre Channel (FC) HBAs :


Prerequisites

- Complete the PowerVault Manager guided system and storage setup process.
- Refer to the cabling diagrams within this guide before attaching a host to the storage system; careful planning ensures a successful deployment.


Attach a Windows host with FC HBAs to the storage system

Perform the following steps to attach the Windows host with Fibre Channel (FC) HBAs to the storage system:

1. Ensure that all HBAs have the latest supported firmware and drivers as described on Dell.com/support. For a list of supported FC HBAs, see the *Dell EMC ME4 Series Storage System Support Matrix* on Dell.com/support.
2. Use the FC cabling diagrams to cable the hosts to the storage system either by using switches or connecting the hosts directly to the storage system.
3. Install MPIO on the FC hosts:
 - a. Open the Server Manager.
 - b. Click **Add Roles and Features**, then click **Next** until you reach the Features page.
 - c. Select **Multipath IO**.
 - d. Click **Next**, click **Install**, click **Close**, and then reboot the host server.
4. Identify and document FC HBA WWNs:
 - a. Open a Windows PowerShell console.
 - b. Type **Get-InitiatorPort** and press Enter.
 - c. Locate and record the FC HBA WWNs. The WWNs are needed to map volumes to the hosts.
5. If the hosts are connected to the storage system using FC switches, implement zoning to isolate traffic for each HBA:

 **NOTE:** Skip this step if hosts are directly connected to the storage system.


 - a. Use the FC switch management interface to create a zone for each server HBA. Each zone must contain only one HBA WWN and all the storage port WWNs.
 - b. Repeat for each FC switch.

 **NOTE:** The ME4 Series storage systems support single initiator/multiple target zones.


Register a Windows host with FC HBAs and create volumes

Perform the following steps to register the Windows host with Fibre Channel (FC) HBAs and create volumes using PowerVault Manager:

1. Log in to the PowerVault Manager.
2. Access the Host Setup wizard:
 - From the Welcome screen, click **Host Setup**.
 - From the Home topic, click **Action > Host Setup**.
3. Confirm that you have met the listed prerequisites, then click **Next**.
4. Type a host name in the **Host Name** field.
5. Using the information documented in step 4 of [Attach a Windows host with FC HBAs to the storage system](#) on page 44, select the FC initiators for the host you are configuring, then click **Next**.
6. Group hosts together with other hosts in a cluster.
 - a. For cluster configurations, group hosts together so that all hosts within the group share the same storage.
 - If this host is the first host in the cluster, select **Create a new host group**, type a name for the host group, and click **Next**.
 - If this host is being added to a host group that exists, select **Add to existing host group**, select the group from the drop-down list, and click **Next**.

 **NOTE:** The host must be mapped with the same access, port, and LUN settings to the same volumes or volume groups as every other initiator in the host group.
 - b. For stand-alone hosts, select the **Do not group this host** option, then click **Next**.
7. On the Attach Volumes page, specify the name, size, and pool for each volume, and click **Next**.

To add a volume, click **Add Row**. To remove a volume, click **Remove**.

 **NOTE:** Dell EMC recommends that you update the name with the hostname to better identify the volumes.
8. On the Summary page, review the host configuration settings, and click **Configure Host**.

If the host is successfully configured, a **Success** dialog box is displayed
9. Click **Yes** to return to the Introduction page of the wizard, or click **No** to close the wizard.

Enable MPIO for the volumes on the Windows host

Perform the following steps to enable MPIO for the volumes on the Windows host:

1. Open the Server Manager.
2. Select **Tools > MPIO**.
3. Click the **Discover Multi-Paths** tab.
4. Select **DellEMC ME4** in the **Device Hardware Id** list.
If **DellEMC ME4** is not listed in the **Device Hardware Id** list:
 - a. Ensure that there is more than one connection to a volume for multipathing.
 - b. Ensure that **Dell EMC ME4** is not already listed in the **Devices** list on the **MPIO Devices** tab.
5. Click **Add** and click **Yes** to reboot the Windows server.

Format volumes on a Windows host

Perform the following steps to format a volume on a Windows host:

1. Open Server Manager.
2. Select **Tools > Computer Management**.
3. Right-click on **Disk Management** and select **Rescan Disks**.
4. Right-click on the new disk and select **Online**.
5. Right-click on the new disk again select **Initialize Disk**.
The **Initialize Disk** dialog box opens.
6. Select the partition style for the disk and click **OK**.
7. Right-click on the unallocated space, select the type of volume to create, and follow the steps in the wizard to create the volume.

Configuring a Windows host with iSCSI network adapters

These instructions document IPv4 configuration with dual switch subnet for network redundancy and failover. These instructions do not cover IPv6 configuration:

Prerequisites

- Complete the PowerVault Manager guided setup process and storage setup process.
- Refer to the cabling diagrams within this guide before attaching a host to the storage system; careful planning ensures a successful deployment.
- Complete a planning worksheet with the iSCSI network IP addresses to be used, per the example in the following table:

Table 8. Example worksheet for host server with dual port iSCSI NICs

Management	IP
Server Management	10.10.96.46
ME4024Controller A Management	10.10.96.128
ME4024 Controller B Management	10.10.96.129
Subnet 1	
Server iSCSI NIC 1	172.1.96.46
ME4024 controller A port 0	172.1.100.128
ME4024 controller B port 0	172.1.200.129
ME4024 controller A port 2	172.1.102.128
ME4024 controller B port 2	172.1.202.129
Subnet Mask	255.255.0.0

Table 8. Example worksheet for host server with dual port iSCSI NICs (continued)

Management	IP
Subnet 2	
Server iSCSI NIC 1	172.2.96.46
ME4024 controller A port 1	172.2.101.128
ME4024 controller B port 1	172.2.201.129
ME4024 controller A port 3	172.2.103.128
ME4024 controller B port 3	172.2.203.129
Subnet Mask	255.255.0.0

i **NOTE:** The following instructions document IPv4 configurations with a dual switch subnet for network redundancy and failover. It does not cover IPv6 configuration.

Attach a Windows host with iSCSI network adapters to the storage system

Perform the following steps to attach the Windows host with iSCSI network adapters to the storage system:

1. Ensure that all network adapters have the latest supported firmware and drivers as described on Dell.com/support.
 - i** **NOTE:** The Dell EMC PowerVault ME4 Series storage system supports only software iSCSI adapters.
2. Use the iSCSI cabling diagrams to connect the hosts to the storage system either by using switches or connecting the hosts directly to the storage system.
3. Install MPIO on the iSCSI hosts:
 - a. Open Server Manager.
 - b. Click **Manage > Add Roles and Features**.
 - c. Click **Next** until you reach the Features page.
 - d. Select **Multipath IO**.
 - e. Click **Next**, click **Install**, and click **Close**.
 - f. Reboot the Windows server.

Assign IP addresses for each network adapter connecting to the iSCSI network

Perform the following steps to assign IP addresses for the network adapter that connects to the iSCSI network:

⚠ CAUTION: IP addresses must match the subnets for each network. Make sure that you assign the correct IP addresses to the NICs. Assigning IP addresses to the wrong ports can cause connectivity issues.

i **NOTE:** If using jumbo frames, they must be enabled and configured on all devices in the data path, adapter ports, switches, and storage system.

1. From the Network and Sharing Center, click **Change adapter settings**.
2. Right-click on the network adapter, then select **Properties**.
3. Select **Internet Protocol Version 4**, then click **Properties**.
4. Select the **Use the following IP address** radio button and type the corresponding IP addresses recorded in the planning worksheet described in the Prerequisites section (ex: 172.1.96.46).
5. Set the netmask .
6. Configure a gateway if appropriate.
7. Click **OK** and **Close**. The settings are applied to the selected adapter.
8. Repeat steps 1-7 for each of the required iSCSI interfaces (NIC 1 and NIC 2 in [Example worksheet for host server with dual port iSCSI NICs](#) on page 45).
9. From the command prompt, ping each of the controller IP addresses to verify host connectivity before proceeding. If ping is not successful, verify connections and the appropriate IP/subnet agreement between interfaces.

Configure the iSCSI Initiator on the Windows host

Perform the following steps to configure the iSCSI Initiator on a Windows host:

1. Open the Server Manager.
2. Select **Tools > iSCSI Initiator**. The **iSCSI Initiator Properties** dialog box opens.
If you are running the iSCSI initiator for the first time, click **Yes** when prompted to have it start automatically when the server reboots.
3. Click the **Discovery** tab, then click **Discover Portal**. The **Discover Target Protocol** dialog box opens.
4. Using the planning worksheet that you created in the Prerequisites section, type the IP address of a port on controller A that is on the first subnet and click **OK**.
5. Repeat steps 3-4 to add the IP address of a port on the second subnet that is from controller B .
6. Click the **Targets** tab, select a discovered target, and click **Connect**.
7. Select the **Enable multi-path** check box and click **Advanced**. The **Advanced Settings** dialog box opens.
 - Select **Microsoft iSCSI initiator** from the **Local adapter** drop-down menu..
 - Select the IP address of NIC 1 from the **Initiator IP** drop-down menu.
 - Select the first IP listed in the same subnet from the **Target portal IP** drop-down menu.
 - Click **OK** twice to return to the **iSCSI Initiator Properties** dialog box.
8. Repeat steps 6-7 for the NIC to establish a connection to each port on the subnet.
i **NOTE:** Step 10 is required for multi-path configurations.
9. Repeat steps 3-8 for the NIC 2, connecting it to the targets on the second subnet.
i **NOTE:** After all connections are made, you can click the **Favorite Targets** tab to see each path. If you click **Details**, you can view specific information the selected path.
10. Click the **Configuration tab** and record the initiator name in the **Initiator Name** field. The initiator name is needed to map volumes to the host.
11. Click **OK** to close the **iSCSI Initiator Properties** dialog box.

Register the Windows host with iSCSI network adapters and create volumes

Perform the following steps to register a Windows host with iSCSI network adapters and create volumes using the PowerVault Manager:

1. Log in to the PowerVault Manager.
2. Access the Host Setup wizard:
 - From the Welcome screen, click **Host Setup**.
 - From the Home topic, select **Action > Host Setup**.
3. Confirm that you have met the listed prerequisites, then click **Next**.
4. Type a host name in the **Host Name** field.
5. Using the information from step 10 of the [Configure the iSCSI Initiator](#), select the iSCSI initiators for the host you are configuring, then click **Next**.
6. Group hosts together with other hosts in a cluster.
 - For cluster configurations, group hosts together so that all hosts within the group share the same storage.
 - If this host is the first host in the cluster, select **Create a new host group**, type a name for the host group, and click **Next**.
 - If this host is being added to a host group that exists, select **Add to existing host group**, select the group from the drop-down list, and click **Next**.
i **NOTE:** The host must be mapped with the same access, port, and LUN settings to the same volumes or volume groups as every other initiator in the host group.
 - For stand-alone hosts, select the **Do not group this host** option, and click **Next**.
7. On the Attach Volumes page, specify the name, size, and pool for each volume, and click **Next**.
To add a volume, click **Add Row**. To remove a volume, click **Remove**.
i **NOTE:** Dell EMC recommends that you update the volume name with the hostname to better identify the volumes.
8. On the Summary page, review the host configuration settings, and click **Configure Host**.
If the host is successfully configured, a **Success** dialog box is displayed.
9. Click **Yes** to return to the Introduction page of the wizard, or click **No** to close the wizard.

Enable MPIO for the volumes on the Windows host

Perform the following steps to enable MPIO for the volumes on a Windows host:

1. Open Server Manager.
2. Select **Tools > MPIO**.
3. Click the **Discover Multi-Paths** tab.
4. Select **DellEMC ME4** in the **Device Hardware Id** list.

If **DellEMC ME4** is not listed in the **Device Hardware Id** list:

- a. Ensure that there is more than one connection to a volume for multipathing.
 - b. Ensure that **Dell EMC ME4** is not already listed in the **Devices** list on the **MPIO Devices** tab.
5. Click **Add** and click **Yes** to reboot the Windows server.

Update the iSCSI initiator on the Windows host

Perform the following steps to update the iSCSI initiator on a Windows host:

1. Open Server Manager.
2. Click **Tools > iSCSI initiator**.
3. Click the **Volumes and Devices** tab.
4. Click **Auto Configure**.
5. Click **OK** to close the **iSCSI Initiator Properties** window.

Format volumes on the Windows host

Perform the following steps to format a volume on a Windows host:

1. Open Server Manager.
2. Select **Tools > Computer Management**.
3. Right-click **Disk Management** and select **Rescan Disks**.
4. Right-click on the new disk and select **Online**.
5. Right-click on the new disk again select **Initialize Disk**. The **Initialize Disk** dialog box opens.
6. Select the partition style for the disk and click **OK**.
7. Right-click on the unallocated space, select the type of volume to create, and follow the steps in the wizard to create the volume.

Configuring a Windows host with SAS HBAs

The following sections describe how to configure a Windows host with SAS HBAs:

Prerequisites

- Complete the PowerVault Manager guided system and storage setup process.
- Refer to the cabling diagrams within this guide before attaching a host to the storage system; careful planning ensures a successful deployment.

Attach a Windows host with SAS HBAs to the storage system

Perform the following steps to attach a Windows host with SAS HBAs to the storage system:

1. Ensure that all HBAs have the latest supported firmware and drivers as described on Dell.com/support. For a list of supported SAS HBAs, see the *Dell EMC ME4 Series Storage System Support Matrix* on Dell.com/support.
2. Use the SAS cabling diagrams to cable the hosts directly to the storage system.
3. Install MPIO on the SAS hosts:
 - a. i. Open Server Manager.
 - ii. Click **Manage > Add Roles and Features**.

- iii. Click **Next** until you reach the Features page.
 - iv. Select **Multipath I/O**.
 - v. Click **Next**, click **Install**, and click **Close**.
 - vi. Reboot the Windows server.
4. Identify and document the SAS HBA WWNs:
 - a. Open a Windows PowerShell console.
 - b. Type **Get-InitiatorPort** and press Enter.
 - c. Locate and record the SAS HBA WWNs . The WWNs are needed to map volumes to the server.

Register a Windows host with SAS HBAs and create volumes

Perform the following steps to register a Windows host with SAS HBAs and create volumes using the PowerVault Manager:

1. Log in to the PowerVault Manager.
2. Access the Host Setup wizard:
 - From the Welcome screen, click **Host Setup**.
 - From the Home topic, click **Action > Host Setup**.
3. Confirm that you have met the listed prerequisites, then click **Next**.
4. Type a host name in the **Host Name** field.
5. Using the information documented in step 4 of [Attach SAS hosts to the storage system](#), select the SAS initiators for the host you are configuring, then click **Next**.
6. Group hosts together with other hosts in a cluster.
 - For cluster configurations, group hosts together so that all hosts within the group share the same storage.
 - If this host is the first host in the cluster, select **Create a new host group**, type a name for the host group, and click **Next**.
 - If this host is being added to a host group that exists, select **Add to existing host group**, select the group from the drop-down list, and click **Next**.
 - **NOTE:** The host must be mapped with the same access, port, and LUN settings to the same volumes or volume groups as every other initiator in the host group.
 - For stand-alone hosts, select the **Do not group this host** option, and click **Next**.
7. On the Attach Volumes page, specify the name, size, and pool for each volume, and click **Next**.
To add a volume, click **Add Row**. To remove a volume, click **Remove**.
- **NOTE:** Dell EMC recommends that you update the volume name with the hostname to better identify the volumes.
8. On the Summary page, review the host configuration settings, and click **Configure Host**.
If the host is successfully configured, a **Success** dialog box is displayed.
9. Click **Yes** to return to the Introduction page of the wizard, or click **No** to close the wizard.

Enable MPIO for the volumes on the Windows host

Perform the following steps to enable MPIO for the volumes on the Windows host:

1. Open Server Manager.
2. Select **Tools > MPIO**.
3. Click the **Discover Multi-Paths** tab.
4. Select **DellEMC ME4** in the **Device Hardware Id** list.
If **DellEMC ME4** is not listed in the **Device Hardware Id** list:
 - a. Ensure that there is more than one connection to a volume for multipathing.
 - b. Ensure that **Dell EMC ME4** is not already listed in the **Devices** list on the **MPIO Devices** tab.
5. Click **Add** and click **Yes** to reboot the Windows server.

Format volumes on the Windows host

Perform the following steps to format a volume on a Windows host:

1. Open Server Manager.

2. Select **Tools > Computer Management**.
3. Right-click on **Disk Management** and select **Rescan Disks**.
4. Right-click on the new disk and select **Online**.
5. Right-click on the new disk again select **Initialize Disk**. The **Initialize Disk** dialog box opens.
6. Select the partition style for the disk and click **OK**.
7. Right-click on the unallocated space, select the type of volume to create, and follow the steps in the wizard to create the volume.

Linux hosts

Ensure that the HBAs or network adapters are installed, the drivers are installed, and the latest supported BIOS is installed.

Configuring a Linux host with FC HBAs

The following sections describe how to configure a Linux host with Fibre Channel (FC) HBAs:


Prerequisites

- Complete the PowerVault Manager guided system and storage setup process.
- Refer to the cabling diagrams within this guide before attaching a host to the storage system; careful planning ensures a successful deployment.
- Administrative or privileged user permissions are required to make system-level changes. These steps assume root level access and that all required software packages are already installed (for example, DM Multipath).

Attach a Linux host with FC HBAs to the storage system

Perform the following steps to attach the Linux host with Fibre Channel (FC) HBAs to the storage system:

1. Ensure that all HBAs have the latest supported firmware and drivers as described on the Dell Support portal . For a list of supported standard FC HBAs, see the *Dell EMC PowerVault ME4 Series storage Matrix* on the Dell website. For OEMs, contact your hardware provider.
2. Use the FC cabling diagrams to cable the host servers either by using switches or attaching them directly to the storage system.
3. Identify Fibre Channel WWNs to connect to the storage system by doing the following:
 - a. Open a terminal session.
 - b. Run the `ls -l /sys/class/fc_host` command.
 - c. Run the `more /sys/class/fc_host/host?/port_name` command and replace the ? with the host numbers that are supplied in the data output.
 - d. Record the WWN numeric name.
4. If the hosts are connected to the storage system by FC switches, implement zoning to isolate traffic for each HBA. Skip this step if hosts are directly connected to the storage system.
 - a. Use the FC switch management interface to create a zone for each server HBA. Each zone must contain only one HBA WWN and all the storage port WWNs.
 - b. Repeat for each FC switch.

 **NOTE:** The ME4 Series storage systems support single initiator/multiple target zones.

Register a Linux host with FC HBAs and create and map volumes

Perform the following steps to register the Linux host with Fibre Channel (FC) HBAs , create volumes, and map volumes:

1. Log in to the PowerVault Manager.
2. Access the Host Setup wizard:
 - From the Welcome screen, click **Host Setup**.
 - From the Home topic, click **Action > Host Setup**.

3. Confirm that you have met the listed prerequisites, then click **Next**.
4. Type a hostname in the **Host Name** field.
5. Using the information from step 3 of [Attach a Linux host with FC HBAs to the storage system](#) on page 50 to identify the correct initiators, select the FC initiators for the host you are configuring, then click **Next**.
6. Group hosts together with other hosts.
 - a. For cluster configurations, group hosts together so that all hosts within the group share the same storage.
 - If this host is the first host in the cluster, select **Create a new host group**, then provide a name and click **Next**.
 - If this host is being added to a host group that exists, select **Add to existing host group**. Select the group from the drop-down list, then click **Next**.

i **NOTE:** The host must be mapped with the same access, port, and LUN settings to the same volumes or volume groups as every other initiator in the host group.
 - b. For stand-alone hosts, select the **Do not group this host** option, then click **Next**.
7. On the Attach Volumes page, specify the name, size, and pool for each volume, and click **Next**.
To add a volume, click **Add Row**. To remove a volume, click **Remove**.

i **NOTE:** Dell EMC recommends that you update the name with the hostname to better identify the volumes.
8. On the Summary page, review the host configuration settings, and click **Configure Host**.
If the host is successfully configured, a **Success** dialog box is displayed
9. Click **Yes** to return to the Introduction page of the wizard, or click **No** to close the wizard.

Enable and configure DM Multipath on Linux hosts

Perform the following steps to enable and configure DM multipath on the Linux host:

- i** **NOTE:** Safeguard and block internal server disk drives from multipath configuration files. These steps are meant as a basic setup to enable DM Multipath to the storage system. It is assumed that DM Multipath packages are installed.
1. Run the `multipath -t` command to list the DM Multipath status.
 2. If no configuration exists, use the information that is listed from running the command in step 1 to copy a default template to the directory `/etc`.
 3. If the DM multipath kernel driver is not loaded:
 - a. Run the `systemctl enable multipathd` command to enable the service to run automatically.
 - b. Run the `systemctl start multipathd` command to start the service.
 4. Run the `multipath` command to load storage devices along with the configuration file.
 5. Run the `multipath -l` command to list the Dell EMC PowerVault ME4 Series storage devices as configured under DM Multipath.

Create a Linux file system on the volumes

Perform the following steps to create and mount an XFS file system:

1. From the `multipath -l` command output, identify the device multipath to target when creating a file system.
In this example, the first time that multipath is configured, the first device is `/dev/mapper/mpatha` and it corresponds to `sg` block devices `/dev/sdb` and `/dev/sdd`.

i **NOTE:** Run the `lsscsi` command to list all SCSI devices from the Controller/Target/Bus/LUN map. This command also identifies block devices per controller.
2. Run the `mkfs.xfs /dev/mapper/mpatha` command to create an xfs type file system.
3. Run the `mkdir /mnt/VolA` command to create a mount point for this file system with a referenced name, such as VolA.
4. Run the `mount /dev/mapper/mpatha /mnt/VolA` command to mount the file system.
5. Begin using the file system as any other directory to host applications or file services.
6. Repeat steps 1–5 for each provisioned volume in PowerVault Manager. For example, the device `/dev/mapper/mpathb` corresponds to `sg` block devices `/dev/sdc` and `/dev/sde`.

Configure a Linux host with iSCSI network adapters

The following sections describe how to configure a Linux host with iSCSI network adapters:

- Complete the PowerVault Manager guided system and storage setup process.
- Refer to the cabling diagrams within this guide before attaching a host to the storage system; careful planning ensures a successful deployment.
- Administrative or privileged user permissions are required to make system-level changes. The following sections assume root level access and that all required software packages are already installed, for example, iSCSI-initiator and DM Multipath.
- Complete a planning worksheet with the iSCSI network IP addresses to be used, per the example in the following table.

Table 9. Example worksheet for single host server with dual port iSCSI NICs

Management	IP
Server Management	10.10.96.46
ME4024 Controller A Management	10.10.96.128
ME4024 Controller B Management	10.10.96.129
Subnet 1	
Server iSCSI NIC 1	172.1.96.46
ME4024 controller A port 0	172.1.100.128
ME4024 controller B port 0	172.1.200.129
ME4024 controller A port 2	172.1.102.128
ME4024 controller B port 2	172.1.202.129
Subnet Mask	255.255.0.0
Subnet 2	
Server iSCSI NIC 1	172.2.96.46
ME4024 controller A port 1	172.2.101.128
ME4024 controller B port 1	172.2.201.129
ME4024 controller A port 3	172.2.103.128
ME4024 controller B port 3	172.2.203.129
Subnet Mask	255.255.0.0


The following instructions document IPv4 configurations with a dual switch subnet for network redundancy and failover. It does not cover IPv6 configuration.

Attach a Linux host with iSCSI network adapters to the storage system

1. Ensure that all network adapters have the latest supported firmware and drivers as described on the Dell Support portal.
2. Use the iSCSI cabling diagrams to cable the host servers to the switches or directly to the storage system.

Assign IP addresses for each network adapter connecting to the iSCSI network

 **CAUTION:** The IP addresses must match the subnets for each network, so ensure that you correctly assign IP addresses to the network adapters. Assigning IP addresses to the wrong ports can cause connectivity issues.

 **NOTE:** If using jumbo frames, they must be enabled and configured on all devices in the data path, adapter ports, switches, and storage system.

For RHEL 7

1. From the server terminal or console, run the **nmtui** command to access the NIC configuration tool (NetworkManager TUI).
2. Select **Edit a connection** to display a list of the Ethernet interfaces installed.
3. Select the iSCSI NIC that you want to assign an IP address to.
4. Change the IPv4 Configuration option to **Manual**.
5. Using the planning worksheet that you created in the “Prerequisites” section, provide the subnet mask by entering the NIC IP address using the format x.x.x.x/16. For example: 172.1.96.46/16
6. Configure a gateway, if appropriate.
7. Select **IGNORE** for the IPv6 Configuration.
8. Check **Automatically connect** to start the NIC when the system boots.
9. Select **OK** to exit Edit connection.
10. Select **Back** to return to the main menu.
11. Select **Quit** to exit NetworkManager TUI.
12. Ping the new network interface and associated storage host ports to ensure IP connectivity.
13. Repeat steps 1-12 for each NIC you are assigning IP addresses to.

For SLES 12

1. From the server terminal or console, run the `yast` command to access the YaST Control Center.
2. Select **System > Network Settings**.
3. Select the iSCSI NIC that you want to assign an IP address to, then select **Edit**.
4. Select **Statically Assigned IP Address**.
5. Using the planning worksheet that you created in the “Prerequisites” section, enter the NIC IP address. For example: 172.1.96.46
6. Using the planning worksheet that you created in the “Prerequisites” section, enter the NIC subnet mask. For example: 255.255.0.0.
7. Select Next.
8. Ping the new network interface and associated storage host ports to ensure IP connectivity.
9. Repeat steps 1-8 for each NIC you are assigning IP addresses to (NIC1 and NIC2 in the planning worksheet you created in the “Prerequisites” section).
10. Select **OK** to exit network settings.
11. Select **OK** to exit YaST.

Configure the iSCSI initiators to connect to the storage system



For RHEL 7

1. From the server terminal or console, run the following `iscsiadm` command to discover targets (port A0):
iscsiadm -m discovery -t sendtargets -p <IP>
Where <IP> is the IP address. For example:
iscsiadm -m discovery -t sendtargets -p 172.1.100.128
2. With the discovery output, log in to each portal by running the `iscsiadm` command:
 - a. Run **iscsiadm -m node -T <full IQN> -p <IP>**
Where <full IQN> is the full IQN listing from the output in step 1 and <IP> is the IP address. For example:
iscsiadm -m node -T iqn.1988-11.com.abcc:01.array.bc305bb0b841-p 172.1.100.128
 - b. Repeat the login for each controller host port using the discovery command output in step 1.
 - c. Reboot the host to ensure that all targets are automatically connected.


For SLES 12

1. From the server terminal or console, use the `yast` command to access YaST Control Center.
2. Select Network **Service > iSCSI Initiator**.
3. On the Service tab, select **When Booting**.
4. Select the **Connected Targets** tab.
5. Select **Add**. The iSCSI Initiator Discovery screen displays.
6. Using the [Example worksheet for single host server with dual port iSCSI NICs](#) you created earlier, enter the IP address for port A0 in the IP Address field, then click **Next**. For example: `172.1.100.128`.
7. Select **Connect**.
8. On iSCSI Initiator Discovery screen, select the next adapter and then select **Connect**.
9. When prompted, select **Continue** to bypass the warning message, "Warning target with TargetName is already connected".
10. Select Startup to Automatic, then click **Next**.
11. Repeat steps 2-10 for all remaining adapters.
12. Once the targets are connected, click **Next > Quit** to exit YaST.
13. Reboot the host to ensure that all targets are automatically connected.

Register the Linux host with iSCSI network adapters and create volumes

1. Log in to the PowerVault Manager.
2. Access the Host Setup wizard:
 - From the Welcome screen, click **Host Setup**.
 - From the Home topic, click **Action > Host Setup**.
3. Confirm that you have met the listed prerequisites, then click **Next**.
4. Type a hostname in the **Host Name** field.
5. Using the information from [Configure the iSCSI initiators to connect to the storage system](#) on page 53, select the iSCSI initiators for the host you are configuring, then click **Next**.
6. Group hosts together with other hosts.
 - a. For cluster configurations, group hosts together so that all hosts within the group share the same storage.
 - If this host is the first host in the cluster, select **create a new host group**, then provide a name and click **Next**.
 - If this host is being added to a host group that exists, select **Add to existing host group**. Select the group from the drop-down list, then click **Next**.
 -  **NOTE:** The host must be mapped with the same access, port, and LUN settings to the same volumes or volume groups as every other initiator in the host group.
 - b. For Stand-alone hosts, select the **Do not group this host option**, then click **Next**.
7. On the Attach Volumes page, specify the name, size, and pool for each volume, and click **Next**.
To add a volume, click **Add Row**. To remove a volume, click **Remove**.
 **NOTE:** Dell EMC recommends that you update the name with the hostname to better identify the volumes.
8. On the Summary page, review the host configuration settings, and click **Configure Host**.
If the host is successfully configured, a **Success** dialog box is displayed
9. Click **Yes** to return to the Introduction page of the wizard, or click **No** to close the wizard.

Enable and configure DM Multipath on the Linux host with iSCSI network adapters

 **NOTE:** Safeguard and block internal server disk drives from multipath configuration files. These steps are meant as a basic setup to enable DM Multipath to the storage system. It is assumed that DM Multipath packages are installed.

1. Run the `multipath -t` command to list the DM Multipath status.
2. If no configuration currently exists, use the command information displayed in step 1 to copy a default template to the directory `/etc`.
3. If the DM multipath kernel driver is not loaded:

- a. Run the `systemctl enable multipathd` command to enable the service to run automatically.
 - b. Run the `systemctl start multipathd` command to start the service.
4. Run the **multipath** command to load storage devices in conjunction with the configuration file.
 5. Run the **multipath -l** command to list the Dell EMC PowerVault ME4 Series storage devices as configured under DM Multipath.

Create a Linux file system on the volumes

Perform the following steps to create and mount an XFS file system:

1. From the `multipath -l` command output above, identify the device multipath to target creating a file system.
In this example, the first time multipath is configured, the first device will be `/dev/mapper/mpatha`, which corresponds to `sg` block devices `/dev/sdb` and `/dev/sdd`.
NOTE: Run the `lsscsi` command to list all SCSI devices from the Controller/Target/Bus/LUN map. This also identifies block devices per controller.
2. Run the `mkfs.xfs/dev/mapper/mpatha` command to create an xfs type file system.
3. Run the `mkdir/mnt/VolA` command to create a new mount point for this file system with a referenced name, such as VolA.
4. Run the `mount /dev/mapper/mpatha /mnt/VolA` command to mount the file system.
5. Begin using the file system as any other directory to host applications or file services.
6. Repeat steps 1-5 for other provisioned volumes from the PowerVault Manager. For example, `/dev/mapper/mpathb` corresponds to `sg` block devices `/dev/sdc` and `/dev/sde`.

SAS host server configuration for Linux

The following sections describe how to configure SAS host servers running Linux:

- Complete the PowerVault Manager guided system and storage setup process.
- Refer to the cabling diagrams within this guide before attaching a host to the storage system; careful planning will ensure a successful deployment.
- Administrative or privileged user permissions are required to make system-level changes. These steps assume root level access and that all required software packages are already installed (for example, DM Multipath).

Attach a Linux host with SAS HBAs to the storage system

Perform the following steps to attach the Linux host with SAS HBAs to the storage system:

1. Ensure that all HBAs have the latest supported firmware and drivers as described on the Dell Support web site. For a list of supported SAS HBAs, see the *Dell EMC ME4 Series Storage System Support Matrix* on the Dell Support web site.
2. Use the SAS cabling diagrams to cable the host servers directly to the storage system.
3. Identify SAS HBA initiators to connect to the storage system by doing the following:
 - a. Open a terminal session.
 - b. Run the `dmesg|grep scsi|grep slot` command.
 - c. Record the WWN numeric name.

Register the host and create and map volumes

1. Log in to the PowerVault Manager.
2. Access the Host Setup wizard:
 - From the Welcome screen, click **Host Setup**.
 - From the Home topic, click **Action > Host Setup**.
3. Confirm that you have met the listed prerequisites, then click **Next**.
4. Type a hostname in the **Host Name** field.
5. Using the information from step 3 of [Attach a Linux host with SAS HBAs to the storage system](#) on page 55, select the SAS initiators for the host you are configuring, then click **Next**.
6. Group hosts together with other hosts.

- a. For cluster configurations, group hosts together so that all hosts within the group share the same storage.
 - If this host is the first host in the cluster, select **create a new host group**, then provide a name and click **Next**.
 - If this host is being added to a host group that exists, select **Add to existing host group**. Select the group from the drop-down list, then click **Next**.

i **NOTE:** The host must be mapped with the same access, port, and LUN settings to the same volumes or volume groups as every other initiator in the host group.

- b. For stand-alone hosts, select the **Do not group this host** option, then click **Next**.
7. On the Attach Volumes page, specify the name, size, and pool for each volume, and click **Next**.

To add a volume, click **Add Row**. To remove a volume, click **Remove**.

i **NOTE:** Dell EMC recommends that you update the name with the hostname to better identify the volumes.

8. On the Summary page, review the host configuration settings, and click **Configure Host**.

If the host is successfully configured, a **Success** dialog box is displayed.

9. Click **Yes** to return to the Introduction page of the wizard, or click **No** to close the wizard.

Enable and configure DM Multipathing

i **NOTE:** Safeguard and block internal server disk drives from multipathing configuration files. These steps are meant as a basic setup to enable DM Multipathing to the storage system. It is assumed that DM Multipathing packages are installed.

1. Run the `multipath -t` command to list the DM Multipathing status.
2. If no configuration exists, use the command information that is listed in step 1 to copy a default template to the directory **/etc**.
3. If the DM multipathing kernel driver is not loaded:
 - a. Run the `systemctl enable multipathd` command to enable the service to run automatically.
 - b. Run the `systemctl start multipathd` command to start the service.
4. Run the `multipath` command to load storage devices along with the configuration file.
5. Run the `multipath -l` command to list the ME4 Series storage devices as configured under DM Multipathing.

Create a Linux file system on the volumes

Perform the following steps to create and mount an XFS file system:

1. From the `multipath -l` command output, identify the device multipathing to target creating a file system.

In this example, the first time that multipathing is configured, the first device is `/dev/mapper/mpatha` which corresponds to `sg` block devices `/dev/sdb` and `/dev/sdd`.

i **NOTE:** Run the `lsscsi` command to list all SCSI devices from the Controller/Target/Bus/LUN map. This command also identifies block devices per controller.

2. Run the `mkfs.xfs/dev/mapper/mpatha` command to create an xfs type file system.
3. Run the `mkdir/mnt/VolA` command to create a mount point for this file system with a referenced name, such as VolA.
4. Run the `mount /dev/mapper/mpatha /mnt/VolA` command to mount the file system.
5. Begin using the file system as any other directory to host applications or file services.
6. Repeat steps 1-5 for other provisioned volumes from the PowerVault Manager. For example, `/dev/mapper/mpathb` corresponds to `sg` block devices `/dev/sdc` and `/dev/sde`.

VMware ESXi hosts

Ensure that the HBAs or network adapters are installed and the latest supported BIOS is installed.

Fibre Channel host server configuration for VMware ESXi

The following sections describe how to configure Fibre Channel host servers running VMware ESXi:

Prerequisites

- Complete the PowerVault Manager guided system and storage setup process.
- Refer to the cabling diagrams within this guide before attaching a host to the storage system; careful planning ensures a successful deployment.
- Install the required version of the VMware ESXi operating system and configure it on the host.

Attach an ESXi host with FC HBAs to the storage system

Perform the following steps to attach the ESXi host with Fibre Channel (FC) HBAs to the storage system:

1. Ensure that all HBAs have the latest supported firmware and drivers as described on the Dell Support portal . For a list of supported standard FC HBAs, see the *Dell EMC ME4 Storage Matrix* on the Dell website. For OEMs, contact your hardware provider.
2. Use the FC cabling diagrams to cable the host servers either by using switches or attaching them directly to the storage system.
3. Log in to the VMware vCenter Server and add the newly configured ESXi host to the appropriate datacenter.
4. On the Configure tab, select **Storage > Storage Adapters**.
5. Verify that the required FC storage adapters are listed, then record the HBA's WWN as listed under Properties.
6. If the hosts are connected to the storage system by FC switches, implement zoning to isolate traffic for each HBA by doing the following (skip this step if hosts are directly connected to the storage system):
 - a. Use the FC switch management interface to create a zone for each server HBA. Each zone must contain only one HBA WWN and all the storage port WWNs.
 - b. Repeat sub-step a for each FC switch.

 **NOTE:** The Dell EMC PowerVault ME4 Series storage systems support single initiator/multiple target zones.

Register an ESXi host with FC HBAs and create and map volumes

Perform the following steps to register the ESXi host with Fibre Channel (FC) HBAs, create volumes, and map volumes storage system:

1. Log in to the PowerVault Manager.
2. Access the Host Setup wizard:
 - From the Welcome screen, click **Host Setup**.
 - From the Home topic, click **Action > Host Setup**.
3. Confirm that you have met the listed prerequisites, then click **Next**.
4. Type a hostname in the **Host Name** field.
5. Using the information from step 5 of [Attach an ESXi host with FC HBAs to the storage system](#) on page 57, select the FC initiators for the host you are configuring, then click **Next**.
6. Group hosts together with other hosts.
 - a. For cluster configurations, group hosts together so that all hosts within the group share the same storage.
 - If this host is the first host in the cluster, select **Create a new host group**, then provide a name and click **Next**.
 - If this host is being added to a host group that exists, select **Add to existing host group**. Select the group from the dropdown list, then click **Next**.

i **NOTE:** The host must be mapped with the same access, port, and LUN settings to the same volumes or volume groups as every other initiator in the host group.

- b. For stand-alone hosts, select the **Do not group this host** option, then click **Next**.
7. On the Attach Volumes page, specify the name, size, and pool for each volume, and click **Next**.
To add a volume, click **Add Row**. To remove a volume, click **Remove**.

i **NOTE:** Dell EMC recommends that you update the name with the hostname to better identify the volumes.

8. On the Summary page, review the host configuration settings, and click **Configure Host**.
If the host is successfully configured, a **Success** dialog box is displayed
9. Click **Yes** to return to the Introduction page of the wizard, or click **No** to close the wizard.

Enable multipathing on an ESXi host with FC volumes

Perform the following steps to enable multipathing on the ESXi host with Fibre Channel (FC) volumes:

1. Log in to the VMware vCenter Server, then click on the ESXi host added.
2. On the Configure tab, select **Storage Devices**.
3. Perform a rescan of the storage devices.
4. Select the FC Disk (Dell EMC Fibre Channel Disk) created in the [Register an ESXi host with FC HBAs and create and map volumes](#) on page 57 procedure, then select the **Properties** tab below the screen.
5. Click **Edit Multipathing**, then select **Round Robin (VMware)** from the drop down list.
6. Click **OK**.
7. Follow steps 4-6 for all volumes presented from the Dell EMC PowerVault ME4 Series Storage system to ESXi host.

Volume rescan and datastore creation for an FC host server

Perform the following steps to rescan storage and create a datastore:

1. Log in to the VMware vCenter Server, then click the configured ESXi host in step 5 of [Attach an ESXi host with FC HBAs to the storage system](#) on page 57.
2. On the Configure tab, select **Storage Adapters**.
3. Select the FC software adapter, and click **Rescan Storage**.
The **Rescan Storage** dialog box opens.
4. Click **OK**.
After a successful rescan, the volumes that are displayed in the [Register an ESXi host with FC HBAs and create and map volumes](#) on page 57 section are visible.
5. Create a VMware datastore file system on the ME4 Series volume.
 - a. On the Configure tab, select **Datastore > Add Storage**.
 - b. Select **VMFS** as the type on New Datastore screen, then click **Next**.
 - c. Enter a name for the datastore, select **right volume/Lun**, and click **Next**.
 - d. Select **VMFS6** as the VMFS version of datastore, then click **OK**.
 - e. On the Partition configuration page, select the default value that is shown, then click **Next**.
 - f. Click **Finish** to complete the creation of the new datastore.

iSCSI host server configuration for VMware ESXi

The following sections describe how to configure iSCSI host servers running VMware ESXi:

Prerequisites

- Complete the PowerVault Manager guided system and storage setup process.
- Refer to the cabling diagrams within this guide before attaching a host to the storage system; careful planning ensures a successful deployment.

- Install the required version of the VMware ESXi operating system and configure it on the host.
- Complete a planning worksheet with the iSCSI network IP addresses to be used, per the example in the following table.

Table 10. Example worksheet for single host server with dual port iSCSI NICs

Management	IP
Server Management	10.10.96.46
ME4024 Controller A Management	10.10.96.128
ME4024 Controller B Management	10.10.96.129
Subnet 1	
Server iSCSI NIC 1	172.1.96.46
ME4024 controller A port 0	172.1.100.128
ME4024 controller B port 0	172.1.200.129
ME4024 controller A port 2	172.1.102.128
ME4024 controller B port 2	172.1.202.129
Subnet Mask	255.255.0.0
Subnet 2	
Server iSCSI NIC 1	172.2.96.46
ME4024 controller A port 1	172.2.101.128
ME4024 controller B port 1	172.2.201.129
ME4024 controller A port 3	172.2.103.128
ME4024 controller B port 3	172.2.203.129
Subnet Mask	255.255.0.0

Attach an ESXi host with network adapters to the storage system

Perform the following steps to attach the ESXi host with network adapters to the storage system:

1. Ensure that all network adapters have the latest supported firmware and drivers as described on the Dell Support portal.
 - NOTE:** The Dell EMC PowerVault ME4 Series storage system supports only software iSCSI adapters.
2. Use the iSCSI cabling diagrams to cable the host servers either by using switches or attaching them directly to the storage system using one-to-one mode. Record the two different IP address ranges for each storage system controller. For example: 172.2.15.x, 172.3.20.x .
3. If the host servers are connected to the storage system by iSCSI switches, configure the switches to use two different IP address ranges/subnets. Configuring the switches with two different IP address ranges/subnets enables high availability.

Configure the VMware ESXi VMkernel

Perform the following steps to configure the VMware ESXi VMkernel:

1. From the VMWare VSphere web client, click **Configure > Networking > Physical adapters**.
2. Locate and document the device name for the NICs used for iSCSI traffic.
3. Click the VMkernel adapters, then click the plus (+) icon to create a VMkernel adapter.
4. On the Select Connection Type page, select **VMkernel Network Adapter > Next**.
5. On the Select Target Device page, select **New standard switch > Next**.

6. On the Create Standard Switch page, click the plus (+) icon, then select **vmnic** > **OK** to connect to the subnet defined in step 4 of the “Attach hosts to the storage system” procedure.
7. Click **Next**.
8. Provide a network label, then update the port properties.
9. On the IPv4 settings page, select **Static IP** and assign an IP using your planning worksheet.
10. Click **Next**.
11. On the Ready to complete page, review the settings and then click **Finish**.
12. Repeat steps 1–11 for each NIC to use for iSCSI traffic.

NOTE: If you are using jumbo frames, they must be enabled and configured on all devices in the data path, adapter ports, switches, and storage system.

Configure the software iSCSI adapter on the ESXi host

Perform the following steps to configure a software iSCSI adapter on the ESXi host:

NOTE: If you plan to use VMware ESXi with 10GBase-T controllers, you must perform *one* of the following tasks:

- Update the controller firmware to the latest version posted on Dell.com/support before connecting the ESXi host to the ME4 Series storage system.

OR

- Run the following ESX CLI command on every ESXi host before connecting it to the ME4 Series storage system:

```
esxcli system settings advanced set --int-value 0 -option /VMFS3 /
HardwareAcceleratedLocking
```

1. Log in to the VMware vCenter Server.
2. On the **Configure** tab, select **Storage** > **Storage Adapters**.
3. Click the plus (+) icon, then select **software iSCSI adapter** > **OK**. The adapter is added to the list of available storage adapters.
4. Select the newly added iSCSI adapter, then click **Targets** > **Add**.
5. Enter the iSCSI IP address that is assigned to the iSCSI host port of storage controller A, then click **OK**.
6. Repeat steps 4–5 for the iSCSI host port of storage controller B.
7. If multiple VMkernels are used on the same subnet, configure the network port binding:
 - a. On the software iSCSI adapter, click the **Network Port Binding** tab, then click the plus (+) icon to add the virtual network port to bind with the iSCSI adapter.

NOTE: This step is required to establish a link between the iSCSI Adapter and the Vmkernel adapters that are created in the *Configure the VMware ESXi Vmkernel* procedure.



If each of the VMkernels used for iSCSI are on separate subnets, skip this step.
 - b. Select the VMKernel adapters that are created in the *Configure the VMware ESXi Vmkernel* procedure, then click **OK**.
 - c. Select **Rescan of storage adapters**.

Register an ESXi host with a configured software iSCSI adapter and create and map volumes

Perform the following steps to register the ESXi host with a software iSCSI adapter, then create volumes, and map volumes:

1. Log in to the PowerVault Manager.
2. Access the Host Setup wizard:
 - From the Welcome screen, click **Host Setup**.
 - From the Home topic, click **Action** > **Host Setup**.
3. Confirm that you have met the listed prerequisites, then click **Next**.
4. Type a hostname in the **Hostname** field.
5. Using the information from [Configure the software iSCSI adapter on the ESXi host](#) on page 60, select the iSCSI initiators for the host you are configuring, then **Next**.
6. Group hosts together with other hosts.

- a. For cluster configurations, group hosts together so that all hosts within the group share the same storage.
 - If this host is the first host in the cluster, select **Create a new host group**, then provide a name and click **Next**.
 - If this host is to be part of a host group that exists, select **Add to existing host group**. Select the group from the dropdown list, then click **Next**.

 **NOTE:** The host must be mapped with the same access, port, and LUN settings to the same volumes or volume groups as every other initiator in the host group.
- b. For stand-alone hosts, select the **Do not group this host** option, then click **Next**.
7. On the Attach Volumes page, specify the name, size, and pool for each volume, and click **Next**.
To add a volume, click **Add Row**. To remove a volume, click **Remove**.
-  **NOTE:** Dell EMC recommends that you update the name with the hostname to better identify the volumes.
8. On the Summary page, review the host configuration settings, and click **Configure Host**.
If the host is successfully configured, a **Success** dialog box is displayed
9. Click **Yes** to return to the Introduction page of the wizard, or click **No** to close the wizard.

Enable multipathing on an ESXi host with iSCSI volumes

Perform the following steps to enable multipathing on the ESXi hosts with iSCSI volumes:

1. Log in to the VMware vCenter Server, then click the ESXi host added.
2. On the Configure tab, select **Storage Devices**.
3. Perform a rescan of the storage devices.
4. Select the iSCSI disk (Dell EMC iSCSI disk) created in the [Register an ESXi host with a configured software iSCSI adapter and create and map volumes](#) on page 60 procedure, then select the **Properties** tab below the screen.
5. Click **Edit Multipathing**, then select **Round Robin (VMware)** from the drop-down list.
6. Click **OK**.
7. Repeat steps 4–6 for all the volumes that are presented from the Dell EMC PowerVault ME4 Series storage system to the ESXi host.

Volume rescan and datastore creation for an ESXi hosts with iSCSI network adapters

Perform the following steps to rescan storage and create datastores on the ESXi host:

1. Log in to the VMware vCenter Server, then click the ESXi host that was configured in step 5 of [Attach an ESXi host with network adapters to the storage system](#) on page 59.
2. On the Configure tab, select **Storage Adapters**.
3. Select the iSCSI software adapter, and click **Rescan Storage**.
The **Rescan Storage** dialog box opens.
4. Click **OK**.
After a successful rescan, the volumes that are displayed in the [Register an ESXi host with a configured software iSCSI adapter and create and map volumes](#) on page 60 section are visible.
5. Create a VMware datastore file system on the ME4 Series volume.
 - a. On the Configure tab, select **Datastore > Add Storage**.
 - b. Select **VMFS** as the type on New Datastore screen, then click **Next**.
 - c. Enter a name for the datastore, select **right volume/Lun**, and click **Next**.
 - d. Select **VMFS6** as the VMFS version of datastore, then click **OK**.
 - e. On the Partition configuration page, select the default value that is shown, then click **Next**.
 - f. Click **Finish** to complete the creation of the new datastore.

SAS host server configuration for VMware ESXi

The following sections describe how to configure SAS host servers running VMware ESXi:


Prerequisites

- Complete the PowerVault Manager guided system and storage setup process.
- Refer to the cabling diagrams within this guide before attaching a host to the storage system; careful planning ensures a successful deployment.
- Install the required version of the ESXi operating system and configure it on the host.

Attach an ESXi host with SAS HBAs to the storage system

Perform the following steps to attach the ESXi host with SAS HBAs to the storage system:


1. Ensure that all HBAs have the latest supported firmware and drivers as described on the Dell Support portal. For a list of supported standard SAS HBAs, see the *Dell EMC ME4 Support Matrix* on the Dell website. For OEMs, contact your hardware provider.
2. Use the SAS cabling diagrams to cable the host servers either by using switches or attaching them directly to the storage system.
3. Log in to the VMware vCenter Server and add the newly configured ESXi host to the Datacenter.
4. On the Configure tab, select **Storage > Storage Adapters**.
5. Verify that the required SAS storage adapters are listed, then record the HBA WWNs as listed under Properties.

 **NOTE:** SAS HBAs have two ports. The World Wide Port Name (WWPN) for port 0 ends in zero and the WWPN for port 1 ends in one.


Register a Linux host with SAS HBAs and create and map volumes

Perform the following steps to register the Linux host with SAS HBAs, create volumes, and map volumes:

1. Log in to the PowerVault Manager.
2. Access the Host Setup wizard by doing one of the following:
 - From the Welcome screen, click **Host Setup**.
 - From the Home topic, click **Action > Host Setup**.
3. Confirm that you have met the listed prerequisites, then click **Next**.
4. Type a hostname in the **Host Name** field.
5. Using the information from step 5 of the [Attach a Linux host with SAS HBAs to the storage system](#) on page 55, select the SAS initiators for the host you are configuring, then click **Next**.
6. Group hosts together with other hosts.
 - a. For cluster configurations, group hosts together so that all hosts within the group share the same storage.
 - If this host is the first host in the cluster, select **Create a new host group**, then provide a name and click **Next**.
 - If this host is being added to a host group that exists, select **Add to existing host group**. Select the group from the drop-down list, then click **Next**.

 **NOTE:** The host must be mapped with the same access, port, and LUN settings to the same volumes or volume groups as every other initiator in the host group.

- b. For stand-alone hosts, select the **Do not group this host** option, then click **Next**.
7. On the Attach Volumes page, specify the name, size, and pool for each volume, and click **Next**.
To add a volume, click **Add Row**. To remove a volume, click **Remove**.

 **NOTE:** Dell EMC recommends that you update the name with the hostname to better identify the volumes.

8. On the Summary page, review the host configuration settings, and click **Configure Host**.
If the host is successfully configured, a **Success** dialog box is displayed

9. Click **Yes** to return to the Introduction page of the wizard, or click **No** to close the wizard.

Enable multipathing on an ESXi host with SAS volumes

Perform the following steps to enable multipathing on the ESXi host with SAS volumes:

1. Log in to the VMware vCenter Server, then click the ESXi host.
2. On the Configure tab, select **Storage > Storage Adapters**.
3. Select the SAS HBA and click **Rescan Storage**.
The **Rescan Storage** dialog box opens.
4. Click **OK**.
5. Select the Dell EMC disk that was added to the ESXi host in [Register a Linux host with SAS HBAs and create and map volumes](#) on page 62.
6. Click the **Properties** tab that is located below the selected disk.
7. Click **Edit Multipathing**.
The **Edit Multipathing Policies** dialog box opens.
8. Select a multipathing policy for the volume from the **Path selection policy** drop-down list and click **OK**.
i **NOTE:** The VMware multipathing policy defaults to **Most Recently Used (VMware)**. Use the default policy for a host with one SAS HBA that has a single path to both controllers. If the host has two SAS HBAs (for example, the host has two paths to each controller), Dell EMC recommends that you change the multipathing policy to **Round Robin (VMware)**.
9. Repeat steps 5–8 for each SAS volume that is attached to the ESXi host.

Volume rescan and datastore creation for a SAS host server

Perform the following steps to rescan storage and create datastores:

1. Log in to the VMware vCenter Server, then click the ESXi host that was configured in step 5 of [Attach an ESXi host with SAS HBAs to the storage system](#) on page 62.
2. On the Configure tab, select **Storage Adapters**.
3. Select the SAS software adapter, and click **Rescan Storage**.
The **Rescan Storage** dialog box opens.
4. Click **OK**.
After a successful rescan, the volumes that are displayed in the [Register a Linux host with SAS HBAs and create and map volumes](#) on page 62 section are visible.
5. Create a VMware datastore file system on the ME4 Series volume.
 - a. On the Configure tab, select **Datastore > Add Storage**.
 - b. Select **VMFS** as the type on New Datastore screen, then click **Next**.
 - c. Enter a name for the datastore, select **right volume/Lun**, and click **Next**.
 - d. Select **VMFS6** as the VMFS version of datastore, then click **OK**.
 - e. On the Partition configuration page, select the default value that is shown, then click **Next**.
 - f. Click **Finish** to complete the creation of the new datastore.

Citrix XenServer hosts

Ensure that the HBAs or network adapters are installed and the latest supported BIOS is installed.

Fibre Channel host server configuration for Citrix XenServer

The following sections describe how to configure Fibre Channel host servers running Citrix XenServer:

Prerequisites


- Complete the PowerVault Manager guided system and storage setup process.
- See the cabling diagrams within this guide before attaching a host to the storage system; careful planning ensures a successful deployment.
- Install and configure the required version of the XenServer operating system on the hosts.
- Install XenCenter on a Windows computer, and connect it to the XenServer hosts.
- Configure the XenServer hosts into a pool.


Attach a XenServer host with FC HBAs to the storage system

Perform the following steps to attach a XenServer host with Fibre Channel (FC) HBAs to the storage system:

1. Ensure that all HBAs have the latest supported firmware and drivers as described on Dell.com/support. For a list of supported FC HBAs, see the *Dell EMC ME4 Series Storage System Support Matrix*.
2. Use the FC cabling diagrams to cable the hosts to the storage system either by using switches or connecting the hosts directly to the storage system.
3. Log in to the console for each XenServer host using SSH or XenCenter.
4. Use the following command to display and record the WWNs for the HBA ports that are connected to the storage system:

```
sysstool -c fc_host -v | grep port_name
```

5. If the hosts are connected to the storage system using FC switches, implement zoning to isolate traffic for each HBA.
 -  **NOTE:** Skip this step if hosts are directly connected to the storage system.
 - a. Use the FC switch management interface to create a zone for each server HBA. Each zone must contain only one HBA WWN and all the storage port WWNs.
 - b. Repeat the previous step for each FC switch.

 **NOTE:** The Dell EMC PowerVault ME4 Series storage systems support single initiator/multiple target zones.

Enable Multipathing on a XenServer host

Perform the following steps to enable Multipathing on a XenServer host using XenCenter:

1. Log in to XenCenter and select the XenServer host.
2. Right-click the host, and select **Enter Maintenance Mode**.
3. On the General tab, click **Properties**.
The **Properties** window is displayed.
4. Click the **Multipathing** tab, and select the **Enable multipathing on this server** checkbox.
5. Click **OK**.
6. Right-click the host, and select **Exit Maintenance Mode**.

Repeat the previous steps for all the hosts in the pool.

Register a XenServer host with FC HBAs and create volumes

Perform the following steps to register a XenServer host with Fibre Channel (FC) HBAs, and create volumes using the PowerVault Manager:

1. Log in to the PowerVault Manager.
2. Access the Host Setup wizard:
 - From the Welcome screen, click **Host Setup**.
 - From the Home topic, click **Action > Host Setup**.
3. Confirm that all the Fibre Channel prerequisites have been met, then click **Next**.
4. Type the hostname in the **Host Name** field.
5. Using the information from step 4 of [Attach a XenServer host with FC HBAs to the storage system](#) on page 64, select the Fibre Channel initiators for the host you are configuring, then click **Next**.
6. Group hosts together with other hosts in a cluster.
 - a. For cluster configurations, group hosts together so that all hosts within the group share the same storage.
 - If this host is the first host in the cluster, select **Create a new host group**, type a name for the host group, and click **Next**.
 - If this host is being added to a host group that exists, select **Add to existing host group**, select the group from the drop-down list, and click **Next**.

NOTE: The host must be mapped with the same access, port, and LUN settings to the same volumes or volume groups as every other initiator in the host group.
 - b. For stand-alone hosts, select the **Do not group this host** option, then click **Next**.
7. On the Attach Volumes page, specify the name, size, and pool for each volume, and click **Next**.
To add a volume, click **Add Row**. To remove a volume, click **Remove**.

NOTE: Dell EMC recommends that you update the name with the hostname to better identify the volumes.
8. On the Summary page, review the host configuration settings, and click **Configure Host**.
If the host is successfully configured, a **Success** dialog box is displayed
9. Click **Yes** to return to the Introduction page of the wizard, or click **No** to close the wizard.

Create a Storage Repository for a volume on a XenServer host with FC HBAs

Perform the following steps to create a Storage Repository (SR) for a volume on a XenServer host with Fibre Channel (FC) HBAs

1. Log in to XenCenter and select the XenServer host.
2. Select the pool in the Resources pane.
3. Click **New Storage**.
The **New Storage Repository** wizard opens.
4. Select **Hardware HBA** as the storage type and click **Next**.
5. Type a name for the new SR in the **Name** field.
6. Click **Next**.
The wizard scans for available LUNs and then displays a page listing all the LUNs found.
7. Select the LUNs from the list of discovered LUNs to use for the new SR.

NOTE: The storage target must be configured to enable every XenServer host in the pool to have access to one or more LUNs.
8. Click **Create**.
The **New Storage Repository** dialog box opens.

NOTE: A warning message is displayed if there are existing SRs on the LUN that you have selected. Review the details, and perform one of the following actions:

 - Click **Reattach** to use the existing SR.
 - Click **Format** to delete the existing SR and to create an SR.
 - If you prefer to select a different LUN, click **Cancel** and select a different LUN from the list.
9. Click **Finish**.

The new SR is displayed in the Resources pane, at the pool level.

iSCSI host server configuration for Citrix XenServer

The following sections describe how to configure iSCSI host servers running Citrix XenServer:

Prerequisites

- Complete the PowerVault Manager guided setup process and storage setup process.
- See the cabling diagrams within this guide before attaching a host to the storage system; careful planning ensures a successful deployment.
- Install and configure the required version of the XenServer operating system on the hosts.
- Install XenCenter on a Windows computer, and connect it to the XenServer hosts.
- Configure the XenServer hosts into a pool.
- Complete a planning worksheet with the iSCSI network IP addresses to be used, per the example in the following table:


Table 11. Example worksheet for single host server with dual port iSCSI NICs

Management	IP
Server Management	10.10.96.46
ME4024 Controller A Management	10.10.96.128
ME4024 Controller B Management	10.10.96.129
Subnet 1	
Server iSCSI network adapter 1	172.1.96.46
ME4024 controller A port 0	172.1.100.128
ME4024 controller B port 0	172.1.200.129
ME4024 controller A port 2	172.1.102.128
ME4024 controller B port 2	172.1.202.129
Subnet Mask	255.255.0.0
Subnet 2	
Server iSCSI network adapter 1	172.2.96.46
ME4024 controller A port 1	172.2.101.128
ME4024 controller B port 1	172.2.201.129
ME4024 controller A port 3	172.2.103.128
ME4024 controller B port 3	172.2.203.129
Subnet Mask	255.255.0.0


Attach a XenServer host with network adapters to the storage system

Perform the following steps to attach a XenServer host with network adapters to the storage system:

1. Ensure that all network adapters have the latest supported firmware and drivers as described on Dell.com/support.

 **NOTE:** The Dell EMC PowerVault ME4 Series storage system supports only software iSCSI adapters.


2. Use the iSCSI cabling diagrams to cable the host servers either by using switches or attaching them directly to the storage system using one-to-one mode. Record the two different IP address ranges for each storage system controller. For example: 172.2.15.x, 172.3.20.x.
3. If the host servers are connected to the storage system by iSCSI switches, configure the switches to use two different IP address ranges/subnets.

 **NOTE:** Configuring the switches with two different IP address ranges/subnets enables high availability.

Configure a software iSCSI adapter on a XenServer host

Perform the following steps to configure a software iSCSI adapter on a XenServer host:

1. Log in to XenCenter and select the XenServer host.
2. Select the pool in the **Resources** pane, and click the **Networking** tab.
3. Identify and document the network name that is used for iSCSI traffic.
4. Click **Configure**.
The **Configure IP Address** dialog box is displayed.
5. Select **Add IP address** in the left pane.
 - a. Type a name for the interface in the **Name** field.
 - b. Select the network identified in step 3 from the **Network** drop-down menu.
 - c. Assign IP addresses to the interface using your planning worksheet.
 - d. Click **OK**.
6. Repeat the previous steps for each network to use for iSCSI traffic.

 **NOTE:** If you are using jumbo frames, they must be enabled and configured on all devices in the data path, adapter ports, switches, and storage system.

Configure the iSCSI IQN on a XenServer host

Perform the following steps to configure the iSCSI IQN on a XenServer host:

1. Log in to XenCenter and select the XenServer host.
2. Select the pool in the **Resources** pane, and click the **General** tab.
3. Click **Properties**.
The **Properties** dialog box is displayed.
4. Type a new value in the **iSCSI IQN** field.
5. Click **OK**.
6. Repeat the previous steps for all the hosts in the pool.

Enable Multipathing on a XenServer host

Perform the following steps to enable Multipathing on a XenServer host using XenCenter:

1. Log in to XenCenter and select the XenServer host.
2. Right-click the host, and select **Enter Maintenance Mode**.
3. On the General tab, click **Properties**.
The **Properties** window is displayed.
4. Click the **Multipathing** tab, and select the **Enable multipathing on this server** checkbox.
5. Click **OK**.
6. Right-click the host, and select **Exit Maintenance Mode**.

Repeat the previous steps for all the hosts in the pool.

Register a XenServer host with a software iSCSI adapter and create volumes

Perform the following steps to register a XenServer host with a software iSCSI adapter, and create volumes using the PowerVault Manager:


1. Log in to the PowerVault Manager.
2. Create initiators for the XenServer hosts.
 - a. From the Hosts topic, select **Action > Create Initiator**


- b. Type the iSCSI IQN that was specified for the XenServer host in [Configure the iSCSI IQN on a XenServer host](#) on page 67.
- c. Type a name for the initiator in the **Initiator Name** field.
3. Select the initiator.
4. Select **Action > Add to Host**.
The **Add to Host** dialog box is displayed.
5. Type a hostname or select a host from the **Host Select** field and click **OK**.
6. Repeat the previous steps for all the XenServer hosts iSCSI IQNs.
7. Group hosts together with other hosts in a cluster.
 - a. Select the host to add to the host group.
 - b. Select **Action > Add to Host Group**.
The **Add to Host Group** dialog box is displayed.
 - c. Type a host group name or select a host group from the **Host Group Select** field and click **OK**.
8. Map the volumes to the host group.
 - a. Click the Volumes topic, select the volume to map.
If a volume does not exist, create a volume.
 - b. Select **Action > Map Volumes**.
The **Map** dialog box is displayed.
 - c. Select the host group from the **Available Host Groups, Host, and Initiators** area.
 - d. If not already selected, select the volumes from the **Available Volume Groups and Volumes** area.
 - e. Click **Map**.
 - f. Click **OK**.


Create a Storage Repository for a volume on a XenServer host with a software iSCSI adapter

Perform the following steps to create a Storage Repository (SR) for a volume on a XenServer host with a software iSCSI adapter:

1. Log in to XenCenter and select the XenServer host.
2. Select the pool in the Resources pane.
3. Click **New Storage**.
The **New Storage Repository** wizard opens.
4. Select **Software iSCSI** as the storage type and click **Next**.
5. Type a name for the new SR in the **Name** field.
6. Type the IP address or hostname of the iSCSI target in the **Target Host** field.

 **NOTE:** The iSCSI storage target must be configured to enable every XenServer host in the pool to have access to one or more LUNs.
7. If you have configured the iSCSI target to use CHAP authentication:
 - a. Select the **Use CHAP** checkbox.
 - b. Type a CHAP username in the **User** field.
 - c. Type the password for the CHAP username in the **Password** field.
8. Click **Discover IQNs** and select the iSCSI target IQN from the **Target IQN** drop-down menu.

 **CAUTION:** The iSCSI target and all servers in the pool must have unique IQNs.
9. Click **Discover LUNs** and select the LUN on which to create the SR from the **Target LUN** drop-down menu.

 **CAUTION:** Each individual iSCSI storage repository must be contained entirely on a single LUN, and may not span more than one LUN. Any data present on the chosen LUN is destroyed.
10. Click **Finish**.
11. Click **Yes** to format the disk.
The new SR is displayed in the Resources pane, at the pool level.

SAS host server configuration for Citrix XenServer

The following sections describe how to configure SAS host servers running Citrix XenServer:

Prerequisites

- Complete the PowerVault Manager guided system and storage setup process.
- See the cabling diagrams within this guide before attaching a host to the storage system; careful planning ensures a successful deployment.
- Install and configure the required version of the XenServer operating system on the hosts.
- Install XenCenter on a Windows computer and connect it to the XenServer hosts.
- Configure the XenServer hosts into a pool.

Attach a XenServer host with SAS HBAs to the storage system

Perform the following steps to attach a XenServer host with SAS HBAs to the storage system:

1. Ensure that all HBAs have the latest supported firmware and drivers as described on Dell.com/support. For a list of supported SAS HBAs, see the *Dell EMC ME4 Series Storage System Support Matrix*.
2. Use the SAS cabling diagrams to cable the hosts to the storage system either by using switches or connecting the hosts directly to the storage system.
3. Log in to the console for each XenServer host using SSH or XenCenter.
4. Use the following command to display and record the initiator ID for the HBA ports that are connected to the storage enclosure:

```
sysstool -c sas_device -v | grep enclosure_identifier
```

NOTE: SAS HBAs have two ports. The World Wide Port Name (WWPN) for port 0 ends in 0 and the WWPN for port 1 ends in 1.

Enable Multipathing on a XenServer host

Perform the following steps to enable Multipathing on a XenServer host using XenCenter:

1. Log in to XenCenter and select the XenServer host.
2. Right-click the host, and select **Enter Maintenance Mode**.
3. On the General tab, click **Properties**.
The **Properties** window is displayed.
4. Click the **Multipathing** tab, and select the **Enable multipathing on this server** checkbox.
5. Click **OK**.
6. Right-click the host, and select **Exit Maintenance Mode**.

Repeat the previous steps for all the hosts in the pool.

Register a XenServer host with SAS HBAs and create volumes

Perform the following steps to register a XenServer host with SAS HBAs, and create volumes using the PowerVault Manager:

1. Log in to the PowerVault Manager.
2. Access the Host Setup wizard:
 - From the Welcome screen, click **Host Setup**.
 - From the Home topic, click **Action > Host Setup**.
3. Confirm that all the SAS prerequisites have been met, then click **Next**.
4. Type a host name in the **Host Name** field.
5. Using the information from step 4 of [Attach a XenServer host with SAS HBAs to the storage system](#) on page 69, select the SAS initiators for the host you are configuring, then click **Next**.

6. Group hosts together with other hosts in a cluster.
 - a. For cluster configurations, group hosts together so that all hosts within the group share the same storage.
 - If this host is the first host in the cluster, select **Create a new host group**, type a name for the host group, and click **Next**.
 - If this host is being added to a host group that exists, select **Add to existing host group**, select the group from the drop-down list, and click **Next**.

i **NOTE:** The host must be mapped with the same access, port, and LUN settings to the same volumes or volume groups as every other initiator in the host group.
 - b. For stand-alone hosts, select the **Do not group this host** option, then click **Next**.
7. On the Attach Volumes page, specify the name, size, and pool for each volume, and click **Next**.
To add a volume, click **Add Row**. To remove a volume, click **Remove**.

i **NOTE:** Dell EMC recommends that you update the name with the hostname to better identify the volumes.
8. On the Summary page, review the host configuration settings, and click **Configure Host**.
If the host is successfully configured, a **Success** dialog box is displayed
9. Click **Yes** to return to the Introduction page of the wizard, or click **No** to close the wizard.

Create a Storage Repository for a volume on a XenServer host with SAS HBAs

Perform the following steps to create a Storage Repository (SR) for a volume on a XenServer host with SAS HBAs:

1. Log in to XenCenter and select the XenServer host.
2. Select the pool in the Resources pane.
3. Click **New Storage**.
The **New Storage Repository** wizard opens.
4. Select **Hardware HBA** as the storage type and click **Next**.
5. Type a name for the new SR in the **Name** field.
6. Click **Next**.
The wizard scans for available LUNs and then displays a page listing all the LUNs found.
7. Select the LUNs from the list of discovered LUNs to use for the new SR.

i **NOTE:** The storage target must be configured to enable every XenServer host in the pool to have access to one or more LUNs.
8. Click **Create**.
The **New Storage Repository** dialog box opens.

i **NOTE:** A warning message is displayed if there are existing SRs on the LUN that you have selected. Review the details, and perform one of the following actions:

 - Click **Reattach** to use the existing SR.
 - Click **Format** to delete the existing SR and to create a SR.
 - If you prefer to select a different LUN, click **Cancel** and select a different LUN from the list.
9. Click **Finish**.
The new SR is displayed in the Resources pane, at the pool level.

Troubleshooting and problem solving

These procedures are intended to be used only during initial configuration for verifying that hardware setup is successful. They are not intended to be used as troubleshooting procedures for configured systems using production data and I/O.

NOTE: For further troubleshooting help after setup, and when data is present, see Dell.com/support.

Topics:

- [Locate the service tag](#)
- [Operators \(Ops\) panel LEDs](#)
- [Initial start-up problems](#)

Locate the service tag

ME4 Series storage systems are identified by a unique Service Tag and Express Service Code.

The Service Tag and Express Service Code can be found on the front of the system by pulling out the information tag. Alternatively, the information might be on a sticker on the back of the storage system chassis. This information is used to route support calls to appropriate personnel.

Operators (Ops) panel LEDs

Each ME4 Series enclosure features an Operators (Ops) panel located on the chassis left ear flange. This section describes the Ops panel for 2U and 5U enclosures.

2U enclosure Ops panel

The front of the enclosure has an Ops panel that is located on the left ear flange of the 2U chassis.

The Ops panel is a part of the enclosure chassis, but is not replaceable on-site.

The Ops panel provides the functions that are shown in the following figure and listed in [Ops panel functions—2U enclosure front panel](#) on page 71.




Figure 30. Ops panel LEDs—2U enclosure front panel

Table 12. Ops panel functions—2U enclosure front panel

No.	Indicator	Status
1	System power	Constant green: at least one PCM is supplying power Off: system not operating regardless of AC present

Table 12. Ops panel functions—2U enclosure front panel (continued)

No.	Indicator	Status
2	Status/Health	Constant blue: system is powered on and controller is ready Blinking blue (2 Hz): Enclosure management is busy Constant amber: module fault present Blinking amber: logical fault (2 s on, 1 s off)
3	Unit identification display	Green (seven-segment display: enclosure sequence)
4	Identity	Blinking blue (0.25 Hz): system ID locator is activated Off: Normal state

System power LED (green)

LED displays green when system power is available. LED is off when system is not operating.

Status/Health LED (blue/amber)

LED illuminates constant blue when the system is powered on and functioning normally. The LED blinks blue when enclosure management is busy, for example, when booting or performing a firmware update. The LEDs help you identify which component is causing the fault. LED illuminates constant amber when experiencing a system hardware fault which could be associated with a Fault LED on a controller module, IOM, or PCM. LED illuminates blinking amber when experiencing a logical fault.

Unit identification display (green)

The UID is a dual seven-segment display that shows the numerical position of the enclosure in the cabling sequence. The UID is also called the enclosure ID.

NOTE: The controller enclosure ID is 0.

Identity LED (blue)

When activated, the Identity LED blinks at a rate of 1 second on, 1 second off to locate the chassis within a data center. The locate function can be enabled or disabled through SES. Pressing the button changes the state of the LED.

NOTE: The enclosure ID cannot be set using the Identity button.

5U enclosure Ops panel

The front of the enclosure has an Ops panel that is located on the left ear flange of the 5U chassis.

The Ops panel is part of the enclosure chassis, but is not replaceable on-site.

The Ops panel provides the functions that are shown in the following figure and listed in [Ops panel functions – 5U enclosure front panel](#) on page 73.




Figure 31. Ops panel LEDs—5U enclosure front panel

Table 13. Ops panel functions – 5U enclosure front panel

No.	Indicator	Status
1	Unit Identification Display (UID)	Green (seven-segment display: enclosure sequence)
2	System Power On/Standby	Constant green: positive indication Constant amber: system in standby (not operational)
3	Module Fault	Constant or blinking amber: fault present
4	Logical Status	Constant or blinking amber: fault present
5	Top Drawer Fault	Constant or blinking amber: fault present in drive, cable, or sideplane
6	Bottom Drawer Fault	Constant or blinking amber: fault present in drive, cable, or sideplane

Unit identification display

The UID is a dual seven-segment display that shows the numerical position of the enclosure in the cabling sequence. The UID is also called the enclosure ID.

NOTE: The controller enclosure ID is 0.

System Power On/Standby LED (green/amber)

LED is amber when only the standby power is available (non-operational). LED is green when system power is available (operational).

Module Fault LED (amber)

LED turns amber when experiencing a system hardware fault. The module fault LED helps you identify the component causing the fault. The module fault LED may be associated with a Fault LED on a controller module, IOM, PSU, FCM, DDIC, or drawer.

Logical Status LED (amber)

This LED indicates a change of status or fault from something other than the enclosure management system. The logical status LED may be initiated from the controller module or an external HBA. The indication is typically associated with a DDIC and LEDs at each disk position within the drawer, which help to identify the DDIC affected.

Drawer Fault LEDs (amber)

This LED indicates a disk, cable, or sideplane fault in the drawer indicate: Top (Drawer 0) or Bottom (Drawer 1).

Initial start-up problems

The following sections describe how to troubleshoot initial start-up problems:

LED colors

LED colors are used consistently throughout the enclosure and its components for indicating status:

- Green: good or positive indication
- Blinking green/amber: non-critical condition
- Amber: critical fault

Troubleshooting a host-side connection with 10Gbase-T or SAS host ports

The following procedure applies to ME4 Series controller enclosures employing external connectors in the host interface ports:

1. Stop all I/O to the storage system. See “Stopping I/O” in the *Dell EMC PowerVault ME4 Series Storage System Owner’s Manual*.
2. Check the host activity LED.

If there is activity, stop all applications that access the storage system.

3. Check the Cache Status LED to verify that the controller cached data is flushed to the disk drives.
 - Solid – Cache contains data yet to be written to the disk.
 - Blinking – Cache data is being written to CompactFlash.
 - Flashing at 1/10 second on and 9/10 second off – Cache is being refreshed by the supercapacitor.
 - Off – Cache is clean (no unwritten data).
4. Reseat the host cable and inspect for damage.

Is the host link status LED on?

- Yes – Monitor the status to ensure that there is no intermittent error present. If the fault occurs again, clean the connections to ensure that a dirty connector is not interfering with the data path.
- No – Proceed to the next step.

5. Move the host cable to a port with a known good link status.

This step isolates the problem to the external data path (host cable and host-side devices) or to the controller module port.

Is the host link status LED on?

- Yes – You now know that the host cable and host-side devices are functioning properly. Return the cable to the original port. If the link status LED remains off, you have isolated the fault to the controller module port. Replace the controller module.
- No – Proceed to the next step.

6. Verify that the switch, if any, is operating properly. If possible, test with another port.
7. Verify that the HBA is fully seated, and that the PCI slot is powered on and operational.
8. Replace the HBA with a known good HBA, or move the host side cable to a known good HBA.

Is the host link status LED on?

- Yes – You have isolated the fault to the HBA. Replace the HBA.
- No – It is likely that the controller module needs to be replaced.

9. Move the host cable back to its original port.

Is the host link status LED on?

- No – The controller module port has failed. Replace the controller module.
- Yes – Monitor the connection. It may be an intermittent problem, which can occur with damaged cables and HBAs.

Isolating controller module expansion port connection faults

During normal operation, when a controller module expansion port is connected to a drive enclosure, the expansion port status LED is green. If the expansion port LED is off, the link is down. Use the following procedure to isolate the fault:

NOTE: Do not perform more than one step at a time. Changing more than one variable at a time can complicate the troubleshooting process.

1. Stop all I/O to the storage system. See “Stopping I/O” in the Dell EMC PowerVault ME4 Series Storage System Owner’s Manual.
2. Check the host activity LED.
If there is activity, stop all applications that access the storage system.
3. Check the Cache Status LED to verify that the controller cached data is flushed to the disk drives.
 - Solid – Cache contains data yet to be written to the disk.
 - Blinking – Cache data is being written to CompactFlash.
 - Flashing at 1/10 second on and 9/10 second off – Cache is being refreshed by the supercapacitor.
 - Off – Cache is clean (no unwritten data).
4. Reseat the expansion cable, and inspect it for damage.
Is the expansion port status LED on?
 - Yes – Monitor the status to ensure that there is no intermittent error present. If the fault occurs again, clean the connections to ensure that a dirty connector is not interfering with the data path.
 - No – Proceed to the next step.
5. Move the expansion cable to a port on the controller enclosure with a known good link status.
This step isolates the problem to the expansion cable or to the controller module expansion port.
Is the expansion port status LED on?
Yes – You now know that the expansion cable is good. Return the cable to the original port. If the expansion port status LED remains off, you have isolated the fault to the controller module expansion port. Replace the controller module.
No – Proceed to the next step.
6. Move the expansion cable back to the original port on the controller enclosure.
7. Move the expansion cable on the drive enclosure to a known good expansion port on the drive enclosure.
Is the expansion port status LED on?
 - Yes – You have isolated the problem to the expansion enclosure port. Replace the expansion module.
 - No – Proceed to the next step.
8. Replace the cable with a known good cable, ensuring the cable is attached to the original ports.
Is the host link status LED on?
 - Yes – Replace the original cable. The fault has been isolated.
 - No – It is likely that the controller module must be replaced.

2U enclosure LEDs

Use the LEDs on the 2U enclosure to help troubleshoot initial start-up problems.

2U PCM LEDs (580 W)

Under normal conditions, the PCM OK LEDs are a constant green.

Table 14. PCM LED states

PCM OK (Green)	Fan Fail (Amber)	AC Fail (Amber)	DC Fail (Amber)	Status
Off	Off	Off	Off	No AC power on any PCM
Off	Off	On	On	No AC power on this PCM only
On	Off	Off	Off	AC present; PCM working correctly
On	Off	Off	On	PCM fan speed is outside acceptable limits
Off	On	Off	Off	PCM fan has failed

Table 14. PCM LED states (continued)

PCM OK (Green)	Fan Fail (Amber)	AC Fail (Amber)	DC Fail (Amber)	Status
Off	On	On	On	PCM fault (over temperature, over voltage, over current)
Off	Blinking	Blinking	Blinking	PCM firmware download is in progress

2U Ops panel LEDs

The Ops panel displays the aggregated status of all the modules. See also [2U enclosure Ops panel](#) on page 71.

Table 15. Ops panel LED states

System Power (Green/Amber)	Module Fault (Amber)	Identity (Blue)	LED display	Associated LEDs/Alarms	Status
On	Off	Off	X		5 V standby power present, overall power has failed or switched off
On	On	On	On		Ops panel power on (5 s) test state
On	Off	Off	X		Power on, all functions good
On	On	X	X	PCM fault LEDs, fan fault LEDs	Any PCM fault, fan fault, over or under temperature
On	On	X	X	SBB module LEDs	Any SBB module fault
On	On	X	X	No module LEDs	Enclosure logical fault
On	Blink	X	X	Module status LED on SBB module	Unknown (invalid or mixed) SBB module type is installed, I ² C bus failure (inter-SBB communications). EBOD VPD configuration error
On	Blink	X	X	PCM fault LEDs, fan fault LEDs	Unknown (invalid or mixed) PCM type is installed or I ² C bus failure (PCM communications)
		X	Blink		Enclosure identification or invalid ID selected

X= Disregard

Actions:

- If the Ops panel Module Fault LED is on, check the module LEDs on the enclosure rear panel to narrow the fault to a CRU, a connection, or both.
- Check the event log for specific information regarding the fault, and follow any Recommended Actions.
- If installing a controller module or IOM CRU:
 - Remove and reinstall the controller module or IOM per the *Dell EMC PowerVault ME4 Series Storage System Owner's Manual*.
 - Check the event log for errors.
- If the CRU Fault LED is on, a fault condition is detected.
 - Restart this controller from the partner controller using the PowerVault Manager or CLI.
 - If the restart does not resolve the fault, remove the controller module or IOM and reinsert it.
- If the previous actions do not resolve the fault, contact Dell EMC for assistance.

2U disk drive carrier module LEDs

A green LED and amber LED mounted on the front of each drive carrier module display the disk drive status.

- In normal operation, the green LED is on and flickers as the drive operates.

- In normal operation, the amber LED is:
 - Off if there is no drive present.
 - Off as the drive operates.
 - On if there is a drive fault.




Figure 32. LEDs: Drive carrier LEDs (SFF and LFF modules) used in 2U enclosures

- | | |
|----------------------|----------------------|
| 1. Disk Activity LED | 2. Disk Fault LED |
| 3. Disk Fault LED | 4. Disk Activity LED |

5U enclosure LEDs

Use the LEDs on the 5U enclosure to help troubleshoot initial start-up problems.

NOTE: When the 5U84 enclosure is powered on, all LEDs are lit for a short period to ensure that they are working. This behavior does not indicate a fault unless LEDs remain lit after several seconds.

5U PSU LEDs

The following table describes the LED states for the PSU:

Table 16. PSU LED states

CRU Fail (Amber)	AC Missing (Amber)	Power (Green)	Status
On	Off	Off	No AC power to either PSU
On	On	Off	PSU present, but not supplying power or PSU alert state. (typically due to critical temperature)
Off	Off	On	Mains AC present, switch on. This PSU is providing power.
Off	Off	Blinking	AC power present, PSU in standby (other PSU is providing power).
Blinking	Blinking	Off	PSU firmware download in progress
Off	On	Off	AC power missing, PSU in standby (other PSU is providing power).
On	On	On	Firmware has lost communication with the PSU module.
On	--	Off	PSU has failed. Follow the procedure in "Replacing a PSU" in the <i>Dell EMC PowerVault ME4 Series Storage System Owner's Manual</i> .

5U FCM LEDs

The following table describes the LEDs on the Fan Cooling Module (FCM) faceplate:

Table 17. FCM LED descriptions

LED	Status/description
Module OK	Constant green indicates that the FCM is working correctly. Off indicates that the fan module has failed. Follow the procedure in "Replacing an FCM" in the <i>Dell EMC PowerVault ME4 Series Storage System Owner's Manual</i> .
Fan Fault	Amber indicates that the fan module has failed. Follow the procedure in "Replacing an FCM" in the <i>Dell EMC PowerVault ME4 Series Storage System Owner's Manual</i> .

5U Ops panel LEDs

The Ops panel displays the aggregated status of all the modules.

Table 18. Ops panel LED descriptions

LED	Status/description
Unit ID display	Usually shows the ID number for the enclosure, but can be used for other purposes, for example, blinking to locate enclosure.
Power On/Standby	Amber if the system is in standby. Green if the system has full power.
Module Fault	Amber indicates a fault in a controller module, IOM, PSU, or FCM. Check the drawer LEDs for indication of a disk fault. See also Drawer Fault LEDs (amber) on page 73.
Logical status	Amber indicates a fault from something other than firmware (usually a disk, an HBA, or an internal or external RAID controller). Check the drawer LEDs for indication of a disk fault. See also 5U drawer LEDs on page 78.
Drawer 0 Fault	Amber indicates a disk, cable, or sideplane fault in drawer 0. Open the drawer and check DDICs for faults.
Drawer 1 Fault	Amber indicates a disk, cable, or sideplane fault in drawer 1. Open the drawer and check DDICs for faults.

5U drawer LEDs

The following table describes the LEDs on the drawers:

Table 19. Drawer LED descriptions

LED	Status/description
Sideplane OK/Power Good	Green if the sideplane card is working and there are no power problems.
Drawer Fault	Amber if a drawer component has failed. If the failed component is a disk, the LED on the failed DDIC lights amber. Follow the procedure in "Replacing a DDIC" in the <i>Dell EMC PowerVault ME4 Series Storage System Owner's Manual</i> . If the disks are OK, contact your service provider to identify the cause of the failure, and resolve the problem.
Logical Fault	Amber (solid) indicates a disk fault. Amber (blinking) indicates that one or more storage systems are in an impacted state.
Cable Fault	Amber indicates the cabling between the drawer and the back of the enclosure has failed. Contact your service provider to resolve the problem.
Activity Bar Graph	Displays the amount of data I/O from zero segments lit (no I/O) to all six segments lit (maximum I/O).

5U DDIC LED

The DDIC supports LFF 3.5" and SFF 2.5" disks as shown in [3.5" disk drive in a DDIC](#) on page 13 and [2.5" drive in a 3.5" DDIC with a hybrid drive carrier adapter](#) on page 14). The following figure shows the top panel of the DDIC as viewed when the disk is aligned for insertion into a drawer slot.




Figure 33. LEDs: DDIC – 5U enclosure disk slot in drawer

1. Slide latch (slides left)
2. Latch button (shown in the locked position)
3. Drive Fault LED

Table 20. DDIC LED descriptions

Fault LED (Amber)	Status/description*
Off	Off (disk module/enclosure)
Off	Not present
Blinking: 1 s on/1 s off	Identify
Any links down: On	Drive link (PHY lane) down
On	Fault (leftover/failed/locked-out)
Off	Available
Off	Storage system: Initializing
Off	Storage system: Fault-tolerant
Off	Storage system: Degraded (non-critical)
Blinking: 3 s on/1 s off	Storage system: Degraded (critical)
Off	Storage system: Quarantined
Blinking: 3 s on/1 s off	Storage system: Offline (dequarantined)
Off	Storage system: Reconstruction
Off	Processing I/O (whether from host or internal activity)

*If multiple conditions occur simultaneously, the LED state behaves as indicated in the previous table.

Each DDIC has a single Drive Fault LED. If the Drive Fault LED is lit amber, a disk fault is indicated. If a disk failure occurs, follow the procedure in "Replacing a DDIC" in the *Dell EMC PowerVault ME4 Series Storage System Owner's Manual*.

5U controller module or IOM LEDs

- For information about controller module LEDs, see [Controller module LEDs](#) on page 80.
- For information about expansion module LEDs, see [IOM LEDs](#) on page 80.

5U temperature sensors

Temperature sensors throughout the enclosure and its components monitor the thermal health of the storage system. Exceeding the limits of critical values causes a notification to occur.

Module LEDs

Module LEDs pertain to controller modules and IOMs.

Controller module LEDs

Use the controller module LEDs on the face plate to monitor the status of a controller module.

Table 21. Controller module LED states

CRU OK (Green)	CRU Fault (Amber)	External host port activity (Green)	Status
On	Off		Controller module OK
Off	On		Controller module fault – see “Replacing a controller module” in the <i>Dell EMC PowerVault ME4 Series Storage System Owner’s Manual</i>
		Off	No external host port connection
		On	External host port connection – no activity
		Blinking	External host port connection – activity
Blinking			System is booting

Actions:

- If the CRU OK LED is blinking, wait for the system to boot.
- If the CRU OK LED is off, and the controller module is powered on, the module has failed.
 - Check that the controller module is fully inserted and latched in place, and that the enclosure is powered on.
 - Check the event log for specific information regarding the failure.
- If the CRU Fault LED is on, a fault condition is detected.
 - Restart the controller module from the partner controller module using the PowerVault Manager or CLI.
 - If the restart does not resolve the fault, remove the controller module and reinsert it.
- If the previous actions do not resolve the fault, contact your supplier for assistance. Controller module replacement may be necessary.

IOM LEDs

Use the IOM LEDs on the face plate to monitor the status of an IOM .

Table 22. IOM LED states

CRU OK (Green)	CRU Fault (Amber)	External host port activity (Green)	Status
On	Off		Controller module OK
Off	On		IOM module fault – see “Replacing an IOM” in the <i>Dell EMC PowerVault ME4 Series Storage System Owner’s Manual</i>
		Off	No external host port connection
		On	HD mini-SAS port connection – no activity
		Blinking	HD mini-SAS port connection – activity
Blinking			EBOD VPD error

Actions:

- If the CRU OK LED is off, and the IOM is powered on, the module has failed.
 - Check that the IOM is fully inserted and latched in place, and that the enclosure is powered on.
 - Check the event log for specific information regarding the failure.

- If the CRU Fault LED is on, a fault condition is detected.
 - Restart this IOM using the PowerVault Manager or CLI.
 - If the restart does not resolve the fault, remove the IOM and reinsert it.
- If the previous actions do not resolve the fault, contact your supplier for assistance. IOM replacement may be necessary.

Troubleshooting 2U enclosures

Common problems that may occur with your 2U enclosure system.

The Module Fault LED on the Ops panel, described in [Ops panel LEDs—2U enclosure front panel](#) on page 71, lights amber to indicate a fault for the problems listed in the following table:


 **NOTE:** All alarms also report through SES.

Table 23. Troubleshooting 2U alarm conditions

Status	Severity	Alarm
PCM alert – loss of DC power from a single PCM	Fault – loss of redundancy	S1
PCM fan fail	Fault – loss of redundancy	S1
SBB module detected PCM fault	Fault	S1
PCM removed	Configuration error	None
Enclosure configuration error (VPD)	Fault – critical	S1
Low warning temperature alert	Warning	S1
High warning temperature alert	Warning	S1
Over-temperature alarm	Fault – critical	S4
I ² C bus failure	Fault – loss of redundancy	S1
Ops panel communication error (I ² C)	Fault – critical	S1
RAID error	Fault – critical	S1
SBB interface module fault	Fault – critical	S1
SBB interface module removed	Warning	None
Drive power control fault	Warning – no loss of disk power	S1
Drive power control fault	Fault – critical – loss of disk power	S1
Drive removed	Warning	None
Insufficient power available	Warning	None

For details about replacing modules, see the *Dell EMC PowerVault ME4 Series Storage System Owner's Manual*.


 **NOTE:** Use the PowerVault Manager to monitor the storage system event logs for information about enclosure-related events, and to determine any necessary recommended actions.

Table 24. Troubleshooting PCM faults

Symptom	Cause	Recommended action
Ops panel Module Fault LED is amber ¹	Any power fault	Verify that AC mains connections to PCM are live
Fan Fail LED is illuminated on PCM ²	Fan failure	Replace PCM

1. See [2U enclosure Ops panel](#) on page 71 for visual reference of Ops panel LEDs.
2. See [2U PCM LEDs \(580 W\)](#) on page 75 for visual reference of PCM LEDs.

The storage enclosure uses extensive thermal monitoring and takes several actions to ensure that component temperatures are kept low, and to also minimize acoustic noise. Air flow is from the front to back of the enclosure.

Table 25. Troubleshooting thermal monitoring and control

Symptom	Cause	Recommended action
<p>If the ambient air is below 25°C (77°F), and the fans increase in speed, some restriction on airflow may be causing the internal temperature to rise.</p> <p>i NOTE: This symptom is not a fault condition.</p>	<p>The first stage in the thermal control process is for the fans to automatically increase in speed when a thermal threshold is reached. This condition may be caused higher ambient temperatures in the local environment, and may be normal condition.</p> <p>i NOTE: The threshold changes according to the number of disks and power supplies fitted.</p>	<ol style="list-style-type: none"> 1. Check the installation for any airflow restrictions at either the front or back of the enclosure. A minimum gap of 25 mm (1") at the front and 50 mm (2") at the rear is recommended. 2. Check for restrictions due to dust build-up. Clean as appropriate. 3. Check for excessive recirculation of heated air from rear to front. Use of the enclosure in a fully enclosed rack is not recommended. 4. Verify that all blank modules are in place. 5. Reduce the ambient temperature.

Table 26. Troubleshooting thermal alarm

Symptom	Cause	Recommended action
<ol style="list-style-type: none"> 1. Ops panel Module Fault LED is amber. 2. Fan Fail LED is illuminated on one or more PCMs. 	<p>Internal temperature exceeds a preset threshold for the enclosure.</p>	<ol style="list-style-type: none"> 1. Verify that the local ambient environment temperature is within the acceptable range. See the technical specifications in the <i>Dell EMC PowerVault ME4 Series Storage System Owner's Manual</i>. 2. Check the installation for any airflow restrictions at either the front or back of the enclosure. A minimum gap of 25 mm (1") at the front and 50 mm (2") at the rear is recommended. 3. Check for restrictions due to dust build-up. Clean as appropriate. 4. Check for excessive recirculation of heated air from rear to front. Use of the enclosure in a fully enclosed rack is not recommended. 5. If possible, shut down the enclosure and investigate the problem before continuing.

Troubleshooting 5U enclosures

Common problems that may occur with your 5U enclosure system.

The Module Fault LED on the Ops panel, described in [Ops panel LEDs—5U enclosure front panel](#) on page 73, lights amber to indicate a fault for the problems listed in the following table:

i **NOTE:** All alarms also report through SES.

Table 27. 5U alarm conditions

Status	Severity
PSU alert – loss of DC power from a single PSU	Fault – loss of redundancy
Cooling module fan failure	Fault – loss of redundancy
SBB I/O module detected PSU fault	Fault
PSU removed	Configuration error
Enclosure configuration error (VPD)	Fault – critical
Low temperature warning	Warning
High temperature warning	Warning
Over-temperature alarm	Fault – critical

Table 27. 5U alarm conditions (continued)

Status	Severity
Under-temperature alarm	Fault – critical
I ² C bus failure	Fault – loss of redundancy
Ops panel communication error (I ² C)	Fault – critical
RAID error	Fault – critical
SBB I/O module fault	Fault – critical
SBB I/O module removed	Warning
Drive power control fault	Warning – no loss of drive power
Drive power control fault	Fault – critical – loss of drive power
Insufficient power available	Warning

For details about replacing modules, see the *Dell EMC PowerVault ME4 Series Storage System Owner's Manual*.

NOTE: Use the PowerVault Manager to monitor the storage system event logs for information about enclosure-related events, and to determine any necessary recommended actions.

Thermal considerations

NOTE: Thermal sensors in the 5U84 enclosure and its components monitor the thermal health of the storage system.

- Exceeding the limits of critical values activates the over-temperature alarm.
- For information about 5U84 enclosure alarm notification, see [5U alarm conditions](#) on page 82.

Fault isolation methodology

ME4 Series Storage Systems provide many ways to isolate faults. This section presents the basic methodology that is used to locate faults within a storage system, and to identify the pertinent CRUs affected.

As noted in [Using guided setup](#) on page 33, use the PowerVault Manager to configure and provision the system upon completing the hardware installation. Configure and enable event notification to be notified when a problem occurs that is at or above the configured severity. See the *Dell EMC PowerVault ME4 Series Storage System Administrator's Guide* for more information.

When you receive an event notification, follow the recommended actions in the notification message to resolve the problem.

Fault isolation methodology basic steps

- Gather fault information, including using system LEDs as described in [Gather fault information](#) on page 84.
- Determine where in the system the fault is occurring as described in [Determine where the fault is occurring](#) on page 84.
- Review event logs as described in [Review the event logs](#) on page 85.
- If required, isolate the fault to a data path component or configuration as described in [Isolate the fault](#) on page 85.

Cabling systems to enable use of the replication feature—to replicate volumes—is another important fault isolation consideration pertaining to initial system installation. See [Host ports and replication](#) on page 91 and [Isolating replication faults](#) on page 94 for more information about troubleshooting during initial setup.

Options available for performing basic steps

When performing fault isolation and troubleshooting steps, select the option or options that best suit your site environment.

Use of any option is not mutually exclusive to the use of another option. You can use the PowerVault Manager to check the health icons/values for the system, or to examine a problem component. If you discover a problem, either the PowerVault Manager or the CLI provides recommended-action text online. Options for performing basic steps are listed according to frequency of use:

- Use the PowerVault Manager
- Use the CLI
- Monitor event notification
- View the enclosure LEDs

Use the PowerVault Manager

The PowerVault Manager uses health icons to show OK, Degraded, Fault, or Unknown status for the system and its components. The PowerVault Manager enables you to monitor the health of the system and its components. If any component has a problem, the system health is in a Degraded, Fault, or Unknown state. Use the PowerVault Manager to find each component that has a problem. Follow actions in the Recommendation field for the component to resolve the problem.

Use the CLI

As an alternative to using the PowerVault Manager, you can run the `show system` CLI command to view the health of the system and its components. If any component has a problem, the system health is in a Degraded, Fault, or Unknown state, and those components are listed as Unhealthy Components. Follow the recommended actions in the component Health Recommendation field to resolve the problem.

Monitor event notification

With event notification configured and enabled, you can view event logs to monitor the health of the system and its components. If a message tells you to check whether an event has been logged, or to view information about an event, use the PowerVault Manager or the CLI. Using the PowerVault Manager, view the event log and then click the event message to see detail about that event. Using the CLI, run the `show events detail` command to see the detail for an event.

View the enclosure LEDs

You can view the LEDs on the hardware to identify component status. If a problem prevents access to the PowerVault Manager or the CLI, viewing the enclosure LEDs is the only option available. However, monitoring/management is often done at a management console using storage management interfaces, rather than relying on line-of-sight to LEDs of racked hardware components.

Performing basic steps

You can use any of the available options that are described in the previous sections to perform the basic steps comprising the fault isolation methodology.

Gather fault information

When a fault occurs, gather as much information as possible. Doing so helps determine the correct action that is needed to remedy the fault.

Begin by reviewing the reported fault:

- Is the fault related to an internal data path or an external data path?
- Is the fault related to a hardware component such as a disk drive module, controller module, or power supply unit?

By isolating the fault to one of the components within the storage system, you are able determine the necessary corrective action more quickly.

Determine where the fault is occurring

When a fault occurs, the Module Fault LED illuminates. Check the LEDs on the back of the enclosure to narrow the fault to a CRU, connection, or both. The LEDs also help you identify the location of a CRU reporting a fault.

Use the PowerVault Manager to verify any faults found while viewing the LEDs. If the LEDs cannot be viewed due to the location of the system, use the PowerVault Manager to determine where the fault is occurring. This web-application provides you with a visual representation of the system and where the fault is occurring. The PowerVault Manager also provides more detailed information about CRUs, data, and faults.

Review the event logs

The event logs record all system events. Each event has a numeric code that identifies the type of event that occurred, and has one of the following severities:

- Critical – A failure occurred that may cause a controller to shut down. Correct the problem immediately.
- Error – A failure occurred that may affect data integrity or system stability. Correct the problem as soon as possible.
- Warning – A problem occurred that may affect system stability, but not data integrity. Evaluate the problem and correct it if necessary.
- Informational – A configuration or state change occurred, or a problem occurred that the system corrected. No immediate action is required.

The event logs record all system events. Review the logs to identify fault and cause of the failure. For example, a host could lose connectivity to a disk group if a user changes channel settings without taking the storage resources that are assigned to it into consideration. In addition, the type of fault can help you isolate the problem to either hardware or software.

Isolate the fault

Occasionally, it might become necessary to isolate a fault. This is true with data paths, due to the number of components comprising the data path. For example, if a host-side data error occurs, it could be caused by any of the components in the data path: controller module, cable, or data host.

If the enclosure does not initialize

It may take up to two minutes for all enclosures to initialize.


If an enclosure does not initialize:

- Perform a rescan
- Power cycle the system
- Make sure that the power cord is properly connected, and check the power source to which it is connected
- Check the event log for errors

Correcting enclosure IDs

When installing a system with expansion enclosures attached, the enclosure IDs might not agree with the physical cabling order. This issue occurs if the controller was previously attached to enclosures in a different configuration, and the controller attempts to preserve the previous enclosure IDs.

To correct this condition, ensure that both controllers are up, and perform a rescan using the PowerVault Manager or the CLI. The rescan reorders the enclosures, but it can take up to two minutes to correct the enclosure IDs.

 **NOTE:** Reordering expansion enclosure IDs only applies to dual-controller mode. If only one controller is available, due to a controller failure, a manual rescan does not reorder the expansion enclosure IDs.

- To perform a rescan using the PowerVault Manager:
 1. Verify that both controllers are operating normally.
 2. In the System tab, click **Action**, and select **Rescan Disk Channels**.
- To perform a rescan using the CLI, type the following command:

```
rescan
```

Host I/O

When troubleshooting disk drive and connectivity faults, stop I/O to the affected disk groups from all hosts as a data protection precaution.

As an extra data protection precaution, it is helpful to conduct regularly scheduled backups of your data. See “Stopping I/O” in the *Dell EMC PowerVault ME4 Series Storage System Owner’s Manual*.

Dealing with hardware faults

Make sure that you have a replacement module of the same type before removing any faulty module. See “Module removal and replacement” in the *Dell EMC PowerVault ME4 Series Storage System Owner’s Manual*.

NOTE: If the enclosure system is powered up and you remove any module, replace it immediately. If the system is used with any modules missing for more than a few seconds, the enclosures can overheat, causing power failure and potential data loss. Such action can invalidate the product warranty.

NOTE: Observe applicable/conventional ESD precautions when handling modules and components, as described in [Electrical safety](#) on page 8. Avoid contact with midplane components, module connectors, leads, pins, and exposed circuitry.

Isolating a host-side connection fault

During normal operation, when a controller module host port is connected to a data host, the port host link status/link activity LED is green. If there is I/O activity, the host activity LED blinks green. If data hosts are having trouble accessing the storage system, but you cannot locate a specific fault or access the event logs, use the following procedures. These procedures require scheduled downtime.

NOTE: Do not perform more than one step at a time. Changing more than one variable at a time can complicate the troubleshooting process.

Host-side connection troubleshooting featuring CNC ports

The following procedure applies to controller enclosures with small form factor pluggable (SFP+) transceiver connectors in 8/16 Gb/s FC or 10 GbE iSCSI host interface ports.

In this procedure, `SFP+ transceiver` and `host cable` is used to refer to any qualified SFP+ transceiver supporting CNC ports used for I/O or replication.

NOTE: When experiencing difficulty diagnosing performance problems, consider swapping out one SFP+ transceiver at a time to see if performance improves.

1. Stop all I/O to the storage system. See “Stopping I/O” in the *Dell EMC PowerVault ME4 Series Storage System Owner’s Manual*.
2. Check the host link status/link activity LED.
If there is activity, stop all applications that access the storage system.
3. Check the Cache Status LED to verify that the controller cached data is flushed to the disk drives.
 - Solid – Cache contains data yet to be written to the disk.
 - Blinking – Cache data is being written to CompactFlash in the controller module.
 - Flashing at 1/10 second on and 9/10 second off – Cache is being refreshed by the supercapacitor.
 - Off – Cache is clean (no unwritten data).
4. Remove the SFP+ transceiver and host cable and inspect for damage.
5. Reseat the SFP+ transceiver and host cable.
Is the host link status/link activity LED on?
 - Yes – Monitor the status to ensure that there is no intermittent error present. If the fault occurs again, clean the connections to ensure that a dirty connector is not interfering with the data path.
 - No – Proceed to the next step.
6. Move the SFP+ transceiver and host cable to a port with a known good link status.

This step isolates the problem to the external data path (SFP+ transceiver, host cable, and host-side devices) or to the controller module port.

Is the host link status/link activity LED on?

- Yes – You now know that the SFP+ transceiver, host cable, and host-side devices are functioning properly. Return the cable to the original port. If the link status LED remains off, you have isolated the fault to the controller module port. Replace the controller module.
- No – Proceed to the next step.

7. Swap the SFP+ transceiver with the known good one.

Is the host link status/link activity LED on?

- Yes – You have isolated the fault to the SFP+ transceiver. Replace the SFP+ transceiver.
- No – Proceed to the next step.

8. Reinsert the original SFP+ transceiver and swap the cable with a known good one.

Is the host link status/link activity LED on?

- Yes – You have isolated the fault to the cable. Replace the cable.
- No – Proceed to the next step.

9. Verify that the switch, if any, is operating properly. If possible, test with another port.

10. Verify that the HBA is fully seated, and that the PCI slot is powered on and operational.

11. Replace the HBA with a known good HBA, or move the host side cable and SFP+ transceiver to a known good HBA.

Is the host link status/link activity LED on?

- Yes – You have isolated the fault to the HBA. Replace the HBA.
- No – It is likely that the controller module needs to be replaced.

12. Move the cable and SFP+ transceiver back to its original port.

Is the host link status/link activity LED on?

- Yes – Monitor the connection for a period of time. It may be an intermittent problem, which can occur with damaged SFP+ transceivers, cables, and HBAs.
- No – The controller module port has failed. Replace the controller module.

Host-side connection troubleshooting featuring 10Gbase-T and SAS host ports

The following procedure applies to ME4 Series controller enclosures employing external connectors in the host interface ports.

The external connectors include 10Gbase-T connectors in iSCSI host ports and 12 Gb SFF-8644 connectors in the HD mini-SAS host ports.

1. Halt all I/O to the storage system. See “Stopping I/O” in the *Dell EMC PowerVault ME4 Series Storage System Owner’s Manual*.

2. Check the host activity LED.

If there is activity, stop all applications that access the storage system.

3. Check the Cache Status LED to verify that the controller cached data is flushed to the disk drives.

- Solid – Cache contains data yet to be written to the disk.
- Blinking – Cache data is being written to CompactFlash in the controller module.
- Flashing at 1/10 second on and 9/10 second off – Cache is being refreshed by the supercapacitor.
- Off – Cache is clean (no unwritten data).

4. Remove the host cable and inspect for damage.

5. Reseat the host cable.

Is the host link status LED on?

- Yes – Monitor the status to ensure that there is no intermittent error present. If the fault occurs again, clean the connections to ensure that a dirty connector is not interfering with the data path.
- No – Proceed to the next step.

6. Move the host cable to a port with a known good link status.

This step isolates the problem to the external data path (host cable and host-side devices) or to the controller module port.

Is the host link status LED on?

- Yes – You now know that the host cable and host-side devices are functioning properly. Return the cable to the original port. If the link status LED remains off, you have isolated the fault to the controller module port. Replace the controller module.
 - No – Proceed to the next step.
7. Verify that the switch, if any, is operating properly. If possible, test with another port.
 8. Verify that the HBA is fully seated, and that the PCI slot is powered on and operational.
 9. Replace the HBA with a known good HBA, or move the host side cable to a known good HBA.

Is the host link status LED on?

 - Yes – You have isolated the fault to the HBA. Replace the HBA.
 - No – It is likely that the controller module needs to be replaced.
 10. Move the host cable back to its original port.


Is the host link status LED on?

 - Yes – Monitor the connection for a period of time. It may be an intermittent problem, which can occur with damaged cables and HBAs.
 - No – The controller module port has failed. Replace the controller module.

Isolating a controller module expansion port connection fault

During normal operation, when a controller module expansion port is connected to an expansion enclosure, the expansion port status LED is green. If the expansion port LED is off, the link is down.

Use the following procedure to isolate the fault. This procedure requires scheduled downtime.

 **NOTE:** Do not perform more than one step at a time. Changing more than one variable at a time can complicate the troubleshooting process.

1. Halt all I/O to the storage system. See “Stopping I/O” in the *Dell EMC PowerVault ME4 Series Storage System Owner’s Manual*.
2. Check the host activity LED.

If there is activity, stop all applications that access the storage system.
3. Check the Cache Status LED to verify that the controller cached data is flushed to the disk drives.
 - Solid – Cache contains data yet to be written to the disk.
 - Blinking – Cache data is being written to CompactFlash in the controller module.
 - Flashing at 1/10 second on and 9/10 second off – Cache is being refreshed by the supercapacitor.
 - Off – Cache is clean (no unwritten data).
4. Remove expansion cable and inspect for damage.
5. Reseat the expansion cable.

Is the expansion port status LED on?

 - Yes – Monitor the status to ensure that there is no intermittent error present. If the fault occurs again, clean the connections to ensure that a dirty connector is not interfering with the data path.
 - No – Proceed to the next step.
6. Move the expansion cable to a port on the controller enclosure with a known good link status.

This step isolates the problem to the expansion cable or to the controller module expansion port.

Is the expansion port status LED on?

 - Yes – You now know that the expansion cable is good. Return the cable to the original port. If the expansion port status LED remains off, you have isolated the fault to the controller module expansion port. Replace the controller module.
 - No – Proceed to the next step.
7. Move the expansion cable back to the original port on the controller enclosure.
8. Move the expansion cable on the expansion enclosure to a known good port on the expansion enclosure.

Is the host link status LED on?

 - Yes – You have isolated the problem to the expansion enclosure port. Replace the IOM in the expansion enclosure.
 - No – Proceed to the next step.
9. Replace the cable with a known good cable, ensuring the cable is attached to the original ports used by the previous cable.

Is the host link status LED on?

- Yes – Replace the original cable. The fault has been isolated.
- No – It is likely that the controller module must be replaced.

Cabling for replication

The following sections describe how to cable storage systems for replication:

Topics:

- [Connecting two storage systems to replicate volumes](#)
- [Host ports and replication](#)
- [Example cabling for replication](#)
- [Isolating replication faults](#)


Connecting two storage systems to replicate volumes

The replication feature performs asynchronous replication of block-level data from a volume in a primary system to a volume in a secondary system.

Replication creates an internal snapshot of the primary volume, and copies the changes to the data since the last replication to the secondary system using FC or iSCSI links.

The two associated standard volumes form a replication set, and only the primary volume (source of data) can be mapped for access by a server. Both systems must be connected through switches to the same fabric or network (no direct attach). The server accessing the replication set is connected to the primary system. If the primary system goes offline, a connected server can access the replicated data from the secondary system.


Systems can be cabled to support replication using CNC-based and 10Gbase-T systems on the same network, or on different networks.

 **NOTE:** SAS systems do not support replication.

As you consider the physical connections of your system, keep several important points in mind:

- Ensure that controllers have connectivity between systems, whether the destination system is colocated or remotely located.
- Qualified Converged Network Controller options can be used for host I/O or replication, or both.
- The storage system does not provide for specific assignment of ports for replication. However, this configuration can be accomplished using virtual LANs for iSCSI and zones for FC, or by using physically separate infrastructure.
- For remote replication, ensure that all ports that are assigned for replication can communicate with the replication system by using the query peer-connection CLI command. See the *ME4 Series Storage System CLI Reference Guide* for more information.
- Allow enough ports for replication permits so that the system can balance the load across those ports as I/O demands rise and fall. If controller A owns some of the volumes that are replicated and controller B owns other volumes that are replicated, then enable at least one port for replication on each controller module. You may need to enable more than one port per controller module depending on replication traffic load.
- For the sake of system security, do not unnecessarily expose the controller module network port to an external network connection.

Conceptual cabling examples are provided addressing cabling on the same network and cabling relative to different networks.

 **NOTE:**
The controller module firmware must be compatible on all systems that are used for replication.

Host ports and replication

ME4 Series Storage System controller modules can use qualified 10Gbase-T connectors or CNC-based ports for replication.

CNC ports must use qualified SFP+ transceivers of the same type, or they can use a combination of qualified SFP+ transceivers supporting different interface protocols. To use a combination of different protocols, configure host ports 0 and 1 to use FC, and configure ports 2 and 3 to use iSCSI. FC and iSCSI ports can be used to perform host I/O or replication, or both.

NOTE: ME4 Series 5U84 enclosures support dual-controller configurations only. ME4 Series 2U controller enclosures support single controller and dual-controller configurations.

- If a partner controller module fails, the storage system fails over and runs on a single controller module until the redundancy is restored.
- In dual-controller module configurations, a controller module must be installed in each slot to ensure sufficient airflow through the enclosure during operation. In single-controller module configurations, a controller module must be installed in slot A, and controller module blank must be installed in slot B.

Example cabling for replication

Simplified versions of controller enclosures are used in the cabling figures to show the host ports that are used for I/O or replication.

- Replication supports FC and iSCSI host interface protocols.
- The 2U enclosure rear panel represents ME4 Series FC and iSCSI host interface ports.
- The 5U84 enclosure rear panel represents ME4 Series FC and iSCSI host interface ports.
- Host ports that are used for replication must use the same protocol (either FC or iSCSI).
- Blue cables show I/O traffic and green cables show replication traffic.

Once the CNC-based systems or 10Gbase-T systems are physically cabled, see the *Dell EMC PowerVault ME4 Series Storage System Administrator's Guide* or online help for information about configuring, provisioning, and using the replication feature.

Single-controller module configuration for replication

Cabling two ME4 Series controller enclosures that are equipped with a single controller module for replication.

Multiple servers, multiple switches, one network

The following diagram shows the rear panel of two controller enclosures with I/O and replication occurring on the same network:




Figure 34. Connecting two storage systems for replication – multiple servers, multiple switches, one network

1. 2U controller enclosures
2. Two switches (I/O)
3. Connection to host servers
4. Switch (Replication)

For optimal protection, use multiple switches for host I/O and replication.

- Connect two ports from the controller module in the left storage enclosure to the left switch.
- Connect two ports from the controller module in the right storage enclosure to the right switch.
- Connect two ports from the controller modules in each enclosure to the middle switch.

Use multiple switches to avoid a single point of failure inherent to using a single switch, and to physically isolate replication traffic from I/O traffic.

Dual-controller module configuration for replication

Cabling two ME4 Series controller enclosures that are equipped with dual-controller modules for replication.

Multiple servers, one switch, one network

Connecting two ME4 Series 2U storage systems for replication – multiple servers, one switch, and one network on page 92 shows the rear panel of two 2U enclosures with I/O and replication occurring on the same network. Connecting two ME4 Series 5U storage systems for replication – multiple servers, one switch, and one network on page 92 shows the rear panel of two 5U enclosures with I/O and replication occurring on the same network.

In the configuration, Virtual Local Area Network (VLAN) and zoning could be employed to provide separate networks for iSCSI and FC. Create a VLAN or zone for I/O and a VLAN or zone for replication to isolate I/O traffic from replication traffic. Either configuration would be displayed physically as a single network, while logically, either configuration would function as multiple networks.




Figure 35. Connecting two ME4 Series 2U storage systems for replication – multiple servers, one switch, and one network

1. 2U controller enclosures
2. Switch (I/O, replication)
3. Connection to host servers




Figure 36. Connecting two ME4 Series 5U storage systems for replication – multiple servers, one switch, and one network

1. 5U controller enclosures
2. Switch (I/O, replication)
3. Connection to host servers

Multiple servers, multiple switches, and one network

Connecting two ME4 Series 2U storage systems for replication – multiple servers, multiple switches, one network on page 93 shows the rear panel of two 2U enclosures with I/O and replication occurring on the same network. Connecting two ME4 Series 5U storage systems for replication – multiple servers, multiple switches, one network on page 93 shows the rear panel of two 5U enclosures with I/O and replication occurring on the same network.

For optimal protection, use multiple switches for host I/O and replication.

- Connect two ports from each controller module in the left storage enclosure to the left switch.
- Connect two ports from each controller module in the right storage enclosure to the right switch.
- Connect two ports from the controller modules in each enclosure to the middle switch.

Use multiple switches to avoid a single point of failure inherent to using a single switch, and to physically isolate replication traffic from I/O traffic.




Figure 37. Connecting two ME4 Series 2U storage systems for replication – multiple servers, multiple switches, one network

- | | |
|-------------------------------|-------------------------|
| 1. 2U controller enclosures | 2. Two switches (I/O) |
| 3. Connection to host servers | 4. Switch (Replication) |

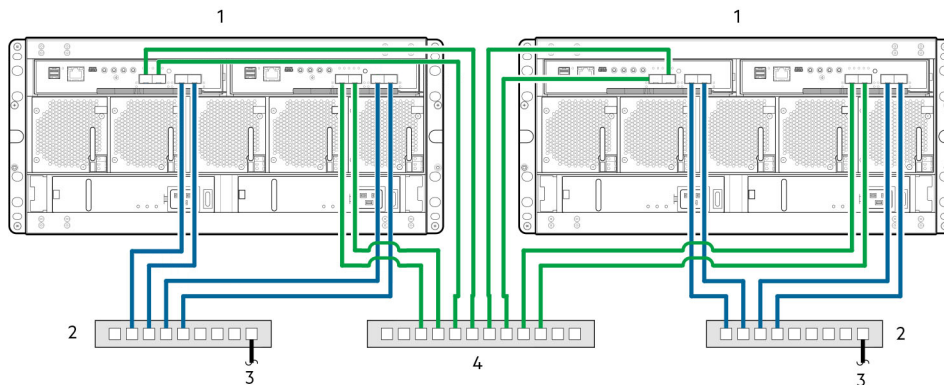


Figure 38. Connecting two ME4 Series 5U storage systems for replication – multiple servers, multiple switches, one network

- | | |
|-------------------------------|-------------------------|
| 1. 5U controller enclosures | 2. Two switches (I/O) |
| 3. Connection to host servers | 4. Switch (Replication) |

Multiple servers, multiple switches, and two networks

Connecting two ME4 Series 2U storage systems for replication– multiple servers, multiple switches, two networks on page 94 shows the rear panel of two 2U enclosures with I/O and replication occurring on different networks. Connecting two ME4 Series 5U storage systems for replication – multiple servers, multiple switches, two networks on page 94 shows the rear panel of two 5U enclosures with I/O and replication occurring on different networks.

- The switch that is on the left supports I/O traffic to local network A.
- The switch that is on the right supports I/O traffic to remote network B.
- The Ethernet WAN in the middle supports replication traffic.

If there is a failure at either the local network or the remote network, you can fail over to the available network.

The following figures represent two branch offices that are cabled for disaster recovery and backup:




Figure 39. Connecting two ME4 Series 2U storage systems for replication– multiple servers, multiple switches, two networks

- | | |
|---|---|
| 1. 2U controller enclosures | 2. Two switches (I/O) |
| 3. Connection to host servers (network A) | 4. Connection to host servers (network B) |
| 5. Ethernet WAN | |




Figure 40. Connecting two ME4 Series 5U storage systems for replication – multiple servers, multiple switches, two networks

- | | |
|---|---|
| 1. 5U controller enclosures | 2. Two switches (I/O) |
| 3. Connection to host servers (network A) | 4. Connection to host servers (network B) |
| 5. Ethernet WAN | |

Isolating replication faults

Replication is a disaster-recovery feature that performs asynchronous replication of block-level data from a volume in a primary storage system to a volume in a secondary storage system.

The replication feature creates an internal snapshot of the primary volume, and copies changes to the data since the last replication to the secondary system using iSCSI or FC connections. The primary volume exists in a primary pool in the primary storage system. Replication can be completed using either the PowerVault Manager or the CLI.

Replication setup and verification

After storage systems are cabled for replication, you can use the PowerVault Manager to prepare for using the replication feature. Alternatively, you can use SSH or telnet to access the IP address of the controller module and access the replication feature using the CLI.

Basic information for enabling the ME4 Series Storage System controller enclosures for replication supplements the troubleshooting procedures that follow.

- Familiarize yourself with replication content provided in the *Dell EMC PowerVault ME4 Series Series Storage System Administrator's Guide*.
- For virtual replication, perform the following steps to replicate an existing volume to a pool on the peer in the primary system or secondary system:

1. Find the port address on the secondary system:

Using the CLI, run the `show ports` command on the secondary system.

2. Verify that ports on the secondary system can be reached from the primary system using either of the following methods:
 - Run the `query peer-connection` CLI command on the primary system, using a port address obtained from the output of the `show ports` command.
 - In the PowerVault Manager **Replications** topic, select **Action > Query Peer Connection**.

3. Create a peer connection.

To create a peer connection, use the `create peer-connection` CLI command or in the PowerVault Manager **Replications** topic, select **Action > Create Peer Connection**.

4. Create a virtual replication set.

To create a replication set, use the `create replication-set` CLI command or in the PowerVault Manager **Replications** topic, select **Action > Create Replication Set**.

5. Replicate.

To initiate replication, use the `replicate` CLI command or in the PowerVault Manager **Replications** topic, select **Action > Replicate**.

- Using the PowerVault Manager, monitor the storage system event logs for information about enclosure-related events, and to determine any necessary recommended actions

NOTE: These steps are a general outline of the replication setup. Refer to the following manuals for more information about replication setup:

- See the *Dell EMC PowerVault ME4 Series Series Storage System Administrator's Guide* for procedures to set up and manage replications.
- See the *Dell EMC PowerVault ME4 Series Series Storage System CLI Guide* for replication commands and syntax.

NOTE: Controller module firmware must be compatible on all systems that are used for replication.

Diagnostic steps for replication setup

The tables in the following section show menu navigation for virtual replication using the PowerVault Manager:

NOTE: SAS controller enclosures do not support replication.

Can you successfully use the replication feature?

Table 28. Diagnostics for replication setup: Using the replication feature

Answer	Possible reasons	Action
Yes	System functioning properly	No action required.
No	Compatible firmware revision supporting the replication feature is not running on each system that is used for replication.	Perform the following actions on each system used for virtual replication: <ul style="list-style-type: none"> • On the System topic, select Action > Update Firmware. The Update Firmware panel opens. The Update Controller Modules tab shows firmware versions that are installed in each controller. • If necessary, update the controller module firmware to ensure compatibility with the other systems. • See the topic about updating firmware in the <i>Dell EMC PowerVault ME4 Series Storage System Administrator's Guide</i> for more information about compatible firmware.
No	Invalid cabling connection. (If multiple enclosures are used, check the cabling for each system.)	Verify controller enclosure cabling: <ul style="list-style-type: none"> • Verify use of proper cables. • Verify proper cabling paths for host connections.

Table 28. Diagnostics for replication setup: Using the replication feature (continued)

Answer	Possible reasons	Action
		<ul style="list-style-type: none"> • Verify cabling paths between replication ports and switches are visible to one another. • Verify that cable connections are securely fastened. • Inspect cables for damage and replace if necessary.
No	A system does not have a pool that is configured.	Configure each system to have a storage pool.

Can you create a replication set?

After valid cabling, and network availability, create a replication set by selecting **Action > Create Replication Set** from the **Replications** topic.

Table 29. Diagnostics for replication setup – Creating a replication set

Answer	Possible reasons	Action
Yes	System functioning properly.	No action required.
No	On controller enclosures equipped with iSCSI host interface ports, replication set creation fails due to use of CHAP.	If using , see the topics about configuring CHAP and working in replications within the <i>Dell EMC PowerVault ME4 Series Storage System Administrator's Guide</i> .
No	Unable to create the secondary volume (the destination volume on the pool to which you replicate data from the primary volume).	<ul style="list-style-type: none"> • Review event logs for indicators of a specific fault in a replication data path component. Follow any Recommended Actions. • Verify valid specification of the secondary volume according to either of the following criteria: <ul style="list-style-type: none"> ○ A conflicting volume does not exist. ○ Available free space in the pool.
No	Communication link is down.	Review event logs for indicators of a specific fault in a host or replication data path component.

Can you replicate a volume?

Table 30. Diagnostics for replication setup – Replicating a volume

Answer	Possible reasons	Action
Yes	System functioning properly.	No action required.
No	Nonexistent .	<ul style="list-style-type: none"> • Determine existence of primary or secondary volumes. • If a replication set has not been successfully created, Action > Create Replication Set on the Replications topic to create a replication. • Review event logs (in the footer, click the events panel and select Show Event List) for indicators of a specific fault in a replication data path component. Follow any Recommended Actions.
No	Network error occurred during in-progress replication.	<ul style="list-style-type: none"> • Review event logs for indicators of a specific fault in a replication data path component. Follow any Recommended Actions. • Click in the Volumes topic, and then click a volume name in the volumes list. Click the Replication Sets tab to display replications and associated metadata. • Replications that enter the suspended state can be resumed manually (see the <i>Dell EMC PowerVault ME4 Series Storage System Administrator's Guide</i> for additional information).

Table 30. Diagnostics for replication setup – Replicating a volume (continued)

Answer	Possible reasons	Action
No	Communication link is down.	Review event logs for indicators of a specific fault in a host or replication data path component.

Has a replication run successfully?

Table 31. Diagnostics for replication setup: Checking for a successful replication

Answer	Possible reasons	Action
Yes	System functioning properly	No action required.
No	Last Successful Run shows N/A	<ul style="list-style-type: none">• In the Volumes topic, click the volume that is a member of the replication set.<ul style="list-style-type: none">○ Select the Replication Sets table.○ Check the Last Successful Run information.• If the replication has not run successfully, use the PowerVault Manager to replicate as described in the topic about working in replications in the <i>Dell EMC PowerVault ME4 Series Storage System Administrator's Guide</i>.
No	Communication link is down	Review event logs for indicators of a specific fault in a host or replication data path component.

SFP+ transceiver for FC/iSCSI ports

This section describes how to install the small form-factor pluggable (SFP+) transceivers ordered with the ME4 Series FC/iSCSI controller module.

Locate the SFP+ transceivers

Locate the SFP+ transceivers that shipped with the controller enclosure, which look similar to the generic SFP+ transceiver that is shown in the following figure:




Figure 41. Install an SFP+ transceiver into the ME4 Series FC/iSCSI controller module

- | | |
|-------------------------------------|----------------------|
| 1. CNC-based controller module face | 2. CNC port |
| 3. SFP+ transceiver (aligned) | 4. Fiber-optic cable |
| 5. SFP+ transceiver (installed) | |

NOTE: Refer to the label on the SFP+ transceiver to determine whether it supports the FC or iSCSI protocol.

Install an SFP+ transceiver

Perform the following steps to install an SFP+ transceiver:

NOTE: Follow the guidelines provided in [Electrical safety](#) on page 8 when installing an SFP+ transceiver.

1. Orient the SFP+ transceiver with the port and align it for insertion.

For 2U controller enclosures, the transceiver is installed either right-side up, or upside down depending upon whether it is installed into controller module A or B,

2. If the SFP+ transceiver has a plug, remove it before installing the transceiver. Retain the plug.
3. Flip the actuator open.


NOTE: The actuator on your SFP+ transceiver may look slightly different than the one shown in [Install an SFP+ transceiver into the ME4 Series FC/iSCSI controller module](#) on page 98

4. Slide the SFP+ transceiver into the port until it locks securely into place.
5. Flip the actuator closed.
6. Connect a qualified fiber-optic interface cable into the duplex jack of the SFP+ transceiver.

If you do not plan to use the SFP+ transceiver immediately, reinsert the plug into the duplex jack of SFP+ transceiver to keep its optics free of dust.

Verify component operation

View the port Link Status/Link Activity LED on the controller module face plate. A green LED indicates that the port is connected and the link is up.

 **NOTE:** To remove an SFP+ transceiver, perform the installation steps in reverse order relative to what is described in [Install an SFP+ transceiver](#) on page 98.

System Information Worksheet

Use the system information worksheet to record the information that is needed to install the ME4 Series Storage System.

ME4 Series Storage System information

Gather and record the following information about the ME4 Series storage system network and the administrator user:

Table 32. ME4 Series Storage System network

Item	Information
Service tag	
Management IPv4 address (ME4 Series Storage System management address)	----- . ----- . ----- . -----
Top controller module IPv4 address (Controller A MGMT port)	----- . ----- . ----- . -----
Bottom controller module IPv4 address (Controller B MGMT port)	----- . ----- . ----- . -----
Subnet mask	----- . ----- . ----- . -----
Gateway IPv4 address	----- . ----- . ----- . -----
Gateway IPv6 address	----- : ----- : ----- : ----- ::-----
Domain name	
DNS server address	----- . ----- . ----- . -----
Secondary DNS server address	----- . ----- . ----- . -----

Table 33. ME4 Series Storage System administrator

Item	Information
Password for the default ME4 Series Storage System Admin user	
Email address of the default ME4 Series Storage System Admin user	

iSCSI network information

For a storage system with iSCSI front-end ports, plan and record network information for the iSCSI network.

i **NOTE:** For a storage system deployed with two Ethernet switches, Dell EMC recommends setting up separate subnets.

Table 34. iSCSI Subnet 1

Item	Information
Subnet mask	----- . ----- . ----- . -----

Table 34. iSCSI Subnet 1 (continued)

Item	Information
Gateway IPv4 address	----- . ----- . ----- . -----
IPv4 address for storage controller module A: port 0	----- . ----- . ----- . -----
IPv4 address for storage controller module B: port 0	----- . ----- . ----- . -----
IPv4 address for storage controller module A: port 2	----- . ----- . ----- . -----
IPv4 address for storage controller module B: port 2	----- . ----- . ----- . -----

Table 35. iSCSI Subnet 2

Item	Information
Subnet mask	----- . ----- . ----- . -----
Gateway IPv4 address	----- . ----- . ----- . -----
IPv4 address for storage controller module A: port 1	----- . ----- . ----- . -----
IPv4 address for storage controller module B: port 1	----- . ----- . ----- . -----
IPv4 address for storage controller module A: port 3	----- . ----- . ----- . -----
IPv4 address for storage controller module B: port 3	----- . ----- . ----- . -----
Gateway IPv6 address	----- :----- :----- :-----:-----

Additional ME4 Series Storage System information

The Network Time Protocol (NTP) and Simple Mail Transfer Protocol (SMTP) server information is optional. The proxy server information is also optional, but it may be required to complete the Discover and Configure Uninitialized text TBD wizard.

Table 36. NTP, SMTP, and Proxy servers

Item	Information
NTP server IPv4 address	----- . ----- . ----- . -----
SMTP server IPv4 address	----- . ----- . ----- . -----
Backup NTP server IPv4 address	----- . ----- . ----- . -----
SMTP server login ID	
SMTP server password	
Proxy server IPv4 address	----- . ----- . ----- . -----

Fibre Channel zoning information

For a storage system with Fibre Channel front-end ports, record the physical and virtual WWNs of the Fibre Channel ports in fabric 1 and fabric 2. This information is displayed on the Review Front-End page of the Discover and Configure Uninitialized wizard. Use this information to configure zoning on each Fibre Channel switch.

Table 37. WWNs in fabric 1

Item	FC switch port	Information
WWN of storage controller A: port 0		
WWN of storage controller B: port 0		
WWN of storage controller A: port 2		
WWN of storage controller B: port 2		
WWNs of server HBAs:		

Table 38. WWNs in fabric 2

Item	FC switch port	Information
WWN of storage controller A: port 1		
WWN of storage controller B: port 1		
WWN of storage controller A: port 3		
WWN of storage controller B: port 3		

Setting network port IP addresses using the CLI port and serial cable

You can manually set the static IP addresses for each controller module. Alternatively, you can specify that IP addresses should be set automatically for both controllers through communication with a Dynamic Host Configuration Protocol (DHCP) server.

In DHCP mode, the network port IP address, subnet mask, and gateway are obtained from a DHCP server. If a DHCP server is not available, the current network addresses are not changed. To determine the addresses that are assigned to the controller modules, use the list of bindings on the DHCP server.

If you did not use DHCP to set network port IP address, you can set them manually using the CLI port and serial cable. You can connect to the controller module using the 3.5mm stereo plug CLI port and the supplied 3.5mm/DB9 serial cable. Alternatively, you can use a generic mini-USB cable (not included) and the USB CLI port. If you plan on using a mini-USB cable, you must enable the USB CLI port for communication.

Network ports on controller module A and controller module B are configured with the following default values:

- **Network port IP address:** 10.0.0.2 (controller A), 10.0.0.3 (controller B)
- **IP subnet mask:** 255.255.255.0
- **Gateway IP address :** 10.0.0.1

If the default IP addresses are not compatible with your network, you must set an IP address for each network port using the CLI.

i NOTE: If you are using the mini-USB CLI port and cable, see [Mini-USB Device Connection](#) on page 106.

- If you are using a host computer running Windows, download and install the USB device driver for the CLI port as described in [Obtaining the USB driver](#) on page 106. Skip this task if you are using a host computer running Windows 10 or Windows Server 2016 and later.
- If you are using a host computer running Linux, prepare the USB port as described in [Linux drivers](#) on page 107.

Use the CLI commands described in the following steps to set the IP address for the network port on each controller module:

i NOTE: When new IP addresses are set, you can change them as needed using the PowerVault Manager. Be sure to change the IP address before changing the network configuration.

1. From your network administrator, obtain an IP address, subnet mask, and gateway address for controller A and another for controller B.
2. Connect the provided 3.5mm/DB9 serial cable from a host computer with a serial port to the 3.5mm stereo plug CLI port on controller A. Alternatively, connect a generic mini-USB cable from a host computer to the USB CLI port on controller A. The mini-USB connector plugs into the USB CLI port as shown in the following figure:




Figure 42. Connecting a USB cable to the CLI port

3. Start a terminal emulator and configure it to use the display settings in [Terminal emulator display settings](#) on page 104 and the connection settings in [Terminal emulator connection settings](#) on page 104.

Table 39. Terminal emulator display settings

Parameter	Value
Terminal emulation mode	VT-100 or ANSI (for color support)
Font	Terminal
Translations	None
Columns	80

Table 40. Terminal emulator connection settings

Parameter	Value
Connector	COM3 (for example) ^{1,2}
Baud rate	115,200
Data bits	8
Parity	None
Stop bits	1
Flow control	None

¹ Your host computer configuration determines which COM port is used for the Disk Array USB Port.

² Verify the appropriate COM port for use with the CLI.


4. Press Enter to display login prompt if necessary.
The CLI displays the system version, Management Controller version, and login prompt.
5. If you are connecting to a storage system with G275 firmware that has not been deployed:
 - a. Type `manage` at the login prompt and press Enter.
 - b. Type `!manage` at the Password prompt and press Enter.
 If you are connecting to a storage system with G275 firmware that has been deployed:
 - a. Type the username of a user with the manage role at the login prompt and press Enter.
 - b. Type the password for the user at the Password prompt and press Enter.
6. If you are connecting to a storage system with G280 firmware that has not been deployed:
 - a. Type `setup` at the login prompt and press Enter.
 - b. Do not type anything at the Password prompt and press Enter.

If you are connecting to a storage system with G280 firmware that has been deployed:

- a. Type the username of a user with the manage role at the login prompt and press Enter.
 - b. Type the password for the user at the Password prompt and press Enter.
7. To use DHCP to set network port IP addresses, type the following command at the prompt:

```
set network-parameters dhcp
```

To use custom static IP addresses, type the following CLI command to set the values you obtained in step 1:

 **NOTE:** Run the command for controller module A first, and then run the command for controller module B.

```
set network-parameters ip address netmask netmask gateway gateway controller a|b
```

where:

- *address* is the IP address of the controller module
- *netmask* is the subnet mask
- *gateway* is the IP address of the subnet router
- *a|b* specifies the controller whose network parameters you are setting

For example:

```
set network-parameters ip 192.168.0.10 netmask 255.255.255.0 gateway 192.168.0.1
controller a
set network-parameters ip 192.168.0.11 netmask 255.255.255.0 gateway 192.168.0.1
controller b
```

8. Type the following CLI command to verify the new IP addresses:

```
show network-parameters
```

The network parameters, including the IP address, subnet mask, and gateway address are displayed for each controller module.

9. Use the CLI ping command to verify connectivity to the gateway address.


For example:

```
ping 192.168.0.1
```

10. Open a command window on the host computer and type the following command to verify connectivity to controller A and controller B:


```
ping controller-IP-address
```

If you cannot access your storage system for at least three minutes after changing the IP address, restart the controllers using the CLI.

 **NOTE:** When you restart a Management Controller, communication with it is temporarily lost until it successfully restarts.

Type the following CLI command to restart the Management Controller in both controllers:

```
restart mc both
```

 **CAUTION:** When configuring an iSCSI storage system or a storage system that uses a combination of Fibre Channel and iSCSI SFPs, do not restart the Management Controller or exit the terminal emulator session until the CNC ports are configured as described [Changing host port settings](#) on page 38

11. Record the IP address for the controller modules to use when connecting to the storage system using PowerVault Manager.
12. When you are done using the CLI, close the terminal emulator.

Topics:

- [Mini-USB Device Connection](#)

Mini-USB Device Connection

The following sections describe the connection to the mini-USB port:

Emulated serial port

When a computer is connected to a controller module using a mini-USB serial cable, the controller presents an emulated serial port to the computer. The name of the emulated serial port is displayed using a *customer vendor ID* and *product ID*. Serial port configuration is unnecessary.

NOTE: Certain operating systems require a device driver or special mode of operation to enable proper functioning of the USB CLI port. See also [Device driver/special operation mode](#) on page 106.

Supported host applications

The following terminal emulator applications can be used to communicate with an ME4 Series controller module:

Table 41. Supported terminal emulator applications

Application	Operating system
PuTTY	Microsoft Windows (all versions)
Minicom	Linux (all versions)

Command-line interface

When the computer detects a connection to the emulated serial port, the controller awaits input of characters from the computer using the command-line interface. To see the CLI prompt, you must press Enter.

NOTE: Directly cabling to the mini-USB port is considered an out-of-band connection. The connection to the mini-USB port is outside of the normal data paths to the controller enclosure.

Device driver/special operation mode

Certain operating systems require a device driver or special mode of operation. The following table displays the product and vendor identification information that is required for certain operating systems:

Table 42. USB identification code

USB identification code type	Code
USB Vendor ID	0x210c
USB Product ID	0xa4a7

Microsoft Windows drivers

Dell EMC provides an ME4 Series USB driver for use in Windows environments.

Obtaining the USB driver

NOTE: If you are using Windows 10 or Windows Server 2016, the operating system provides a native USB serial driver that supports the mini-USB port. However, if you are using an older version of Windows, you should download and install the USB driver.

1. Go to Dell.com/support and search for ME4 Series USB driver.

2. Download the ME4 Series Storage Array USB Utility file from the Dell EMC support site.
3. Follow the instructions on the download page to install the ME4 Series USB driver.

Known issues with the CLI port and mini-USB cable on Microsoft Windows

When using the CLI port and cable for setting network port IP addresses, be aware of the following known issue on Windows:

Problem

The computer might encounter issues that prevent the terminal emulator software from reconnecting after the controller module restarts or the USB cable is unplugged and reconnected.

Workaround

To restore a connection that stopped responding when the controller module was restarted:

1. If the connection to the mini-USB port stops responding, disconnect and quit the terminal emulator program.
 - a. Using Device Manager, locate the COM n port that is assigned to the mini-USB port.
 - b. Right-click on the **Disk Array USB Port (COM n)** port, and select **Disable device**.
2. Right-click on the **Disk Array USB Port (COM n)** port, and select **Enable device**.
3. Start the terminal emulator software and connect to the COM port.


 **NOTE:** On Windows 10 or Windows Server 2016, the XON/XOFF setting in the terminal emulator software must be disabled to use the COM port.

Linux drivers

Linux operating systems do not require the installation of an ME4 Series USB driver. However, certain parameters must be provided during driver loading to enable recognition of the mini-USB port on an ME4 Series controller module.

- Type the following command to load the Linux device driver with the parameters that are required to recognize the mini-USB port:

```
# modprobe usbserial vendor=0x210c product=0xa4a7 use_acm=1
```

 **NOTE:** Optionally, this information can be incorporated into the `/etc/modules.conf` file.