

Dell PowerVault ME5 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
ME5 Series system requirements.....	6
Web browser requirements.....	6
SupportAssist and APEX AIOps Infrastructure Observability requirements.....	7
Secure Connect Gateway.....	9
Unpack the enclosure.....	9
Unpacking a 2U Enclosure.....	9
Unpacking a 5U enclosure.....	10
Safety guidelines.....	11
Safe handling.....	11
Safe operation.....	12
Electrical safety.....	12
Rack system safety precautions.....	13
Installation checklist.....	14
Planning for installation.....	14
Preparing for installation.....	15
Preparing the site and host server.....	15
Required tools.....	16
Requirements for rackmount installation.....	16
Disk drive module.....	16
Drive carrier module in 2U chassis.....	16
Drive status indicators.....	17
Blank drive carrier modules.....	17
DDIC in a 5U enclosure.....	17
Populating drawers with DDICs.....	18
 Chapter 2: Mount the enclosures in the rack.....	 20
Rackmount rail kit.....	20
Install the 2U enclosure using toolless rails.....	20
Install the 2U enclosure front bezel.....	23
Install the 5U84 enclosure.....	24
Connect optional expansion enclosures.....	25
Cable requirements for expansion enclosures.....	25
 Chapter 3: Connect to the management network.....	 28
 Chapter 4: Cable host servers to the storage system.....	 29
Cabling considerations.....	29
Connecting the enclosure to hosts.....	29
Fibre Channel protocol.....	29
iSCSI protocol.....	30
SAS protocol.....	32
Host connection.....	32
32 Gb Fibre Channel host connection.....	32

25 GbE iSCSI host connection.....	32
10GBase-T host connection.....	32
12 Gb HD mini-SAS host connection.....	33
Connecting direct attach configurations.....	33
Single-controller module configurations.....	33
Dual-controller module configurations.....	33
Chapter 5: Connect power cables and power on the storage system.....	39
Power cable connection.....	39
Chapter 6: Perform system and storage setup.....	41
Prerequisites.....	41
Record storage system information.....	41
About guided setup.....	41
Access the PowerVault Manager.....	42
System configuration.....	42
Configuring controller network ports.....	42
Set the date and time.....	44
Set up users.....	45
Notifications.....	46
Configure iSCSI ports.....	47
Set up SupportAssist and APEX AIOps Infrastructure Observability.....	48
Storage configuration.....	48
Set up virtual storage.....	48
Set up linear storage.....	49
Provisioning.....	50
Set up hosts.....	50
Set up volumes.....	50
Chapter 7: Setting up hosts.....	52
Host system requirements.....	52
Windows hosts.....	53
Configuring a Windows host with FC HBAs.....	53
Configuring a Windows host with iSCSI network adapters.....	55
Configuring a Windows host with SAS HBAs.....	58
Linux hosts.....	60
Configuring a Linux host with FC HBAs	60
Configuring a Linux host with iSCSI network adapters.....	62
Configuring a SAS host server for Linux.....	66
VMware ESXi hosts.....	68
Configuring a Fibre Channel host server for VMware ESXi.....	68
Configuring an ESXi host with an iSCSI network adapter.....	70
Configuring a SAS host server for VMware ESXi.....	73
Citrix XenServer hosts.....	75
Configuring a Fibre Channel host server for Citrix XenServer.....	75
Configuring an iSCSI host server for Citrix XenServer.....	77
Configuring a SAS host for Citrix XenServer.....	80
Chapter 8: Troubleshooting and problem solving.....	83

Fault isolation methodology.....	83
Options available for performing basic steps.....	83
Performing basic steps.....	84
Host I/O.....	85
2U enclosure LEDs.....	85
2U enclosure Ops panel.....	85
2U enclosure PCM LEDs.....	86
2U enclosure Ops panel LEDs.....	86
2U enclosure disk drive carrier module LEDs.....	87
IO Module LEDs.....	88
12 Gb/s controller module LEDs.....	88
5U84 enclosure LEDs.....	91
5U enclosure Ops panel.....	91
ME5084 PSU LEDs.....	91
ME5084 FCM LEDs.....	92
ME5084 Ops panel LEDs.....	92
ME5084 drawer LEDs.....	92
ME5084 DDIC LEDs.....	93
5U84 controller module and IOM LEDs	94
Initial start-up problems.....	94
Troubleshooting 2U enclosures.....	94
Troubleshooting 5U enclosures.....	94
If the enclosure does not initialize.....	95
Correcting enclosure IDs.....	95
Troubleshooting hardware faults.....	96
Appendix A: Cabling for replication.....	99
Connecting two storage systems to replicate volumes.....	99
Example cabling for replication.....	99
Single-controller module configuration for replication.....	100
Dual-controller module configuration for replication.....	100
Isolating replication faults.....	102
Diagnostic steps for replication setup.....	103
Appendix B: SFP transceiver for FC/iSCSI ports.....	106
Appendix C: System Information Worksheet.....	107
Appendix D: Setting network port IP addresses using the CLI port.....	110
Set a network port IP address using the micro USB port.....	110
Micro-USB device connection.....	112
Microsoft Windows drivers.....	113
Linux drivers.....	113
Appendix E: Technical specifications.....	114

Before you begin

This document describes the initial hardware setup for Dell PowerVault ME5 Series storage systems.

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

Topics:

- [ME5 Series system requirements](#)
- [Unpack the enclosure](#)
- [Safety guidelines](#)
- [Installation checklist](#)
- [Planning for installation](#)
- [Preparing for installation](#)
- [Disk drive module](#)
- [Populating drawers with DDICs](#)

ME5 Series system requirements

The following sections list browser and network requirements for the ME5 Series system.

Web browser requirements

PowerVault Manager supports the browsers that are listed below.

- Apple Safari 11 and newer (Mac)
- Google Chrome 70 and newer
- Microsoft Internet Explorer 11
- Mozilla Firefox 68 and newer

For best results, follow these guidelines:

- The recommended resolution for the page display area in the browser is 1360 x 768 pixels.
- To optimize the display, use a color monitor and set its color quality to the highest setting.
- To navigate beyond the sign-in page (with a valid user account):
 - If the PowerVault Manager is configured to use HTTPS, ensure that your browser is set to use TLS 1.2.
 - Verify that the browser is set to allow cookies, at least for the IP addresses of the storage system network ports.
 - For Internet Explorer, set the local-intranet security option in the browser to medium or medium-low.
 - For Internet Explorer, add the network IP address of each controller module as a trusted site.
- To see the help window, enable pop-up windows.

i NOTE: By default, your system is loaded with self-signed certificates. Generate new self-signed certificates on each controller using the `create_certificate` CLI command. Browser messages warning you about security or privacy concerns due to self-signed or untrusted certificates or invalid certificate authorities are expected, and warnings can be safely bypassed if you trust that you are contacting the correct controller within your network. Depending on the browser and its settings, once you navigate through the browser warning, a security exception may be created and, the warning may not appear. Your browser address bar still indicates that the connection is not trusted or not secure. You can safely ignore the indication if you trust you are accessing the correct controller within your network.

SupportAssist and APEX AI Ops Infrastructure Observability requirements

SupportAssist enhances your support experience by sending configuration and diagnostic information to technical support at regular intervals. APEX AI Ops Infrastructure Observability 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.

- To use APEX AI Ops Infrastructure Observability, you need a ProSupport contract. See the [APEX AI Ops Infrastructure Observability](#) product page for more information.
- To connect with SupportAssist, you need to do the following in advance:
 - Create a business account on dell.com
 - Create a PIN and generate an access key
 - Set up a direct connection to SupportAssist

The following procedures describe how to meet these SupportAssist requirements.

Creating a business account

Administrators who manage Dell storage equipment must create a Dell business account and an access key to use SupportAssist with the ME5 Series storage system.

Steps

1. Go to the [Dell account registration page](#).
2. In the **Create an Account** section, enter information in all the required fields to create a business account. You must use your business email address.
3. Click **Create Account**.
4. On the **Additional Access Rights Required** page, select **I'm an employee of an organization that has purchased Dell EMC enterprise products** and then click **Submit**.

 **NOTE:** If you are an employee of a Dell partner company or want to become a partner, choose that option and follow the instructions that are provided.

The **Business account registration** page opens.

5. In the select your relationship page, choose **Yes, my company has already purchased Dell EMC infrastructure solutions** and then click **Next**.
The **Submit your organization information** page opens.
6. Enter the business name and other information for your organization and click **Submit**.
Your request for an account is sent to Dell. When your account has been validated, Dell sends a confirmation email with information about accessing your account to complete your registration. This confirmation may take up to 24 hours.
7. Access your account and enter the validation code that Dell support provided in the **Validation Code** field.
When you are logged in to your account, a checkmark appears next to your username in the menu bar. The color indicates the account status:
 - Black—the account is a validated business account.
 - Green—registration is not yet complete. Follow the instructions in your confirmation email from Dell to complete registration.
 - Yellow—the account is not a validated business account. If you think this status is in error, contact Dell technical support.

When the business account is validated and a black checkmark appears next to your username, you may generate an access key as described in the next section.

Generating an access key and PIN

To support the increased focus on security in corporate environments worldwide, and in order to prevent device spoofing Dell has introduced a two-factor authentication method to permit devices to connect to Dell secure remote servers. Administrators must generate an access key and PIN to facilitate this key exchange. This section describes how to generate a PIN and an access key.

Prerequisites

- Dell business account
- ME5 product Service Tag number

Steps

1. Go to the [Dell account registration page](#).

2. In the **Sign in** section, enter you Email and Password.

A checkmark appears next to your username in the menu bar. Ensure that the checkmark is black, indicating that your account is valid. When the business account is validated and a black checkmark appears next to your username, you may generate an access key.

3. In the **Identify your product** search box, enter your Service Tag or your product model.

 **NOTE:** If you receive the message *We could not find the site you were searching for, please verify if you have access to the site*, contact [Dell customer support](#).

4. On the **Overview** tab for your product, under **Quick links** click **Generate Access key**.

5. Select your product from the **Select a Product ID or Service tag** list.

6. Enter a four-digit PIN in the **Create PIN** box.
Record the PIN for future use.

7. Click **Generate Access Key** and then click **Done** to close the window.

The Dell Services Connectivity Team <ServicesConnectivity_Noreply@dell.com> sends you an email that contains your Access Key details. Retain this information for future use.

 **NOTE:** The access key is valid for seven days. If the key has expired and you must configure SupportAssist, use PowerVault Manager (**Maintenance > Support**) or follow these steps to generate a new access key.

Related links

- [KB 000198043 How to generate an access key](#)
- [KB 000197465 Unable to generate access key and PIN](#)

SupportAssist direct connection requirements

The following network requirements must be met to use a direct connection for SupportAssist.

- A minimum of one DNS server must be configured.
- The local system must connect to the following destinations to ensure connectivity to the Global servers:
 - esrs3-core.emc.com
 - esrs3-coredr.emc.com

Use the following command to verify the connections:

```
# check support-assist-connection mode direct
```

If the connection is working, it returns an HTTP Status 200 message:

```
SupportAssist Connection
-----
Connection State: Connected

Endpoint
-----
Mode: direct
Endpoint: https://esrs3-core.emc.com
Status: success
HTTP Status: 200
Message: OK
Status Detail: Success
Proxy Type: none

Endpoint
-----
Mode: direct
Endpoint: https://esrs3-coredr.emc.com
Status: success
```

```
HTTP Status: 200
Message: OK
Status Detail: Success
Proxy Type: none
Success: Command completed successfully. (2022-01-08 18:04:00)
```

To ensure connection integrity, proxy servers and devices external to your demilitarized zone (DMZ) must not perform any method of SSL decryption on outbound or inbound traffic for the Dell secure remote servers. SSL decryption that is performed on outbound communication causes a loss of connectivity to the backend. SSL decryption includes decryption by your firewall, proxies, web traffic filtering appliances or cloud services, web traffic shaping or load balancing, certificate verification, certificate proxy, or Intrusion Detection Services (IDS).

In case the SSL decryption is enabled on the proxy servers and other devices, ensure that `esrs3-core.emc.com` and `esrs3-coredr.emc.com` are added to the SSL decryption exclusion list on the proxy servers and devices.

Table 1. Port requirements

Type of connection	Ports that must be open	Protocol used	Communication
Connect Directly	443	TCP	Outbound
Connect via Gateway Server	9443	TCP	Outbound

Secure Connect Gateway

ME5 Series storage systems support a Secure Connect Gateway (SCG) virtual appliance for SupportAssist.

The default connection to the SCG virtual appliance is `https://<SCG.IP.address>:9443`.

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.

Unpacking a 2U Enclosure

About this task

- 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 FC or iSCSi controller modules that require SFP transceivers, see [SFP transceiver for FC/iSCSi ports](#).

Steps

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 the following figure. 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 |

Unpacking a 5U enclosure

About this task

- 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](#).
- For enclosures configured with FC or iSCSI controller modules that require SFP transceivers, see [SFP transceiver for FC/iSCSI ports](#).

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.

Steps

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 the following figure. 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 |

Safety guidelines

Always follow these safety guidelines to avoid injury and damage to ME5 Series components.

If you use this equipment in a manner that is not specified by Dell, 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 PowerVault ME5 Series Storage System Getting Started Guide* for product safety and regulatory information. Warranty information is included as a separate document.

Safe handling

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

CAUTION: Use this equipment in a manner specified by Dell. 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.
- 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 that is attached to the storage enclosure.**

Safe operation

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

- 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 or replacement of a PCM or PSU, disconnect supply power from the module to be replaced. See the *Dell PowerVault ME5 Series Storage System Owner's Manual*.
 - Follow the instructions in the hazardous voltage warning label attached to the 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 attached 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 cable must have a safe electrical grounding connection. Check the grounding connection of the enclosure before you turn on the power supply.
- 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 attached to the 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.**

ESD precautions

Before you begin any of the procedures, review the following precautions and preventive measures.

Preventing electrostatic discharge

To prevent electrostatic discharge (ESD) from damaging the system, be aware of the precautions to consider when setting up the system or handling parts. A discharge of static electricity from a finger or other conductor may damage system boards or other static-sensitive devices. This type of damage may reduce the life expectancy of the device.

 **CAUTION:** Parts can be damaged by electrostatic discharge. Follow these precautions:

- **Avoid hand contact by transporting and storing products in static-safe containers.**
- **Keep electrostatic-sensitive parts in their containers until they arrive at static-protected workstations.**
- **Place parts in a static-protected area before removing them from their containers.**
- **Avoid touching pins, leads, or circuitry.**
- **Always be properly grounded when touching a static-sensitive component or assembly.**
- **Remove clutter (plastic, vinyl, foam) from the static-protected workstation.**

Grounding methods to prevent electrostatic discharge

Several methods are used for grounding. Adhere to the following precautions when handling or installing electrostatic-sensitive parts.

 **CAUTION:** Parts can be damaged by electrostatic discharge. Use proper anti-static protection:

- **Keep the replacement CRU in the ESD bag until needed; and when removing a CRU from the enclosure, immediately place it in the ESD bag and anti-static packaging.**
- **Wear an ESD wrist strap connected by a ground cord to a grounded workstation or unpainted surface of the computer chassis. Wrist straps are flexible straps with a minimum of 1 megohm (± 10 percent) resistance in the ground cords. To provide proper ground, wear the strap snug against the skin.**
- **If an ESD wrist strap is unavailable, touch an unpainted surface of the chassis before handling the component.**
- **Use heel straps, toe straps, or boot straps at standing workstations. Wear the straps on both feet when standing on conductive floors or dissipating floor mats.**
- **Use conductive field service tools.**
- **Use a portable field service kit with a folding static-dissipating work mat.**

If you do not have any of the suggested equipment for proper grounding, have an authorized technician install the part. For more information about static electricity or assistance with product installation, contact customer support. For additional information, see www.dell.com/support.

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 2. Installation checklist

Step	Task	Where to find procedure
1	Unpack the enclosure.	Unpack the enclosure
2	Install the controller enclosure and optional expansion enclosures in the rack. ¹	• Required tools • Requirements for rackmount installation • Install the 2U enclosure • Install the 5U84 enclosure
3	Populate drawers with disks (DDICs) in 5U84 enclosure; 2U enclosures ship with disks installed.	Populating drawers with DDICs
4	Cable the optional expansion enclosures.	Connect optional expansion enclosures
5	Connect the management ports.	Connect to the management network
6	Cable the controller host ports. ²	Connecting the enclosure to hosts
7	Connect the power cords and power on the system.	Power cable connection
8	Perform system and storage setup.	Accessing the storage manager
9	Perform host setup: • Attach the host servers. • Install the required host software.	• Host system requirements • Windows hosts • Linux hosts • VMware ESXi hosts • Citrix XenServer hosts
10	Perform the initial configuration tasks. ³	Using guided setup

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 PowerVault ME5 Series Storage System Administrator's Guide*.

³ The PowerVault Manager is introduced in [Accessing the storage manager](#). See the *Dell PowerVault ME5 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 3. 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 .
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	- 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 200–240 VAC for operation. See the [Technical specifications](#) 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](#) 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 ME5 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 482 mm (19 in.) cabinet capable of holding components that occupy 2U or 5U of rack space.

- Minimum depth: 707 mm (27.83 in.) from rack posts to maximum extremity of enclosure (includes back panel cabling and cable bend radii).
- Weight:
 - 2U enclosures—weigh up to 32 kg (71 lb) depending on the configuration.
 - 5U enclosures—weigh up to 128 kg (282 lb) depending on the configuration.
- 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 ME5 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 3.5 in. disk drive in its carrier. The disk drives are installed horizontally in the chassis. A carrier adapter is available to accommodate 2.5 in. disk drives.
- Each 2U24 drive slot holds a single low profile 2.5 in. disk drive in its carrier. The disk drives are installed vertically in the chassis.
- The carriers have mounting locations for Direct doc SAS drives.
- A sheet steel carrier holds each drive which physically protects the drive. The carrier also provides thermal conduction, radio frequency, and electro-magnetic induction protection.
- The front cap also has an ergonomic handle which gives the following functions:
 - Secures the location of the carrier into and out of drive slots
 - Provides positive spring-loading of the drive (midplane) connector
- The carrier can use a dual-path direct dock Serial Attached SCSI interface.

The following figures show the supported drive carrier modules.




Figure 3. Dual path LFF 3.5 in. drive carrier module




Figure 4. Dual path SFF 2.5 in. drive carrier module




Figure 5. 2.5 in. to 3.5 in. 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 in. (2U12) and 2.5 in. (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 in. drive slot (left); 2.5 in. 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.

The DDIC features a slide latch button with directional arrow. The slide latch secures the DDIC into the disk slot within the drawer. The slide latch also enables you to disengage the DDIC from its slot to 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 in. disk drive.




Figure 7. 3.5 in. disk drive in a DDIC

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




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

Populating drawers with DDICs

The 5U84 enclosure does not ship with DDICs installed. Follow these guidelines when populating disks into the drawer.

- 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 alternate between the top drawer and the bottom drawer. For example, install first in 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 in. disks can be intermixed with DDICs holding 2.5 in. disks in the enclosure. However, each row should be populated with disks of the same form factor (all 3.5 in. disks or 2.5 in. disks).

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




Figure 9. 5U84 enclosure drawer fully populated with DDICs

Mount the enclosures in the rack

This section describes how to unpack the ME5 Series Storage System equipment, prepare for installation, and safely mount the enclosures into the rack.

Topics:

- [Rackmount rail kit](#)
- [Install the 2U enclosure using toolless rails](#)
- [Install the 2U enclosure front bezel](#)
- [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 to ensure that suitable mounting rails are available for the rack you plan to use.

Install the 2U enclosure using toolless rails

The 2U enclosure is delivered with the drives installed.

Prerequisites

- Remove the rail kit from the accessories box, and examine for damage.
- Ensure that the preassembled rails are the correct length for the rack. The rail lengths adjust from 60 cm (24 in.) to 86 cm (34 in.).
- Ensure that the rail pins on the front and back of the rails are correct pins for the type of rack that you are using. The rails are shipped from the factory with pins for racks with universal square holes or standard round holes. If you are using a rack with smaller holes, remove the existing pins from the rails and install the pins that correspond to your rack.




Figure 10. Toolless rails

About this task

When installing the rails, ensure that the rail pins are installed in the same rack holes at the rear and front of the rack so that the rails are level.

CAUTION: Lifting enclosures above a height of 20U is not recommended. Mechanical assistance is required to lift an enclosure above this height.

Steps

1. Identify the rack holes to use when installing the rails in the racks.
2. Insert the rail pins on the rear of the left rail into the rack holes of the rear rack post.
Ensure that the snap latch locks into place on the rear rack post.
3. Extend the left rail to fit between the front and rear rack posts.
4. Insert the rail pins onto the front of the left rail into the rack holes of the front rack post.
Ensure that the snap latch locks into place on the front rack post.




Figure 11. Insert the left rail in rack

5. Repeat the previous steps to install the right rail in the rack.
6. Use two people to lift the enclosure and align it with the installed rails.
i **NOTE:** Ensure that the enclosure remains level while inserting it in the rack.
7. Carefully insert the inner rails on each side of the enclosure into the rails.




Figure 12. Insert the enclosure in the rails

8. Push the enclosure fully into the rack.
9. Secure the enclosure to the rack using the enclosure fastening screws in the rack mount ears on the left and right side of the enclosure.




Figure 13. Secure the enclosure to the rack

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 14. 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 the detail view.

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 back 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 (26 in.) to 840 mm (33 in.). 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 15. Install the rail in the rack (left side 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)

Table 4. Install the rail in the rack (continued)

Item	Description	Item	Description
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. Secure the front and rear of the enclosure using the enclosure fastening screws.

Reinsert the back panel modules and install the DDICs into the drawers. See the instructions in the *Dell PowerVault ME5 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

ME5 Series storage systems 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 off the controller enclosure before connecting the expansion enclosures.

- ME5 Series 2U controller enclosures support up to ten 2U enclosures (including the controller enclosure), or a maximum of 240 disk drives.
- ME5 Series 5U controller enclosures support up to four 5U enclosures (including the controller enclosure), or a maximum of 336 disk drives.
- ME5 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 provides continued access to other enclosures if any drive enclosure fails or is removed. Fault tolerance and performance requirements determine whether to optimize the configuration for high availability or high performance when cabling.

Cable requirements for expansion enclosures

ME5 Series storage systems support 2U12, 2U24, and 5U84 expansion enclosures, each of which can be configured as a controller enclosure or an expansion enclosure.

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

- When connecting SAS cables to IOMs, use only supported HD mini-SAS x4 cables.
- Use qualified HD mini-SAS to HD mini-SAS 0.5 m (1.64 ft) cables to connect cascaded enclosures in the rack.
- The maximum enclosure cable length that is allowed in any configuration is 2 m (6.56 ft).
- Using more than two expansion enclosures may require additional cables depending upon the number of enclosures and the cabling method used.
- Reverse-cabling for a fault-tolerant configuration may require additional or longer cables.

In the following cabling diagrams, the controller enclosure is shown at the top of the stack of connected expansion enclosures. 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](#) for more detail.

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

The ME5 Series identifies controller modules and IOMs by enclosure ID and IOM ID.

The following figure shows the cabling configuration for a 2U controller enclosure with 2U expansion enclosures. 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.




Figure 16. Cabling connections between a 2U controller enclosure and 2U expansion enclosures

The following figure shows the cabling configuration for a 5U84 controller enclosure with 5U84 expansion enclosures. 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 (3B), 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.




Figure 17. Cabling connections between a 5U controller enclosure and 5U expansion enclosures

Label the back-end cables

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 recommends using separate VLANs.




Figure 18. Connect a 2U controller enclosure to the management network




Figure 19. Connect a 5U controller enclosure to the management network

NOTE: See also the topic about configuring network ports on controller modules in the *Dell PowerVault ME5 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 ME5 Series controller enclosures can connect to respective hosts using direct-attach or switch-attach methods.

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

Use only Dell cables for host connections:

Table 5. Cable and transceiver options

Controller protocol	Speed	Transceiver	Cable
FC32 Gb	32 Gb	FC SFP+	Multi-mode OM3/4 LC-LC cables. Choice of quantity and length up to: • OM3 100 meters maximum • OM4 150 meters maximum
FC32 Gb	16 Gb	FC SFP+	
iSCSI 25 Gb	25 Gb	iSCSI SFP28 - <i>or</i> - SFP28 Direct Attach Copper (DAC) cable	
iSCSI 25 Gb	10 Gb	iSCSI SFP+ - <i>or</i> - SFP+ DAC cable	
10GBaseT	10 Gb	N/A	RJ45 CAT-6 Copper Patch lead
12 Gb SAS	12 Gb	N/A	12Gb HD Mini-SAS to HD Mini-SAS cables. Choice of quantity and length up to 4 meters maximum

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 ME5 Series controller enclosures and shows some common cabling configurations. ME5 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.

Fibre Channel protocol

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 or 32 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](#).

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

When you connect to the storage system using FC switches, use the switch management interface to create zones for isolating traffic for each HBA.

Use the PowerVault Manager to set FC port speed and options. See [Setting up hosts](#). 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

The controller supports 25 GbE iSCSI ports. These ports can be used for:

- Connecting to 25 GbE iSCSI hosts directly, or through a switch used for the 25 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 PowerVault ME5 Series Storage System Administrator's Guide*.

Use the PowerVault Manager to set iSCSI port options. See [Setting up hosts](#). 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 create a VLAN to isolate iSCSI traffic from the rest of the switch traffic.

Example iSCSI port address assignments

The following figure and the supporting table provides 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 20. Two subnet switch example (IPv4)

In the following table, the last three columns explain the addressing method for the last two octets of the IP address.

Table 6. Two subnet switch IP addressing example

No.	Device	IP Address	3rd octet	4th octet	4th octet
			Subnet	Port (middle digit)	Controller (last digit)
1	A0	192.68.10.200	10	0	0
2	A1	192.68.11.210	11	1	0
3	A2	192.68.10.220	10	2	0
4	A3	192.68.11.230	11	3	0
5	B0	192.68.10.205	10	0	5
6	B1	192.68.11.215	11	1	5
7	B2	192.68.10.225	10	2	5
8	B3	192.68.11.235	11	3	5
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 PowerVault ME5 Series Storage System Administrator's Guide*.

SAS protocol

ME5 Series SAS models use 12 Gb/s host interface protocol and qualified cable options for host connection.

The 12 Gb SAS controller enclosures support two controller modules which contain four SFF-8644 HD mini-SAS host ports each. These host ports support data rates up to 12 Gb/s. The HD mini-SAS host ports connect directly to SAS hosts. The host computer must support SAS and optionally, multipath I/O. Use a qualified cable when connecting to a host.

Host connection

ME5 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.

32 Gb Fibre Channel host connection

Connect each FC port on the controller to a switch that is connected to the host ports as shown in the cabling diagram examples. You can also connect the storage system directly to the host server.

Connect to either a 16 Gb or 32 Gb host from the 32 Gb FC controller by using the applicable transceiver. Match the transceiver to your host/cable speed:

- For 16 Gb connections—use a FC SFP transceiver
- For 32 Gb connections—use a FC SFP+ transceiver

Use multi-mode OM3 or OM4 cables of the appropriate speed and length up to the following maximum:

- OM3—100 meters
- OM4—150 meters

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).

For information about configuring FC HBAs, see the FC topics under [Setting up hosts](#).

In addition to providing a host connection, these cables can be used for connecting two storage systems through a switch to facilitate use of the optional replication feature. See [Cabling for replication](#).

See the *Dell PowerVault ME5 Series Storage System Support Matrix* for supported Fibre Channel HBAs.

25 GbE iSCSI host connection

Connect each iSCSI port on the controller to a switch that is connected to the host ports as shown in the cabling diagram examples. You can also connect the storage system directly to the host server.

Connect to either a 10 Gb or 25 Gb host from the 25 GbE iSCSI controller by using the applicable transceiver and/or cable:

- For 10 Gb connections —use either an SFP+ transceiver and 10 Gb cable, or a 10 G SFP+ DAC cable
- For 25 Gb connections—use either an SFP28 transceiver and 25 Gb cable, or an SFP28 DAC cable

For information about configuring iSCSI initiators/HBAs, see the iSCSI topics under [Setting up hosts](#).

See the *Dell PowerVault ME5 Series Storage System Support Matrix* for supported iSCSI HBAs.

10GBase-T host connection

Connect each iSCSI port on the controller to a switch that is connected to the host ports as shown in the cabling diagram examples. You can also connect the storage system directly to the host server. To connect a 10GBase-T controller, use an RJ45 CAT-6 cable with a copper patch lead connector.

For information about configuring network adapters and iSCSI HBAs, see the iSCSI topics under [Setting up hosts](#).

12 Gb HD mini-SAS host connection

Connect each SAS port on the controller to a switch that is connected to the host ports as shown in the cabling diagram examples. To connect controller modules with HD mini-SAS host interface ports to a server HBA, use a 12 Gb HD mini-SAS to HD mini-SAS cable.

The ME5 Series storage controller supports cable lengths up to 4 meters.

For information about configuring SAS HBAs, see the SAS topics under [Setting up hosts](#).

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 SAS and 10Gbase-T host connections for ME5 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.
 - Single controller configurations are not supported in ME5084 systems.




Figure 21. Connecting hosts: ME5 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 ME5 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 22. Connecting hosts: ME5 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 23. Connecting hosts: ME5 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 24. Connecting hosts: ME5 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 25. Connecting hosts: ME5 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 26. Connecting hosts: ME5 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 27. Connecting hosts: ME5 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.

NOTE: About switch-attached configurations:

- See the recommended switch-attached examples for host connection in the *Setting Up Your Dell PowerVault ME5 Series Storage System* poster that is provided with your controller enclosure.
- See [Two subnet switch example \(IPv4\)](#) for an example showing host port and controller port addressing on an IPv4 network.




Figure 28. Connecting hosts: ME5 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 29. Connecting hosts: ME5 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](#) instructions up to this point. After you connect the power cables and power on the system, 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 a PDU (power distribution unit) in the rack.

The power cables must be connected to at least two separate and independent power supplies to ensure redundancy.




Figure 30. Typical AC power cable connection from PCM to PDU (2U)




Figure 31. Typical AC power cable connection from PSU to PDU (5U)

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

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 [Technical specifications](#). 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.

- 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.

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.

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](#) for details about the 2U Ops panel LEDs and related fault conditions.
- See [5U enclosure Ops panel](#) for details about the 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

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

You can configure your system using the guided setup described in this section. The system configuration can also be modified or completed using the features within PowerVault Manager.

Topics:

- [Prerequisites](#)
- [About guided setup](#)
- [Access the PowerVault Manager](#)
- [System configuration](#)
- [Set up SupportAssist and APEX AIOps Infrastructure Observability](#)
- [Storage configuration](#)
- [Provisioning](#)

Prerequisites

Before starting the guided setup, make sure that:

- Hosts are configured and connected to the storage system
- Initiators are available/identified
- Switch zoning is complete (for Fibre channel protocol)
- System and network information is recorded
- A business account is set up on dell.com and validated
- Access key and PIN are available

Record storage system information

Use the [System Information Worksheet](#) to record the information that you need to install the ME5 Series storage system.

About guided setup

After you log in to PowerVault Manager, the setup wizard guides you through process of configuring your system.

The guided setup includes the following tasks:

- System Configuration—Network settings, set date and time, add users, set up notifications, and if applicable, add iSCSI connectivity.
- SupportAssist Configuration—Accept license, set up connection, and add support contacts
- Storage Configuration—set storage type and set up storage pools
- Provisioning—Add groups, set up hosts, and add volumes

Access the PowerVault Manager

Start the initial configuration by logging in to the PowerVault Manager, changing the password, and verifying the firmware version.

About this task

 **NOTE:** To avoid IP conflicts, do not turn on more than one unconfigured controller enclosure at a time.

Steps

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. 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.
 2. To read the license agreement click **EULA**, review the EULA and click **Close**.
 3. Click **Get Started**.
 4. Type a new user name for the storage system in the **Username** field. A username is case sensitive and can have a maximum of 29 bytes. The name cannot already exist in the system, include spaces, or include any of the following: " , < \
 5. Type password for the new username in the **Password** and **Confirm Password** fields. A password is case sensitive and can have from 8 to 32 characters. If the password contains only printable ASCII characters, then it must contain at least one uppercase character, one lowercase character, one numeric character, and one non-alphanumeric character. A password can include printable UTF-8 characters except for the following: a space or " ' , < > \
 6. Click **Apply And Continue**.
 7. The storage system creates the user and displays the **Update Firmware** panel.
 8. For the initial configuration, click **Use Current Firmware Bundle**.
-  **NOTE:** For future firmware updates, you can upload and install new firmware from the **Maintenance > Firmware** panel in the PowerVault Manager. 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.

The System Configuration Main Page opens.

System configuration

The system configuration setup includes network configuration, setting the date and time, adding users, setting your notification preferences, and if applicable, setting up iSCSI connectivity. The system configuration can be changed if needed after initial setup using the **Settings** menu in the PowerVault Manager.

For more details about configuration settings, see the *Dell PowerVault ME5 Series Administrator's Guide*.

Click **Start** to begin the system configuration.

Configuring controller network ports

The system provides concurrent support for IPv4 and IPv6 protocols. Both protocols can be set up at the same time by configuring the network parameters.

You can manually set static IP address parameters for network ports, or you can specify that IP values be set automatically, using DHCP (Dynamic Host Configuration Protocol) for IPv4 or DHCPv6 or SLAAC (Stateless address auto-configuration) for IPv6.

 **NOTE:** SLAAC relies on Neighbor Discovery Protocol (NDP), and is the simplest way to provide an IPv6 address to a client.

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

addressing. Given that network parameter settings are independent between the two protocols, you can set them as needed for IP addressing on controller B.

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, the system will use its default values. You must have some means of determining what addresses have been assigned, such as the list of bindings on the DHCP server. You can retrieve the DHCP assigned IP addresses either through the USB serial console login page, which lists IPv4 and IPv6; via CLI commands; or from the DHCP server list of MAC address to IP address leases. When using Auto mode, addresses are retrieved from both DHCP and 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

 **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, and 169.254.255.4. Because these addresses are routable, do not use them anywhere in 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 is able to 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 address settings can cause management hosts to lose access to the storage system after the changes are applied in the confirmation step.

Network Settings

The **Network Settings** panel provides options for you to configure IPv4 and IPv6 network-port settings and configure a DNS server. The network settings can be changed if needed after initial setup using the **Settings > Network** panel in the PowerVault Manager.

Steps

1. On the **System Configuration Main Page**, click **Start** under **System Configuration**.
2. Select the network settings to configure:
 - **IPv4**
 - **IPv6**
 - **DNS**—selected automatically
 - **Skip this step**
3. Click **Continue**.

Set IPv4 addresses for network ports

Prerequisites

The IPv4 option was selected in the guided setup and the **Network Settings: IPv4** page is displayed.

Steps

1. In the Source section, select the type of IP address settings to use for each controller:
 - Select **Manual** to specify static IP addresses.
 - Select **DHCP** to allow the system to automatically obtain IP addresses from a DHCP server.
2. If you selected Manual, type the IP address, IP mask, and Gateway addresses for each controller



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.

3. If you selected DHCP, and the controllers successfully obtained IP addresses from the DHCP server, the new IP addresses are displayed
4. When settings for both controllers are complete, click **Apply and Continue**.
The **Network Settings: DNS** page opens.

Set IPv6 addresses for network ports

Prerequisites

The IPv6 option was selected in the guided setup and the **Network Settings: IPv6** page is displayed.

Steps

1. In the Source section, select the type of IP address settings to use for each controller:
 - Select **Manual** to specify static IP addresses.
 - Select **Auto** to allow the system to automatically obtain IP addresses.
2. If you selected Manual, type the Gateway and Static IP addresses for each controller. You can set up to four static IP addresses for each controller.
3. If you selected Auto, the address obtained by the system is typically displayed as the Link-Local Address. If SLAAC is used for automatic discovery, the SLAAC IP address will be displayed.
4. When settings for both controllers are complete, click **Apply and Continue**.
The **Network Settings: DNS** page opens.

Set DNS settings for network ports

Prerequisites

The DNS option was selected in the guided setup and the **Network Settings: DNS** page is displayed.

Steps

1. (Optional) Change the **Host Name**.
2. (Optional) Search for a domain in the **Search Domain** box. If a Search Domain has not been added to the network previously, you can add one by clicking on the **Add Another Search Domain** link.
3. (Optional) Add up to two more DNS servers.
4. When the settings for both controllers are complete, click **Apply And Continue**.

Set the date and time

You can set the date and time manually or configure the system to use Network Time Protocol (NTP) to obtain date and time from an available network-attached server. Using NTP allows multiple storage devices, hosts, log files, and so on to be synchronized. The NTP server address value can be an IPv4 address, IPv6 address, or FQDN. If NTP is enabled but no NTP server is present, the date and time are maintained as if NTP was not enabled.

Prerequisites

The **Set Date and Time** page is displayed in the guided setup.

About this task

The date and time settings can be changed if needed after initial setup using the **Settings > System > Date and Time** panel in the PowerVault Manager or by clicking on the date and time displayed in the banner.

Steps

1. Select either **Network Time Protocol (NTP)** or **Manual**.
 - For Manual setting, enter the current (local) date and time.
 - For NTP setting, enter the NTP server address and the NTP time zone offset .
2. Click **Apply And Continue**.

Set up users

When you logged in to the PowerVault Manager for the first time, you created a user for managing the system. You can add users to the system through the guided setup. The user settings can be changed and new users can be added after initial setup using the **Settings > Users** panel in the PowerVault Manager.

Prerequisites

The **User Settings** page is displayed in the guided setup.

Steps

1. Select the type of user to set up:
 - **Local**
 - **SNMPv3**
 - **Skip this step**
2. Click **Continue**.

Set up local users

The user settings can be changed if needed after initial setup using the **Settings > Users** panel in the PowerVault Manager.

Prerequisites

The **Local User** option was selected in the guided setup and the **User Settings: Local** page is displayed.

Steps

1. Click **Add New User**.
2. Enter information for the new user:
 - **Username**—A username is case sensitive and can have a maximum of 29 bytes. The name cannot already exist in the system, include spaces, or include any of the following: " , < \
 - **Password**—A password is case sensitive and can have from 8 to 32 printable characters. If the password contains only ASCII characters, it must contain at least one uppercase character, one lowercase character, one numeric character, and one non-alphanumeric character. A password can include UTF-8 characters except for the following: a space or " ' , < > \
 - **Interfaces**—Select one or more of the following interfaces:
 - **WBI**. Enables access to the PowerVault Manager.
 - **CLI**. Enables access to the command-line interface.
 - **FTP**. Enables access to the FTP interface or the SFTP interface, which can be used instead of the PowerVault Manager to install firmware updates and to download logs.
 - **Roles**—Select one or more of the following roles:
 - **Manage**. Enables the user to change system settings.
 - **Monitor**. Enables the user to view but not change system status and settings.
 - **Language**—Select a display language for the user. The default is English. Installed language sets include Chinese-Simplified, English, French, German, Japanese, Korean, and Spanish. The locale determines the character used for the decimal (radix) point. The locale setting is determined by the Language setting, which can be accessed by selecting the pencil icon for any user in the table.
 - **Temperature Preference**—Select whether to use the Celsius or Fahrenheit scale for display of temperatures. The default is Celsius.
 - **Timeout**—Select the amount of time that the user session can be idle before the user is automatically signed out (from 2 to 720 minutes). The default is 30 minutes.
3. Click **Create New User**.

4. Continue to add new users, and click **Apply And Continue** when complete.

Set up SNMPv3 users

SNMPv3 users must exist on the system to add them to PowerVault Manager.

Prerequisites

- SNMPv3 users are set up on the network.
- The **SNMPv3** option was selected in the guided setup and the **User Settings: SNMP** page is displayed.

Steps

1. Click **Create SNMPv3 User**.
2. Enter information for the new user:
 - **Username**—A username is case sensitive and can have a maximum of 29 bytes. The name cannot already exist in the system, include spaces, or include any of the following: " , < \
 - **Password**—A password is case sensitive and can have from 8 to 32 printable characters. If the password contains only ASCII characters, it must contain at least one uppercase character, one lowercase character, one numeric character, and one non-alphanumeric character. A password can include UTF-8 characters except for the following: a space or " ' , < > \
 - **Authentication Type**. Select whether to use MD5 or SHA (SHA-1) authentication, or no authentication. If authentication is enabled, the password set in the Password and Confirm Password fields must include a minimum of 8 characters and follow the other SNMPv3 privacy password rules.
 - **Privacy Type**. Select whether to use DES or AES encryption, or no encryption. To use encryption you must also set a privacy password and enable authentication.
 - **Privacy Password**. If the privacy type is set to use encryption, specify an encryption password. This password is case sensitive and can have from 8 to 32 characters. If the password contains only printable ASCII characters, then it must contain at least one uppercase character, one lowercase character, one numeric character, and one non-alphabetic character. A password can include printable UTF-8 characters except for the following: a space or " ' , < > \
 - **Trap Host Address**. Specify the network address of the host system that will receive SNMP traps. The value can be an IPv4 address, IPv6 address, or FQDN.
3. Click **Create SNMPv3 User**.
4. Continue to add new users, and click **Apply And Continue** when complete.

Notifications

The Notifications panel provides options to send system alert notifications to users through email, SNMP trap hosts, or a remote syslog server. The notification settings can be changed if needed after initial setup using the **Settings > Notifications** panel in the PowerVault Manager.

About this task

Enable at least one notification service to monitor the system.

Steps

1. Select the type of notification to set up:
 - **Email**
 - **SNMP**
 - **Syslog**
 - **Skip this step**
2. Click **Continue**.

Set up Email notifications

Use the Email Notifications panel to choose to be notified by email when system alerts occur. Alert notifications can be sent to a maximum of three email addresses. Weekly alerts concerning system health issues will also be sent until corrective action has

been taken and the system health value has returned to OK. Enter information in the text boxes to receive alert notifications. For details about panel options, see the on-screen tool tips.

Set up SNMP notifications

Use the SNMP panel to set options for sending alert notifications to SNMP trap hosts. You must enable SNMP for the system to send alert notifications to SNMP users. Enter information in the text boxes to receive alert notifications. For details about panel options, see the on-screen tool tips.

Set up syslog notifications

Use the Syslog panel to set remote syslog notifications to allow alerts to be logged by the syslog of a specified host computer. Syslog is a protocol for sending alert messages across an IP network to a logging server. This feature supports User Datagram Protocol (UDP), but not Transmission Control Protocol (TCP). For details about panel options, see the on-screen tool tips.

Configure iSCSI ports

If your system uses iSCSI ports, the guided setup assists you in configuring the iSCSI ports. The iSCSI settings can be changed or set after initial setup using the **Settings > iSCSI** panel in the PowerVault Manager

Steps

1. On the **iSCSI Settings** panel, configure the following settings:
 - **IP Version.** Select whether to use IPV4 or IPV6. IPV4 uses 32-bit addresses. IPV6 uses 128-bit addresses.
 - **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.
 - **CHAP Authentication.** Enables or disables use of Challenge Handshake Authentication Protocol. Enabling or disabling CHAP in this panel updates the setting in the Configure CHAP panel.
 - **ISNS.** Enables or disables registration with a specified Internet Storage Name Service server, which provides name-to-IP-address mapping. If selected, then specify the IP address of an iSNS server and an alternate iSNS address. The alternate address can be on a different subnet.
 **CAUTION:** Changing IP settings can cause data hosts to lose access to the storage system.
2. Click **Continue**.
3. On the **Host port addresses** panel, set the IP address, netmask, and gateway for each port on both controllers.
4. Click **Continue**.
If you selected CHAP Authentication, the **CHAP Authentication** panel opens.
5. Configure the CHAP settings:
 - **Initiator Name.** Enter a name for the initiator.
 - **Mutual CHAP.** Select to require that the storage system authenticate itself to the host. Without mutual CHAP, only the initiator is authenticated to the target.
 - **Initiator Authentication Secret.** Enter a secret for the initiator. The secret is a text string that is known to both the initiator and the storage array. It must have 12-16 characters, and include spaces and printable UTF-8 characters except: " or <
 - **Target Authentication Secret.** Enter a secret for the target. The secret is a text string that is known to both the initiator and the storage array. It must have 12-16 characters, and include spaces and printable UTF-8 characters except: " or <
6. Click **Continue**.

Set up SupportAssist and APEX AI Ops Infrastructure Observability

SupportAssist provides an enhanced support experience for ME5 Series storage systems by sending configuration and diagnostic information to technical support at regular intervals. APEX AI Ops Infrastructure Observability provides storage monitoring and proactive service, with access to near real-time analytics, and the ability to monitor storage systems from anywhere at any time.

Prerequisites

- You have a [business account](#) with Dell.
- You have an [access key](#).
- Network requirements are met as described in [SupportAssist direct connection requirements](#).
- You have a ProSupport contract to use [APEX AI Ops Infrastructure Observability](#).

Steps

1. On the **System Configuration Main Page**, click **Start** under **SupportAssist Configuration**.
 2. In the **License Agreement** panel, read through the agreement, and then acknowledge it by selecting **I accept this agreement**.
 3. Click **ACCEPT AND CONTINUE**.
 4. Choose the support and monitoring features to use:
 - **SupportAssist**—select to send configuration and diagnostic information to technical support at regular intervals.
 - **Connect to APEX AI Ops Infrastructure Observability**—select to use APEX AI Ops Infrastructure Observability for storage monitoring and proactive service.
-  **NOTE:**
5. In the **Connection Information** panel, select your connectivity options:
 - **Connection Type**. Select whether to connect directly or through a gateway.
 - **Proxy Details**. If applicable, select **Use a Proxy Server** and then enter the server settings.
 - **Access key and PIN**. Enter the information requested. If you do not have the access key or PIN, click **request a new Access Key and PIN** and follow the prompts to have new key information emailed to you.
 6. Click **Test and Enable Connectivity**.
Test results are displayed, you can either go back and reenter information, or click **Continue** to proceed.
 7. In the **Contact Information** panel, enter the primary contact information and select the preferred contact settings. You can also enter information for a secondary contact.
 8. Click **Continue**.

Storage configuration

The **Storage Configuration** setup provides options for you to configure storage on your system.

Steps

1. On the **System Configuration Main Page**, click **Start** under **Storage Configuration**.
2. In the **Select Storage Type** panel, review the option descriptions for Virtual and Linear storage. You can also choose to skip this step and configure storage later using **Maintenance > Settings > Storage** in the PowerVault Manager.
 - **Virtual**
 - **Linear**
 - **Skip this step**
3. Click **Continue**.

Set up virtual storage

When you set up virtual storage, you create pools. A pool is an aggregation of one or more disk groups that serves as a container for volumes. You can have the system set up your pools automatically, or you can create pools by manually adding

disk groups. You can also add or edit the storage settings after initial setup using the **Maintenance > Storage** panel in the PowerVault Manager.

Prerequisites

The **Storage Type > Pools** panel is displayed in the setup wizard.

About this task

For detailed information about pools and disk groups, see the *ME5 Series Administrator's Guide*.

Steps

1. To automatically set up storage, click **Auto Storage Setup**.
 - a. Verify that the Disk Scan results indicate that the system is Healthy.
 - b. Review the pool configuration and if it meets your needs, click **Apply Configuration**, otherwise click **Cancel** and set up storage manually.
After applying the configuration, the system configures the pools and spares and displays a success message when complete.
 - c. Click **Ok**.
2. To manually set up storage, expand **Pool A** and click **Add Disk Group**. The Add Disk Group panel opens.
 - a. In the **Configuration** section, choose the **Protection Level** (RAID) from the drop down box.
 - b. In the **Available Disks** section, select the disks to include in the pool.
 **NOTE:** It is recommended that disks and provisioning are balanced between Pool A and Pool B.
 - c. Review the **Summary**, and click **Add Disk Group**.
After applying the configuration, the system configures the pools and spares and displays a success message when complete.
 - d. Click **Ok**.
 - e. Repeat these steps for **Pool B**.
3. In the **Storage Type > Pools** panel click **Continue**.

Set up linear storage

When you set up linear storage, you create pools. A pool is an aggregation of one or more disk groups that serves as a container for volumes. You can add or edit the storage settings after initial setup using the **Maintenance > Storage** panel in the PowerVault Manager

Prerequisites

The **Storage Type > Pools** panel is displayed in the setup wizard.

About this task

For detailed information about pools and disk groups, see the *ME5 Series Administrator's Guide*.

Steps

1. To manually set up storage, click **Add Disk Group**. The Add Disk Group panel opens.
2. In the **Configuration** section, choose the **Protection Level** (RAID) from the drop down box.
 - a. **Name**. Enter a name for the disk group.
 - b. **Assigned Controller**. Select either controller A or controller B, or select Auto to have the system choose where to write data for the most efficient storage.
 - c. **Protection Level**. Select the RAID level for this disk group.
 - d. **Chunk Size**. The amount of contiguous data that is written to a disk group member before moving to the next member of the disk group. Select from 64 KB, 128 KB, 256 KB, or 512 KB.
 - e. **Online Initialization**. Select to make the pool accessible before the initialization is complete.
3. In the **Available Disks** section, select the disks to include in the pool.

 **NOTE:** It is recommended that disks and provisioning are balanced between Pool A and Pool B.

4. Review the **Summary**, and click **Add Disk Group**.
After applying the configuration, the system configures the pools and spares and displays a success message when complete.
5. Click **Ok**.
6. Repeat these steps for **Pool B**.
7. In the **Storage Type > Pools** panel click **Continue**.

Provisioning

The Provisioning setup guides you through the process to connect to hosts and create initial volumes. Hosts must be configured and attached to the ME5 Series storage system to complete provisioning.

About this task

See [Setting up hosts](#) for information about configuring host servers.

Steps

1. On the **System Configuration Main Page**, click **Start** under **Provisioning**.
The introduction panel opens, describing the provisioning process.
2. Click **Continue**.

Set up hosts

When you set up hosts, the system must find initiators that are previously set up on the network. You create hosts from the initiators that were found. You can also add or edit the host settings after initial setup using the **Provisioning > Hosts** panel in PowerVault Manager.

Prerequisites

- Hosts are configured as described in [Host setup](#).
- The **Provisioning > Hosts** panel is displayed in the setup wizard.

Steps

1. Select **Create a New Host**.
2. Enter a **Host Name**.
3. Select an initiator from the list to assign to this host.
The host is displayed in the **New Hosts** list.
4. Click **Continue**.

Set up volumes

Next the setup wizard guides you through the process of setting up volumes. You can also add or edit volumes after initial setup using the **Provisioning > Volumes** panel in the PowerVault Manager.

Prerequisites

The **Provisioning > Volumes** panel is displayed in the setup wizard.

Steps

1. If you want to attach volumes now, select **Attach host or host groups to volumes**. You can skip this step and set up volumes later if you prefer.
2. If you are attaching volumes now, select whether to create new volumes or select existing volumes to attach to the host.
3. Click **Continue**.

4. Select the pool for the new volume and enter a **Volume Name**.
5. Enter the **Volume Size** and select the units of measure. Optionally, you can choose to use the remaining space for the volume.
6. Click **Add Volume**.
7. Review the volume parameters. From this panel you can:
 - Delete the volume ()
 - **Add New Volume**
 - Click **Continue** to proceed.The provisioning summary is displayed.
8. Review the provisioning configuration and click **Continue** to proceed, or **Back** to return to make changes to the provisioning.
9. Click **OK** at the Success prompt.
The final provisioning panel is displayed. From you can configure additional hosts or click **Continue** to return to the **System Configuration Main Page** and exit to the PowerVault Manager Dashboard.

Next steps

Set up multipathing and configure volumes on the host as described in [Host setup](#)

Setting up hosts

This section describes the end-to-end process to set up hosts and add volumes for Dell PowerVault ME5 Series storage systems. You can also set up hosts and volumes using the guided setup.

For more information, see the topics about initiators, hosts, and host groups, and attaching hosts and volumes in the *Dell PowerVault ME5 Series Storage System Administrator's Guide*.

Topics:

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

Host system requirements

Host server requirements include considerations for multipathing, switch zoning, and HBA identification.

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

Multipathing I/O

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.
- ME5 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 host servers to use multipath I/O (MPIO).

Fibre channel switch zoning

If the hosts are connected to the storage system by FC switches, implement zoning to isolate traffic for each HBA. Use the FC switch management interface to create a zone for each server HBA. Each zone must contain only one HBA.

See the [Fibre Channel zoning information](#) in the System information worksheet for more information.

Host adapters

For a list of supported HBAs, see the *ME5 Series Storage System Support Matrix* on the Dell support site.

- Ensure that all HBAs have the latest supported firmware and drivers as described on Dell.com/support.
- For Fibre channel and SAS protocols, identify and record the WWNs of the HBAs that are used by the ME5 Series Storage System.
- For iSCSI protocols, identify and record the IP addresses of the HBAs that are used by the ME5 Series Storage System.

Windows hosts

ME5 Series storage systems support Windows host servers using Fibre Channel, iSCSI, or SAS protocol.

Configuring a Windows host with FC HBAs

The following steps describe the end-to-end process for setting up hosts and provisioning volumes. This process can be done after the guided setup.

Prerequisites

- Ensure that all HBAs are installed and have the latest supported firmware and drivers as described on Dell.com/support. For a list of supported HBAs, see the *Dell ME5 Series Storage System Support Matrix* on the Dell support site.
- Cable the host servers as described in [Cable host servers to the storage system](#).

Install MPIO on the host

Perform the following steps to install MPIO on the Windows server.

Steps

1. Open the Server Manager.
2. Click **Add Roles and Features**.
3. Click **Next** until you reach the Features page.
4. Select **Multipath IO**.
5. Click **Next**, click **Install**, and then click **Close**.
6. Reboot the host server.

Identify the FC WWNs and set up switch zoning

Record the FC HBA WWNs on the System Information Worksheet and set up switch zoning as required.

Steps

1. Identify and record 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.
2. 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 ME5 Series storage systems support single initiator/multiple target zones.

Create a host and attach volumes in PowerVault Manager

If you did not set up hosts during the guided setup, or if you want to add new hosts, use the PowerVault Manager create hosts and attach volumes.

Steps

1. In the PowerVault Manager Dashboard, go to **Provisioning > Hosts**.
The Hosts panel opens with the **Hosts and Host Groups** table selected.
2. Click **Create Host**.
3. In the Create Host panel, select the **Create a New Host** radio button.
4. Enter a **Host Name**.
5. Select the initiator to assign to this host, using your worksheet as a guide to map the WWN and the Initiator ID.
6. (Optional) Enter a nickname for the this host initiator that clearly identifies the initiator for that particular host.
7. Click **Add Initiators To Host**.
The host is displayed in the **New Hosts** list.
8. Click **Continue**.
9. If you want to attach volumes now, select **Attach host or host groups to volumes**. You can skip this step and set up volumes later if you prefer.
10. If you are attaching volumes now, select whether to create new volumes or select existing volumes to attach to the host.
11. Click **Continue**.
12. If you are creating new volumes:
 - a. Select the pool for the new volume and enter a **Volume Name**. Use a name that indicates how the volume is used, such as `{host name}_Host1_Vol1`.
 - b. Enter the **Volume Size** and select the units of measure. Optionally, you can choose to use the remaining space for the volume.
 - c. Click **Add Volume**.Review the volume parameters. From this panel you can:
 - Delete the volume ()
 - **Add New Volume**
13. If you are using an existing volume, select the volume or volumes to attach to the host.
14. Click **Continue** to proceed.
The provisioning summary is displayed.
15. Review the provisioning configuration and click **Continue** to proceed, or **Back** to return to make changes to the provisioning.
16. Click **OK** at the Success prompt and return to the PowerVault Manager Dashboard.

Enable MPIO for the volumes on the Windows host

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

Steps

1. Open the Server Manager.
2. Select **Tools > MPIO**.
3. Click the **Discover Multi-Paths** tab.
4. Select **DellEMC ME5** in the **Device Hardware Id** list.
If **DellEMC ME5** is not listed in the **Device Hardware Id** list:
 - a. Ensure that the host has a minimum of one logical connection to both controllers.
 - b. Ensure that **DellEMC ME5** 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:

Steps

1. Open Server Manager.
2. Select **Tools > Computer Management**.
3. Right-click on **Disk Management** and then select **Rescan Disks**.
4. Right-click on the new disk and then select **Online**.
5. Right-click on the new disk and then 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 **New Simple Volume**, and follow the steps in the wizard to create the volume.

Configuring a Windows host with iSCSI network adapters

The following steps describe the end-to-end process for setting up hosts and provisioning volumes. This process can be done after the guided setup. These instructions document IPv4 configuration with dual switch subnet for network redundancy and failover. These instructions do not cover IPv6 configuration.

Prerequisites

- Ensure that the latest host operating system is installed and configured on the server.
- Ensure that all HBAs are installed and have the latest supported firmware and drivers as described on Dell.com/support. For a list of supported HBAs, see the *Dell ME5 Series Storage System Support Matrix* on the Dell support site.
- Cable the host servers as described in [Cable host servers to the storage system](#).
- Record the IP addresses assigned to each port as shown in the following example.

Table 7. Example worksheet for IP addresses

Port	IP Address
Subnet 1	
Host server 1, Port 0	192.68.10.20
Host server 2, Port 0	192.68.10.21
ME5 controller A port 0	192.68.10.200
ME5 controller A port 2	192.68.10.220
ME5 controller B port 0	192.68.10.205
ME5 controller B port 2	192.68.10.225
Subnet 2	
Host server 1, Port 1	192.68.11.20
Host server 2, Port 1	192.68.11.21
ME5 controller A port 1	192.68.11.210
ME5 controller A port 3	192.68.11.230
ME5 controller B port 1	192.68.11.215
ME5 controller B port 3	192.68.11.235

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:

About this task

 **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.

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

Steps

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.
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 (Subnet 1 and Subnet 2 in the example worksheet above).
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.

Install MPIO on the host

Perform the following steps to install MPIO on the Windows server.

Steps

1. Open the Server Manager.
2. Click **Add Roles and Features**.
3. Click **Next** until you reach the Features page.
4. Select **Multipath IO**.
5. Click **Next**, click **Install**, and then click **Close**.
6. Reboot the host server.

Configure the iSCSI Initiator on the Windows host

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

Steps

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.
- Repeat steps 6-7 for the NIC to establish a connection to each port on the subnet.
 - Repeat steps 3-8 for the NIC 2, connecting it to the targets on the second subnet.
-  **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.
-  **NOTE:** This step is required for multi-path configurations.
- Click the **Configuration** tab and record the initiator name specified in the **Initiator Name** field. The initiator name is needed to map volumes to the host.
- Click **OK** to close the **iSCSI Initiator Properties** dialog box.

Create a host and attach volumes in PowerVault Manager

If you did not set up hosts during the guided setup, or if you want to add new hosts, use the PowerVault Manager create hosts and attach volumes.

Steps

- In the PowerVault Manager Dashboard, go to **Provisioning > Hosts**.
The Hosts panel opens with the **Hosts and Host Groups** table selected.
- Click **Create Host**.
- In the Create Host panel, select the **Create a New Host** radio button.
- Enter a **Host Name**.
- Select the initiator to assign to this host, using your worksheet as a guide to map the WWN and the Initiator ID.
- (Optional) Enter a nickname for the this host initiator that clearly identifies the initiator for that particular host.
- Click **Add Initiators To Host**.
The host is displayed in the **New Hosts** list.
- Click **Continue**.
- If you want to attach volumes now, select **Attach host or host groups to volumes**. You can skip this step and set up volumes later if you prefer.
- If you are attaching volumes now, select whether to create new volumes or select existing volumes to attach to the host.
- Click **Continue**.
- If you are creating new volumes:
 - Select the pool for the new volume and enter a **Volume Name**. Use a name that indicates how the volume is used, such as *{host name}_Host1_Vol1*.
 - Enter the **Volume Size** and select the units of measure. Optionally, you can choose to use the remaining space for the volume.
 - Click **Add Volume**.

Review the volume parameters. From this panel you can:

 - Delete the volume ()
 - **Add New Volume**
- If you are using an existing volume, select the volume or volumes to attach to the host.
- Click **Continue** to proceed.
The provisioning summary is displayed.
- Review the provisioning configuration and click **Continue** to proceed, or **Back** to return to make changes to the provisioning.
- Click **OK** at the Success prompt and return to the PowerVault Manager Dashboard.

Enable MPIO for the volumes on the Windows host

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

Steps

1. Open the Server Manager.
2. Select **Tools > MPIO**.
3. Click the **Discover Multi-Paths** tab.
4. Select **DellEMC ME5** in the **Device Hardware Id** list.
If **DellEMC ME5** is not listed in the **Device Hardware Id** list:
 - a. Ensure that the host has a minimum of one logical connection to both controllers.
 - b. Ensure that **DellEMC ME5** 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:

Steps

1. Open Server Manager.
2. Select **Tools > Computer Management**.
3. Right-click on **Disk Management** and then select **Rescan Disks**.
4. Right-click on the new disk and then select **Online**.
5. Right-click on the new disk and then 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 **New Simple Volume**, and follow the steps in the wizard to create the volume.

Update the iSCSI initiator on the Windows host

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

Steps

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.

Configuring a Windows host with SAS HBAs

The following steps describe the end-to-end process for setting up hosts and provisioning volumes. This process can be done after the guided setup.

Prerequisites

- Ensure that all HBAs are installed and have the latest supported firmware and drivers as described on Dell.com/support. For a list of supported HBAs, see the *Dell ME5 Series Storage System Support Matrix* on the Dell support site.
- Cable the host servers as described in [Cable host servers to the storage system](#).

Install MPIO on the host

Perform the following steps to install MPIO on the Windows server.

Steps

1. Open the Server Manager.
2. Click **Add Roles and Features**.
3. Click **Next** until you reach the Features page.
4. Select **Multipath IO**.
5. Click **Next**, click **Install**, and then click **Close**.
6. Reboot the host server.

Identify SAS HBAs on a Windows server

Follow these steps to identify the SAS HBA initiators to connect to the storage system.

Steps

1. Open a Windows PowerShell console.
2. Type `Get-InitiatorPort` and press Enter.
3. Locate and record the SAS HBA WWNs .

Create a host and attach volumes in PowerVault Manager

If you did not set up hosts during the guided setup, or if you want to add new hosts, use the PowerVault Manager create hosts and attach volumes.

Steps

1. In the PowerVault Manager Dashboard, go to **Provisioning > Hosts**.
The Hosts panel opens with the **Hosts and Host Groups** table selected.
2. Click **Create Host**.
3. In the Create Host panel, select the **Create a New Host** radio button.
4. Enter a **Host Name**.
5. Select the initiator to assign to this host, using your worksheet as a guide to map the WWN and the Initiator ID.
6. (Optional) Enter a nickname for the this host initiator that clearly identifies the initiator for that particular host.
7. Click **Add Initiators To Host**.
The host is displayed in the **New Hosts** list.
8. Click **Continue**.
9. If you want to attach volumes now, select **Attach host or host groups to volumes**. You can skip this step and set up volumes later if you prefer.
10. If you are attaching volumes now, select whether to create new volumes or select existing volumes to attach to the host.
11. Click **Continue**.
12. If you are creating new volumes:
 - a. Select the pool for the new volume and enter a **Volume Name**. Use a name that indicates how the volume is used, such as `{host name}_Host1_Vol1`.
 - b. Enter the **Volume Size** and select the units of measure. Optionally, you can choose to use the remaining space for the volume.
 - c. Click **Add Volume**.Review the volume parameters. From this panel you can:
 - Delete the volume ()
 - **Add New Volume**
13. If you are using an existing volume, select the volume or volumes to attach to the host.
14. Click **Continue** to proceed.

The provisioning summary is displayed.

15. Review the provisioning configuration and click **Continue** to proceed, or **Back** to return to make changes to the provisioning.
16. Click **OK** at the Success prompt and return to the PowerVault Manager Dashboard.

Enable MPIO for the volumes on the Windows host

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

Steps

1. Open the Server Manager.
2. Select **Tools > MPIO**.
3. Click the **Discover Multi-Paths** tab.
4. Select **DellEMC ME5** in the **Device Hardware Id** list.
If **DellEMC ME5** is not listed in the **Device Hardware Id** list:
 - a. Ensure that the host has a minimum of one logical connection to both controllers.
 - b. Ensure that **DellEMC ME5** 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:

Steps

1. Open Server Manager.
2. Select **Tools > Computer Management**.
3. Right-click on **Disk Management** and then select **Rescan Disks**.
4. Right-click on the new disk and then select **Online**.
5. Right-click on the new disk and then 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 **New Simple Volume**, and follow the steps in the wizard to create the volume.

Linux hosts

ME5 Series storage systems support Linux host servers using Fibre Channel, iSCSI, or SAS protocol.

Configuring a Linux host with FC HBAs

The following steps describe the end-to-end process for setting up hosts and provisioning volumes. This process can be done after the guided setup.

Prerequisites

- Ensure that all HBAs are installed and have the latest supported firmware and drivers as described on Dell.com/support. For a list of supported HBAs, see the *Dell ME5 Series Storage System Support Matrix* on the Dell support site.
- Cable the host servers as described in [Cable host servers to the storage system](#).
- 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).

Identify FC HBAs on a Linux server

Perform the following steps to identify the Fibre Channel HBAs on a Linux host.

Steps

1. 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.
2. 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 ME5 Series storage systems support single initiator/multiple target zones.

Create a host and attach volumes in PowerVault Manager

If you did not set up hosts during the guided setup, or if you want to add new hosts, use the PowerVault Manager create hosts and attach volumes.

Steps

1. In the PowerVault Manager Dashboard, go to **Provisioning > Hosts**.
The Hosts panel opens with the **Hosts and Host Groups** table selected.
2. Click **Create Host**.
3. In the Create Host panel, select the **Create a New Host** radio button.
4. Enter a **Host Name**.
5. Select the initiator to assign to this host, using your worksheet as a guide to map the WWN and the Initiator ID.
6. (Optional) Enter a nickname for the this host initiator that clearly identifies the initiator for that particular host.
7. Click **Add Initiators To Host**.
The host is displayed in the **New Hosts** list.
8. Click **Continue**.
9. If you want to attach volumes now, select **Attach host or host groups to volumes**. You can skip this step and set up volumes later if you prefer.
10. If you are attaching volumes now, select whether to create new volumes or select existing volumes to attach to the host.
11. Click **Continue**.
12. If you are creating new volumes:
 - a. Select the pool for the new volume and enter a **Volume Name**. Use a name that indicates how the volume is used, such as `{host name}_Host1_Vol1`.
 - b. Enter the **Volume Size** and select the units of measure. Optionally, you can choose to use the remaining space for the volume.
 - c. Click **Add Volume**.
Review the volume parameters. From this panel you can:
 - Delete the volume ()
 - **Add New Volume**
13. If you are using an existing volume, select the volume or volumes to attach to the host.
14. Click **Continue** to proceed.
The provisioning summary is displayed.

15. Review the provisioning configuration and click **Continue** to proceed, or **Back** to return to make changes to the provisioning.
16. Click **OK** at the Success prompt and return to the PowerVault Manager Dashboard.

Enable and configure DM Multipath on Linux hosts

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

About this task

 **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.

Steps

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 PowerVault ME5 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:

Steps

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`.

 **NOTE:** Run the `ls SCSI` 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`.

Configuring a Linux host with iSCSI network adapters

The following steps describe the end-to-end process for setting up hosts and provisioning volumes. This process can be done after the guided setup.

Prerequisites

- Ensure that the latest host operating system is installed and configured on the server.
- Ensure that all HBAs are installed and have the latest supported firmware and drivers as described on [Dell.com/support](https://www.dell.com/support). For a list of supported HBAs, see the *Dell ME5 Series Storage System Support Matrix* on the Dell support site.
- Cable the host servers as described in [Cable host servers to the storage system](#).
- Record the IP addresses assigned to each port as shown in the following example.

Table 8. Example worksheet for IP addresses

Port	IP Address
Subnet 1	
Host server 1, Port 0	192.68.10.20
Host server 2, Port 0	192.68.10.21
ME5 controller A port 0	192.68.10.200
ME5 controller A port 2	192.68.10.220
ME5 controller B port 0	192.68.10.205
ME5 controller B port 2	192.68.10.225
Subnet 2	
Host server 1, Port 1	192.68.11.20
Host server 2, Port 1	192.68.11.21
ME5 controller A port 1	192.68.11.210
ME5 controller A port 3	192.68.11.230
ME5 controller B port 1	192.68.11.215
ME5 controller B port 3	192.68.11.235

Assign IP addresses and configure iSCSI initiators

Follow the steps to assign IP addresses and configure the initiators for Linux host that connects to your storage system.

 **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.

Assign IP addresses for RHEL 7

Steps

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 to assign an IP address.
4. Change the IPv4 Configuration option to **Manual**.
5. Provide the subnet mask using the NIC IP address in the format x.x.x.x/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.

Configure RHEL 7 iSCSI initiators to connect to the storage system

Steps

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 192.68.10.200`
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.
 - 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.

Assign IP addresses for SLES 12

Steps

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 previously, enter the NIC IP address and subnet mask.
6. Select **Next**.
7. Ping the new network interface and associated storage host ports to ensure IP connectivity.
8. Repeat steps 1 through 7 for each NIC IP address.
9. Select **OK** to exit network settings.
10. Select **OK** to exit YaST.

Configure SLES 12 iSCSI initiators to connect to the storage system

Steps

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 you created earlier, enter the IP address for port A0 in the IP Address field, then click **Next**.
For example: `192.68.10.200`.
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.

Create a host and attach volumes in PowerVault Manager

If you did not set up hosts during the guided setup, or if you want to add new hosts, use the PowerVault Manager create hosts and attach volumes.

Steps

1. In the PowerVault Manager Dashboard, go to **Provisioning > Hosts**.
The Hosts panel opens with the **Hosts and Host Groups** table selected.
2. Click **Create Host**.
3. In the Create Host panel, select the **Create a New Host** radio button.
4. Enter a **Host Name**.
5. Select the initiator to assign to this host, using your worksheet as a guide to map the WWN and the Initiator ID.
6. (Optional) Enter a nickname for the this host initiator that clearly identifies the initiator for that particular host.
7. Click **Add Initiators To Host**.
The host is displayed in the **New Hosts** list.
8. Click **Continue**.
9. If you want to attach volumes now, select **Attach host or host groups to volumes**. You can skip this step and set up volumes later if you prefer.
10. If you are attaching volumes now, select whether to create new volumes or select existing volumes to attach to the host.
11. Click **Continue**.
12. If you are creating new volumes:
 - a. Select the pool for the new volume and enter a **Volume Name**. Use a name that indicates how the volume is used, such as `{host name}_Host1_Vol1`.
 - b. Enter the **Volume Size** and select the units of measure. Optionally, you can choose to use the remaining space for the volume.
 - c. Click **Add Volume**.Review the volume parameters. From this panel you can:
 - Delete the volume ()
 - **Add New Volume**
13. If you are using an existing volume, select the volume or volumes to attach to the host.
14. Click **Continue** to proceed.
The provisioning summary is displayed.
15. Review the provisioning configuration and click **Continue** to proceed, or **Back** to return to make changes to the provisioning.
16. Click **OK** at the Success prompt and return to the PowerVault Manager Dashboard.

Enable and configure DM Multipath on Linux hosts

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

About this task

 **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.

Steps

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 PowerVault ME5 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:

Steps

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`.
 **NOTE:** Run the `ls SCSI` 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`.

Configuring a SAS host server for Linux

The following steps describe the end-to-end process for setting up hosts and provisioning volumes. This process can be done after the guided setup.

Prerequisites

- Ensure that all HBAs are installed and have the latest supported firmware and drivers as described on [Dell.com/support](https://www.dell.com/support). For a list of supported HBAs, see the *Dell ME5 Series Storage System Support Matrix* on the Dell support site.
- Cable the host servers as described in [Cable host servers to the storage system](#).
- 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).

Identify SAS HBAs on a Linux host

Follow these steps to identify the SAS HBA initiators to connect to the storage system.

Steps

1. Open a terminal session.
2. Run the `dmesg | grep SCSI | grep slot` command.
3. Record the WWN numeric name.

Create a host and attach volumes in PowerVault Manager

If you did not set up hosts during the guided setup, or if you want to add new hosts, use the PowerVault Manager create hosts and attach volumes.

Steps

1. In the PowerVault Manager Dashboard, go to **Provisioning > Hosts**.
The Hosts panel opens with the **Hosts and Host Groups** table selected.
2. Click **Create Host**.
3. In the Create Host panel, select the **Create a New Host** radio button.

4. Enter a **Host Name**.
5. Select the initiator to assign to this host, using your worksheet as a guide to map the WWN and the Initiator ID.
6. (Optional) Enter a nickname for the this host initiator that clearly identifies the initiator for that particular host.
7. Click **Add Initiators To Host**.
The host is displayed in the **New Hosts** list.
8. Click **Continue**.
9. If you want to attach volumes now, select **Attach host or host groups to volumes**. You can skip this step and set up volumes later if you prefer.
10. If you are attaching volumes now, select whether to create new volumes or select existing volumes to attach to the host.
11. Click **Continue**.
12. If you are creating new volumes:
 - a. Select the pool for the new volume and enter a **Volume Name**. Use a name that indicates how the volume is used, such as *{host name}_Host1_Vol1*.
 - b. Enter the **Volume Size** and select the units of measure. Optionally, you can choose to use the remaining space for the volume.
 - c. Click **Add Volume**.
Review the volume parameters. From this panel you can:
 - Delete the volume ()
 - **Add New Volume**
13. If you are using an existing volume, select the volume or volumes to attach to the host.
14. Click **Continue** to proceed.
The provisioning summary is displayed.
15. Review the provisioning configuration and click **Continue** to proceed, or **Back** to return to make changes to the provisioning.
16. Click **OK** at the Success prompt and return to the PowerVault Manager Dashboard.

Enable and configure DM Multipath on Linux hosts

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

About this task

 **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.

Steps

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 PowerVault ME5 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:

Steps

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`.

NOTE: Run the `ls SCSI` 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`.

VMware ESXi hosts

ME5 Series storage systems support ESXi host servers using Fibre Channel, iSCSI, or SAS protocol.

Configuring a Fibre Channel host server for VMware ESXi

The following steps describe the end-to-end process for setting up hosts and provisioning volumes. This process can be done after the guided setup.

Prerequisites

- Ensure that all HBAs are installed and have the latest supported firmware and drivers as described on Dell.com/support. For a list of supported HBAs, see the *Dell ME5 Series Storage System Support Matrix* on the Dell support site.
- Cable the host servers as described in [Cable host servers to the storage system](#).
- Install the required version of the VMware ESXi operating system and configure it on the host.

Identify FC HBAs on an ESXi server

Perform the following steps to identify the Fibre Channel HBAs on an ESXi server.

Steps

1. Login to the vSphere Client.
2. Add the newly configured ESXi host to the appropriate datacenter and select it in the inventory list.
3. On the **Configure** tab, select **Storage > Storage Adapters**.
4. Verify that the required FC storage adapters are listed.
5. Select each adapter to see the WWN under **Properties** for the HBA. Record the WWN for each HBA.
6. 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 ME5 Series storage systems support single initiator/multiple target zones.

Create a host and attach volumes in PowerVault Manager

If you did not set up hosts during the guided setup, or if you want to add new hosts, use the PowerVault Manager create hosts and attach volumes.

Steps

1. In the PowerVault Manager Dashboard, go to **Provisioning > Hosts**.
The Hosts panel opens with the **Hosts and Host Groups** table selected.
2. Click **Create Host**.
3. In the Create Host panel, select the **Create a New Host** radio button.
4. Enter a **Host Name**.
5. Select the initiator to assign to this host, using your worksheet as a guide to map the WWN and the Initiator ID.
6. (Optional) Enter a nickname for the this host initiator that clearly identifies the initiator for that particular host.
7. Click **Add Initiators To Host**.
The host is displayed in the **New Hosts** list.
8. Click **Continue**.
9. If you want to attach volumes now, select **Attach host or host groups to volumes**. You can skip this step and set up volumes later if you prefer.
10. If you are attaching volumes now, select whether to create new volumes or select existing volumes to attach to the host.
11. Click **Continue**.
12. If you are creating new volumes:
 - a. Select the pool for the new volume and enter a **Volume Name**. Use a name that indicates how the volume is used, such as `{host name}_Host1_Vol1`.
 - b. Enter the **Volume Size** and select the units of measure. Optionally, you can choose to use the remaining space for the volume.
 - c. Click **Add Volume**.Review the volume parameters. From this panel you can:
 - Delete the volume ()
 - **Add New Volume**
13. If you are using an existing volume, select the volume or volumes to attach to the host.
14. Click **Continue** to proceed.
The provisioning summary is displayed.
15. Review the provisioning configuration and click **Continue** to proceed, or **Back** to return to make changes to the provisioning.
16. Click **OK** at the Success prompt and return to the PowerVault Manager Dashboard.

Rescan volume and create a datastore on the host

Use the vSphere Client to rescan storage and create a VMFS datastore.

Steps

1. Log in to the vSphere Client, and then select the ESXi host that you created.
2. On the Configure tab, select **Storage > Storage Adapters**.
3. Click **Rescan Storage**.
The **Rescan Storage** dialog box opens.
4. Click **OK**.
After a successful rescan, the volumes you created on the new host are visible in vCenter as a new disk or volume.
5. Create a VMware datastore file system on the ME5 Series volume.
Right-click on the ESXi host and select **Storage > New Datastore**. Continue through the wizard, using the following settings:
 - Datastore type: **VMFS**
 - Datastore name: Enter a name and then select the disk or LUN to use for provisioning the datastore.
 - Datastore version: **VMFS6**

- Partition configuration: select the default settings
6. Review the datastore information and click **Finish**.

Enable multipathing on an ESXi host

Use the vSphere Client to perform the following steps to enable multipathing on the ESXi host.

Steps

1. Log in to the vSphere Client, and then select the new datastore.
2. Select the **Configure** tab and click **Connectivity and Multipathing**.
3. Select the host to use for multipathing.
4. Under Multipathing Policies click **Actions** and select **Edit Multipathing**.
5. In the **Edit Multipathing Policies** window, select **Round Robin (VMware)** from the **Path selection policy** drop-down list.
6. Click **OK**.
7. Repeat this process for each volume that is presented from the ME5 Series Storage system to ESXi host.

 **NOTE:** For detailed steps and custom commands for MPIO see the *Dell PowerVault ME5 Series: VMware vSphere Best Practices* guide.

Configuring an ESXi host with an iSCSI network adapter

The following steps describe the end-to-end process for setting up hosts and provisioning volumes. This process can be done after the guided setup.

Prerequisites

- Ensure that the latest host operating system is installed and configured on the server.
- Ensure that all HBAs are installed and have the latest supported firmware and drivers as described on Dell.com/support. For a list of supported HBAs, see the *Dell ME5 Series Storage System Support Matrix* on the Dell support site.
- Cable the host servers as described in [Cable host servers to the storage system](#).
- Record the IP addresses assigned to each port as shown in the following example.

Table 9. Example worksheet for IP addresses

Port	IP Address
Subnet 1	
Host server 1, Port 0	192.68.10.20
Host server 2, Port 0	192.68.10.21
ME5 controller A port 0	192.68.10.200
ME5 controller A port 2	192.68.10.220
ME5 controller B port 0	192.68.10.205
ME5 controller B port 2	192.68.10.225
Subnet 2	
Host server 1, Port 1	192.68.11.20
Host server 2, Port 1	192.68.11.21
ME5 controller A port 1	192.68.11.210
ME5 controller A port 3	192.68.11.230
ME5 controller B port 1	192.68.11.215

Table 9. Example worksheet for IP addresses (continued)

Port	IP Address
ME5 controller B port 3	192.68.11.235

Configure the virtual switch

Perform the following steps to configure a virtual switch.

Steps

1. From the vSphere Client, click **Configure > Networking > Physical adapters**.
2. Locate and document the device name for the NICs used for iSCSI traffic.
3. Select **Virtual switches** and click **Add Networking** to launch the Add Networking wizard.
4. On the Select Connection Type page, select **VMkernel Network Adapter** and click **Next**.
5. On the Select Target Device page, select **New standard switch**, specify MTU as needed, and click **Next**.
6. On the Create Standard Switch page, click the plus (+) icon, to add an adapter.
7. Select the NIC that is on the first network to connect to the subnet defined previously. Click **OK** and then click **Next**.
8. In the Port properties, change the **Network label** to something that helps identify the purpose of the NIC, such as *iSCSI switch 1*.
9. On the IPv4 settings page, select **Use static IP settings** and assign the IP from your planning worksheet that corresponds to the port that will connect to this adapter.
10. Click **Next**.
11. On the Ready to complete page, review the settings and then click **Finish**.
The new Virtual switch is shown in the Virtual switches pane.
12. Repeat this process to configure the second NIC on the other subnet.

 **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

If iSCSI has not been configured, you may have to add a new software iSCSI adapter in vSphere.

About this task

 **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 ME5 Series storage system.

OR

- Run the following ESX CLI command on every ESXi host before connecting it to the ME5 Series storage system:

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

Steps

1. Log in to the vSphere Client.
2. On the **Configure** tab, select **Storage > Storage Adapters**.
3. Select the newly added iSCSI adapter, then select the **Dynamic Discovery** tab and click **Add**.
4. Enter the iSCSI IP address that is assigned to the iSCSI host port of storage controller A, then click **OK**.
5. Repeat steps 4-5 for the iSCSI host port of storage controller B.
6. 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 **Add** 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 were created in the Configure the VMware ESXi Vmkernel procedure, then click **OK**.
- c. Select **Rescan of storage adapters**.

Create a host and attach volumes in PowerVault Manager

If you did not set up hosts during the guided setup, or if you want to add new hosts, use the PowerVault Manager create hosts and attach volumes.

Steps

1. In the PowerVault Manager Dashboard, go to **Provisioning > Hosts**.
The Hosts panel opens with the **Hosts and Host Groups** table selected.
2. Click **Create Host**.
3. In the Create Host panel, select the **Create a New Host** radio button.
4. Enter a **Host Name**.
5. Select the initiator to assign to this host, using your worksheet as a guide to map the WWN and the Initiator ID.
6. (Optional) Enter a nickname for the this host initiator that clearly identifies the initiator for that particular host.
7. Click **Add Initiators To Host**.
The host is displayed in the **New Hosts** list.
8. Click **Continue**.
9. If you want to attach volumes now, select **Attach host or host groups to volumes**. You can skip this step and set up volumes later if you prefer.
10. If you are attaching volumes now, select whether to create new volumes or select existing volumes to attach to the host.
11. Click **Continue**.
12. If you are creating new volumes:
 - a. Select the pool for the new volume and enter a **Volume Name**. Use a name that indicates how the volume is used, such as `{host name}_Host1_Vol1`.
 - b. Enter the **Volume Size** and select the units of measure. Optionally, you can choose to use the remaining space for the volume.
 - c. Click **Add Volume**.

Review the volume parameters. From this panel you can:

 - Delete the volume ()
 - **Add New Volume**
13. If you are using an existing volume, select the volume or volumes to attach to the host.
14. Click **Continue** to proceed.
The provisioning summary is displayed.
15. Review the provisioning configuration and click **Continue** to proceed, or **Back** to return to make changes to the provisioning.
16. Click **OK** at the Success prompt and return to the PowerVault Manager Dashboard.

Rescan volume and create a datastore on the host

Use the vSphere Client to rescan storage and create a VMFS datastore.

Steps

1. Log in to the vSphere Client, and then select the ESXi host that you created.
2. On the Configure tab, select **Storage > Storage Adapters**.
3. Click **Rescan Storage**.
The **Rescan Storage** dialog box opens.
4. Click **OK**.

After a successful rescan, the volumes you created on the new host are visible in vCenter as a new disk or volume.

5. Create a VMware datastore file system on the ME5 Series volume.

Right-click on the ESXi host and select **Storage > New Datastore**. Continue through the wizard, using the following settings:

- Datastore type: **VMFS**
- Datastore name: Enter a name and then select the disk or LUN to use for provisioning the datastore.
- Datastore version: **VMFS6**
- Partition configuration: select the default settings

6. Review the datastore information and click **Finish**.

Enable multipathing on an ESXi host

Use the vSphere Client to perform the following steps to enable multipathing on the ESXi host.

Steps

1. Log in to the vSphere Client, and then select the new datastore.
2. Select the **Configure** tab and click **Connectivity and Multipathing**.
3. Select the host to use for multipathing.
4. Under Multipathing Policies click **Actions** and select **Edit Multipathing**.
5. In the **Edit Multipathing Policies** window, select **Round Robin (VMware)** from the **Path selection policy** drop-down list.
6. Click **OK**.
7. Repeat this process for each volume that is presented from the ME5 Series Storage system to ESXi host.

 **NOTE:** For detailed steps and custom commands for MPIO see the *Dell PowerVault ME5 Series: VMware vSphere Best Practices* guide.

Configuring a SAS host server for VMware ESXi

The following steps describe the end-to-end process for setting up hosts and provisioning volumes. This process can be done after the guided setup.

Prerequisites

- Ensure that all HBAs are installed and have the latest supported firmware and drivers as described on Dell.com/support. For a list of supported HBAs, see the *Dell ME5 Series Storage System Support Matrix* on the Dell support site.
- Cable the host servers as described in [Cable host servers to the storage system](#).
- Install the required version of the VMware ESXi operating system and configure it on the host.

Identify SAS HBAs on an ESXi server

Perform the following steps to identify the SAS HBAs on the ESXi server.

Steps

1. Log in to the ESXi host.
2. Click on **Storage** and then select the **Adapters** tab.
3. Record the WWN for the SAS HBA.

 **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.

Create a host and attach volumes in PowerVault Manager

If you did not set up hosts during the guided setup, or if you want to add new hosts, use the PowerVault Manager create hosts and attach volumes.

Steps

1. In the PowerVault Manager Dashboard, go to **Provisioning > Hosts**.
The Hosts panel opens with the **Hosts and Host Groups** table selected.
2. Click **Create Host**.
3. In the Create Host panel, select the **Create a New Host** radio button.
4. Enter a **Host Name**.
5. Select the initiator to assign to this host, using your worksheet as a guide to map the WWN and the Initiator ID.
6. (Optional) Enter a nickname for the this host initiator that clearly identifies the initiator for that particular host.
7. Click **Add Initiators To Host**.
The host is displayed in the **New Hosts** list.
8. Click **Continue**.
9. If you want to attach volumes now, select **Attach host or host groups to volumes**. You can skip this step and set up volumes later if you prefer.
10. If you are attaching volumes now, select whether to create new volumes or select existing volumes to attach to the host.
11. Click **Continue**.
12. If you are creating new volumes:
 - a. Select the pool for the new volume and enter a **Volume Name**. Use a name that indicates how the volume is used, such as `{host name}_Host1_Vol1`.
 - b. Enter the **Volume Size** and select the units of measure. Optionally, you can choose to use the remaining space for the volume.
 - c. Click **Add Volume**.Review the volume parameters. From this panel you can:
 - Delete the volume ()
 - **Add New Volume**
13. If you are using an existing volume, select the volume or volumes to attach to the host.
14. Click **Continue** to proceed.
The provisioning summary is displayed.
15. Review the provisioning configuration and click **Continue** to proceed, or **Back** to return to make changes to the provisioning.
16. Click **OK** at the Success prompt and return to the PowerVault Manager Dashboard.

Rescan volume and create a datastore on the host

Use the vSphere Client to rescan storage and create a VMFS datastore.

Steps

1. Log in to the vSphere Client, and then select the ESXi host that you created.
2. On the Configure tab, select **Storage > Storage Adapters**.
3. Click **Rescan Storage**.
The **Rescan Storage** dialog box opens.
4. Click **OK**.
After a successful rescan, the volumes you created on the new host are visible in vCenter as a new disk or volume.
5. Create a VMware datastore file system on the ME5 Series volume.
Right-click on the ESXi host and select **Storage > New Datastore**. Continue through the wizard, using the following settings:
 - Datastore type: **VMFS**
 - Datastore name: Enter a name and then select the disk or LUN to use for provisioning the datastore.
 - Datastore version: **VMFS6**

- Partition configuration: select the default settings
6. Review the datastore information and click **Finish**.

Enable multipathing on an ESXi host with SAS volumes

If you have more than one SAS HBA connection to each ME5 controller, perform the following steps to enable multipathing. If you have only one SAS HBA connection to each ME5 controller, you do not need to modify the multipathing policy.

Steps

1. Select the newly created datastore.
2. Select the **Configure** tab and click **Connectivity and Multipathing**.
3. Select the host to use for multipathing.
4. Under Multipathing Policies click **Actions** and select **Edit Multipathing**.
5. In the **Edit Multipathing Policies** window, select **Round Robin (VMware)** from the **Path selection policy** drop-down list.

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 recommends that you change the multipathing policy to **Round Robin (VMware)**.

6. Repeat this process for each SAS volume that is attached to the ESXi host.

NOTE: For detailed steps and custom commands for MPIO see the *Dell PowerVault ME5 Series: VMware vSphere Best Practices* guide.

Citrix XenServer hosts

ME5 Series storage systems support Citrix XenServer host servers using Fibre Channel, iSCSI, or SAS protocol.

Configuring a Fibre Channel host server for Citrix XenServer

The following steps describe the end-to-end process for setting up hosts and provisioning volumes. This process can be done after the guided setup.

Prerequisites

- Ensure that all HBAs are installed and have the latest supported firmware and drivers as described on Dell.com/support. For a list of supported HBAs, see the *Dell ME5 Series Storage System Support Matrix* on the Dell support site.
- Cable the host servers as described in [Cable host servers to the storage system](#).
- 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.

Identify FC HBAs on a XenServer

Perform the following steps to identify the FC HBAs on a XenServer.

Steps

1. Log in to the console for each XenServer host using SSH or XenCenter.
2. Use the following command to display and record the WWNs for the HBA ports that are connected to the storage system:

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

3. 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 ME5 Series storage systems support single initiator/multiple target zones.

Enable Multipathing on a XenServer

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

Steps

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** check box.
5. Click **OK**.
6. Right-click the host, and select **Exit Maintenance Mode**.
7. Repeat the previous steps for all the hosts in the pool.

Create a host and attach volumes in PowerVault Manager

If you did not set up hosts during the guided setup, or if you want to add new hosts, use the PowerVault Manager create hosts and attach volumes.

Steps

1. In the PowerVault Manager Dashboard, go to **Provisioning > Hosts**.
The Hosts panel opens with the **Hosts and Host Groups** table selected.
2. Click **Create Host**.
3. In the Create Host panel, select the **Create a New Host** radio button.
4. Enter a **Host Name**.
5. Select the initiator to assign to this host, using your worksheet as a guide to map the WWN and the Initiator ID.
6. (Optional) Enter a nickname for the this host initiator that clearly identifies the initiator for that particular host.
7. Click **Add Initiators To Host**.
The host is displayed in the **New Hosts** list.
8. Click **Continue**.
9. If you want to attach volumes now, select **Attach host or host groups to volumes**. You can skip this step and set up volumes later if you prefer.
10. If you are attaching volumes now, select whether to create new volumes or select existing volumes to attach to the host.
11. Click **Continue**.
12. If you are creating new volumes:
 - a. Select the pool for the new volume and enter a **Volume Name**. Use a name that indicates how the volume is used, such as `{host name}_Host1_Vol1`.
 - b. Enter the **Volume Size** and select the units of measure. Optionally, you can choose to use the remaining space for the volume.
 - c. Click **Add Volume**.

Review the volume parameters. From this panel you can:

- Delete the volume ()
- **Add New Volume**

13. If you are using an existing volume, select the volume or volumes to attach to the host.
14. Click **Continue** to proceed.
The provisioning summary is displayed.
15. Review the provisioning configuration and click **Continue** to proceed, or **Back** to return to make changes to the provisioning.
16. Click **OK** at the Success prompt and return to the PowerVault Manager Dashboard.

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

Steps

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.

Configuring an iSCSI host server for Citrix XenServer

The following steps describe the end-to-end process for setting up hosts and provisioning volumes. This process can be done after the guided setup.

Prerequisites

- Ensure that all HBAs are installed and have the latest supported firmware and drivers as described on Dell.com/support. For a list of supported HBAs, see the *Dell ME5 Series Storage System Support Matrix* on the Dell support site.
- Cable the host servers as described in [Cable host servers to the storage system](#).
- 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.

Identify iSCSI adapters on a XenServer

Perform the following steps to identify the iSCSI network adapters on the XenServer.

About this task

 **NOTE:** The Dell PowerVault ME5 Series storage system supports only software iSCSI adapters.

Steps

1. Record the two different IP address ranges for each storage system controller.
2. 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:

Steps

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:

Steps

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

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

Steps

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** check box.
5. Click **OK**.
6. Right-click the host, and select **Exit Maintenance Mode**.
7. Repeat the previous steps for all the hosts in the pool.

Create a host and attach volumes in PowerVault Manager

If you did not set up hosts during the guided setup, or if you want to add new hosts, use the PowerVault Manager create hosts and attach volumes.

Steps

1. In the PowerVault Manager Dashboard, go to **Provisioning > Hosts**.
The Hosts panel opens with the **Hosts and Host Groups** table selected.
2. Click **Create Host**.
3. In the Create Host panel, select the **Create a New Host** radio button.
4. Enter a **Host Name**.
5. Select the initiator to assign to this host, using your worksheet as a guide to map the WWN and the Initiator ID.
6. (Optional) Enter a nickname for the this host initiator that clearly identifies the initiator for that particular host.
7. Click **Add Initiators To Host**.
The host is displayed in the **New Hosts** list.
8. Click **Continue**.
9. If you want to attach volumes now, select **Attach host or host groups to volumes**. You can skip this step and set up volumes later if you prefer.
10. If you are attaching volumes now, select whether to create new volumes or select existing volumes to attach to the host.
11. Click **Continue**.
12. If you are creating new volumes:
 - a. Select the pool for the new volume and enter a **Volume Name**. Use a name that indicates how the volume is used, such as `{host name}_Host1_Vol1`.
 - b. Enter the **Volume Size** and select the units of measure. Optionally, you can choose to use the remaining space for the volume.
 - c. Click **Add Volume**.Review the volume parameters. From this panel you can:
 - Delete the volume ()
 - **Add New Volume**
13. If you are using an existing volume, select the volume or volumes to attach to the host.
14. Click **Continue** to proceed.
The provisioning summary is displayed.
15. Review the provisioning configuration and click **Continue** to proceed, or **Back** to return to make changes to the provisioning.
16. Click **OK** at the Success prompt and return to the PowerVault Manager Dashboard.

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:

Steps

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.

Configuring a SAS host for Citrix XenServer

The following steps describe the end-to-end process for setting up hosts and provisioning volumes. This process can be done after the guided setup.

Prerequisites

- Ensure that all HBAs are installed and have the latest supported firmware and drivers as described on Dell.com/support. For a list of supported HBAs, see the *Dell ME5 Series Storage System Support Matrix* on the Dell support site.
- Cable the host servers as described in [Cable host servers to the storage system](#).
- 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.

Identify SAS HBAs on a XenServer

Perform the following steps to identify SAS HBAs on a XenServer.

Steps

1. Log in to the console for each XenServer host using SSH or XenCenter.

2. Use the following command to display and record the initiator ID for the HBA ports that are connected to the storage enclosure:

```
systool -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

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

Steps

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** check box.
5. Click **OK**.
6. Right-click the host, and select **Exit Maintenance Mode**.
7. Repeat the previous steps for all the hosts in the pool.

Create a host and attach volumes in PowerVault Manager

If you did not set up hosts during the guided setup, or if you want to add new hosts, use the PowerVault Manager create hosts and attach volumes.

Steps

1. In the PowerVault Manager Dashboard, go to **Provisioning > Hosts**.
The Hosts panel opens with the **Hosts and Host Groups** table selected.
2. Click **Create Host**.
3. In the Create Host panel, select the **Create a New Host** radio button.
4. Enter a **Host Name**.
5. Select the initiator to assign to this host, using your worksheet as a guide to map the WWN and the Initiator ID.
6. (Optional) Enter a nickname for the this host initiator that clearly identifies the initiator for that particular host.
7. Click **Add Initiators To Host**.
The host is displayed in the **New Hosts** list.
8. Click **Continue**.
9. If you want to attach volumes now, select **Attach host or host groups to volumes**. You can skip this step and set up volumes later if you prefer.
10. If you are attaching volumes now, select whether to create new volumes or select existing volumes to attach to the host.
11. Click **Continue**.
12. If you are creating new volumes:
 - a. Select the pool for the new volume and enter a **Volume Name**. Use a name that indicates how the volume is used, such as *{host name}_Host1_Vol1*.
 - b. Enter the **Volume Size** and select the units of measure. Optionally, you can choose to use the remaining space for the volume.
 - c. Click **Add Volume**.Review the volume parameters. From this panel you can:
 - Delete the volume ()
 - **Add New Volume**
13. If you are using an existing volume, select the volume or volumes to attach to the host.

14. Click **Continue** to proceed.
The provisioning summary is displayed.
15. Review the provisioning configuration and click **Continue** to proceed, or **Back** to return to make changes to the provisioning.
16. Click **OK** at the Success prompt and return to the PowerVault Manager Dashboard.

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:

- [Fault isolation methodology](#)
- [2U enclosure LEDs](#)
- [5U84 enclosure LEDs](#)
- [Initial start-up problems](#)

Fault isolation methodology

ME5 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.

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 PowerVault ME5 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. In addition, see the following topics for troubleshooting guidance:

- [Options available for performing basic steps](#)
- [Performing basic steps](#)
- [Host I/O](#)

Options available for performing basic steps

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

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

Use the PowerVault Manager to monitor the health of the system and its components. If any component has a problem, PowerVault Manager shows the system health as Degraded, Fault, or Unknown. Use the PowerVault Manager to find failed or unhealthy components and 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 CLI shows the system health as Degraded, Fault, or

Unknown, 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.

Performing basic steps

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

Gather fault information

When a fault occurs, gather as much information as possible. Doing so helps determine the correct action that is required 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 or if the LEDs cannot be viewed due to the location of the system. The **Maintenance > Hardware** view provides a visual representation of the system and shows faults when they occur. 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 identifies the severity:

- 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 has 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 faults and the cause of the failure. For example, a host could lose connectivity to a disk group if a user changes the channel settings without considering the storage resources that are assigned. 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 because of the data paths and the number of components in the data path. For example, any of the components in the data path could cause a host-side data error: controller module, cable, or data host.

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.

2U enclosure LEDs

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

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.




Figure 32. Ops panel LEDs—2U enclosure front panel

Table 10. Ops panel functions—2U enclosure front panel

Indicator	Description	Color	Status
	System power	Green	On steady: at least one PCM is supplying power Off: system not operating regardless of AC present
	Status/Health	Blue	On steady: system is powered on and controller is ready Blinking (2 Hz): Enclosure management is busy (for example, when booting or performing a firmware update)
		Amber	On steady: module fault present (may be associated with a Fault LED on a controller module, IOM, or PCM) Blinking: logical fault (2 s on, 1 s off)
	Unit identification display (UID)	Green	Dual seven-segment display that shows the numerical position of the enclosure in the cabling sequence. The UID is also called the enclosure ID. The controller enclosure ID is 0.
	Identity	Blue	Blinking (0.25 Hz): system ID locator is activated to assist in locating the chassis within a data center. Off: Normal state

2U enclosure PCM LEDs

Under normal conditions, the power cooling module (PCM) OK LEDs will be a constant green.

Table 11. 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
Off	On	On	On	PCM fault (above temperature, above voltage, above current)
Off	Blinking	Blinking	Blinking	PCM firmware download is in progress

2U enclosure Ops panel LEDs

The Ops panel displays the aggregated status of all the modules. The following table describes the Ops panel LED states.

Table 12. Ops panel LED states

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

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 an IOM CRU:
 - Remove and reinstall the IOM.
 - 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 IOM and reinsert it.

2U enclosure disk drive carrier module LEDs

Disk drive status is monitored by a green LED and an amber LED mounted on the front of each drive carrier module, as shown in the following figure.

The drive module LEDs are identified in the figure, and the LED behavior is described in the table following the figure.

- In normal operation, the green LED are on, and flicker as the drive operates.
- In normal operation the amber LED will be:
 - Off if there is no drive present.
 - Off as the drive operates.
 - On if there is a drive fault.




Figure 33. 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 |

Table 13. Drive carrier LED states

Activity LED (Green)	Fault LED (Amber)	Status/condition*
Off	Off	Off (disk module/enclosure)
Off	Off	Not present
Blink off with activity	Blinking: 1s on /1s off	Identify
● 1 down: Blink with activity ● 2 down: Off	On	Drive link (PHY lane) down
On	On	Fault (leftover/failed/locked-out)
Blink off with activity	Off	Available
Blink off with activity	Off	Storage system: Initializing
Blink off with activity	Off	Storage system: Fault-tolerant
Blink off with activity	Off	Storage system: Degraded (not critical)
Blink off with activity	Blinking: 3s on/ 1s off	Storage system: Degraded (critical)
On	Off	Storage system: Quarantined
Blink off with activity	Blinking: 3s on/ 1s off	Storage system: Offline (dequarantined)
Blink off with activity	Off	Storage system: Reconstruction
Blink off with activity	Off	Processing I/O (whether from host or internal activity)

Table 13. Drive carrier LED states (continued)

Activity LED (Green)	Fault LED (Amber)	Status/condition*
*If multiple conditions occur simultaneously, the LED state behaves as indicated by the condition listed earliest in the table, as rows are read from top to bottom.		

IO Module LEDs

IOM status is monitored by the LEDs located on the face plate. The following table describes LED behaviors for expansion enclosure IOMs.

Table 14. Expansion enclosure IOM LEDs

LED	Description	Color	Status
	Module fault	Amber	On - Ops panel undergoing 5s test - Rear panel area module fault: IOM, fan, PSU, when paired with module fault LED - Drive module hardware fault, paired with drive fault LED
			Flashing - Unknown, invalid, or mixed module type, such as drive module or PSU - Vital product data (VPD) configuration error or I2C bus failure
			Off—IOM functioning properly
	Power on or standby	Green	On—IOM power is on
		Amber	On—Part of standby sequence as IOM comes online
		None	Off—IOM power is off
	Unit identification (UID)	White	On—UID active to locate or identify for service activity
			Off—UID not active
	12 Gb/s SAS port	Green	On—Connected, link is up
			Off—Not connected or link is down
		Amber	On—Critical SAS cable fault
			Fast flash (1s:1s)—SAS UID active
			Slow flash (3s:1s)—Non-critical SAS cable fault
	Ethernet port	---	Off—SAS expansion port functioning properly
			Ethernet port is disabled

12 Gb/s controller module LEDs

The diagrams with tables that immediately follow provide descriptions for the different controller modules that can be installed into the rear panel of the controller enclosures.

 **NOTE:** Consider the following when viewing the controller module diagrams on the following pages:

- In each diagram, the controller module is oriented for insertion into the top slot (A) of 2U enclosures. When oriented for use in the bottom slot (B) of 2U enclosures, the controller module labels appear upside down.
- In each diagram, the controller module is oriented for insertion into either slot of 5U84 enclosures.
- Alternatively, you can configure the 2U controller enclosure with a single controller module. Install the controller module in slot A, and install a blank plate in slot B.




Figure 34. ME5 Series controller module

Table 15. Common controller module LEDs

LED	Description	Color	Status
✓	Hardware normal	Green	On—Controller module is functioning properly
			Flashing—Part of standby sequence as the controller module comes online
			Off—Controller module power is off, controller module is offline, or controller module has a fault condition
⚠	Hardware fault	Amber	On—Controller module hardware fault
			Off—Controller module functioning properly
🗑	OK to remove	White	On—Ready for removal, the cache is clear
			Off—Do not remove the controller module, cache still contains unwritten data
🔍	Identify	Blue	On—Unit identification (UID) is active
			Off—Normal state, UID is not active
CACHE	Cache status	Green	On—Cache contains unwritten data, controller module is functioning properly
			Fast flash (1s:1s)—Cache is active, cache flush in progress
			Slow flash (3s: 1s)—Cache self-refresh in progress after cache flush
			Off—Cache is clear or system is coming online
🔗	Ethernet management port speed	Amber	On—1000Base-T negotiated rate
			Off—10/100Base-T negotiated rate
ACT	Ethernet management port link activity	Green	On—Ethernet link is up
			Off—Ethernet link is down
🔗	12 Gb/s SAS expansion port status	Green	On—Connected, link is up
		Green or amber	Flashing—Link activity
		Amber	On—Connected, partial link is up
		None	Off—Not connected or link is down

The following figure shows the host port LEDs on a 32Gb/s Fibre Channel controller module:




Figure 35. 32Gb/s Fibre Channel ports

LED	Description	Color	Status
	Fibre Channel link activity	Green	On—Connected, link is up
			Flashing—Link activity
			Off—Not connected or link is down

The following figure shows the host port LEDs on a 25 GbE iSCSI controller module:




Figure 36. 25 GbE iSCSI ports

LED	Description	Color	Status
	iSCSI link activity	Green	On—Connected, link is up
			Flashing—Link activity
			Off—Not connected or link is down

The following figure shows host port LEDs on a 10Gbase-T iSCSI controller module:




Figure 37. 10Gbase-T iSCSI ports

LED	Description	Color	Status
	iSCSI 10Gbase-T link speed	Green	On—10GbE link speed
		Amber	On—1GbE link speed
		None	Off—Not connected or link is down
ACT	iSCSI 10Gbase-T link activity	Green	On—Connected, link is up
			Flashing—Link activity
			Off—Not connected or link is down

The following figure shows the host port LEDs on a 12 Gb/s SAS controller module:



Figure 38. 12 Gb/s SAS ports

LED	Description	Color	Status
	SAS port status	Green	On—Connected, link is up
		Green or amber	Flashing—Link activity
		Amber	On—Connected, partial link is up
		None	Off—Not connected or link is down

5U84 enclosure LEDs

When the 5U84 enclosure is powered on, all LEDs turn on for a short period to ensure that they are working.

NOTE: This behavior does not indicate a fault unless LEDs remain lit after several seconds.

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.




Figure 39. Ops panel LEDs—5U enclosure front panel

Table 16. Ops panel functions – 5U enclosure front panel

Indicator	Description	Color	Status
	Unit Identification Display (UID)	Green	Dual seven-segment display that shows the numerical position of the enclosure in the cabling sequence. The UID is also called the enclosure ID. The controller enclosure ID is 0.
	System Power On/Standby	Green	On steady: system power is available (operational)
		Amber	Constant amber: system in standby (not operational)
	Module Fault	Amber	Constant or blinking: system hardware 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	Constant or blinking: fault present 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.
	Top Drawer Fault	Amber	Constant or blinking: fault present in drive, cable, or sideplane (drawer 0)
	Bottom Drawer Fault	Amber	Constant or blinking: fault present in drive, cable, or sideplane (drawer 10)

ME5084 PSU LEDs

The following table describes the LED states for the PSU:

Table 17. 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. (usually 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. .

ME5084 FCM LEDs

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

Table 18. FCM LED states

LED	Status/description
Module OK	Constant green indicates that the FCM is working correctly. Off indicates the fan module has failed.
Fan Fault	Amber indicates the fan module has failed. .

ME5084 Ops panel LEDs

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

Table 19. Ops panel LED states

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.
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 ME5084 drawer LEDs .
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.

 **CAUTION:** The sideplanes on the enclosure drawers are not hot swappable or customer serviceable.

ME5084 drawer LEDs


The following table describes the LEDs on the drawers:

Table 20. Drawer LED states

LED	Status/description
Sideplane OK/Power Good	Green if the sideplane 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 will light amber.
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.
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).

ME5084 DDIC LEDs

The DDIC supports LFF 3.5" and SFF 2.5" disks. The following figure shows the top panel of the DDIC as viewed when the disk is aligned for insertion into a drawer slot.

**Figure 40. 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 21. DDIC LED states

Fault LED (Amber)	Status/description*
Off	Off (disk module/enclosure)
Off	Not present
Blinking: 1s on/1s 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: 3s on/1s off	Storage system: Degraded (critical)
Off	Storage system: Quarantined
Blinking: 3s on/1s 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 will behave as indicated by the condition listed earliest in the table, as rows are read from top to bottom.

Each DDIC has a single Drive Fault LED. A disk drive fault is indicated if the Drive Fault LED is lit amber. In the event of a disk failure, replace the DDIC.

5U84 controller module and IOM LEDs

Controller module and IOM CRUs are common to the 2U and 5U84 enclosures.

- For information about controller module LEDs, see [12 Gb/s controller module LEDs](#).
- For information about IOM LEDs, see [IO Module LEDs](#).

Initial start-up problems

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

Troubleshooting 2U enclosures

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

The Module Fault LED on the Ops panel lights amber to indicate a fault for the problems listed in the following table:

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

Table 22. 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

Troubleshooting 5U enclosures

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

The Module Fault LED on the Ops panel lights amber to indicate a fault for the problems listed in the following table:

NOTE: All alarms also report through SES.

Table 23. 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
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

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.

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.

About this task

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.

Steps

1. To perform a rescan using the PowerVault Manager:
 - a. Verify that both controllers are operating normally.
 - b. Select **Maintenance > Hardware**.
 - c. Select **Actions > Rescan All Disks**
 - d. Click **Rescan**.
2. To perform a rescan using the CLI, type the following command:
`rescan`

Troubleshooting hardware faults

Make sure that you have a replacement module of the same type before removing any faulty module.

-  **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. 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 10Gbase-T and SAS host ports

The following procedure applies to ME5 Series controller enclosures employing external connectors in the host interface ports.

About this task

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

Steps

1. Halt all I/O to the storage system.
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 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.

About this task

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.

Steps

1. Halt all I/O to the storage system.

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 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

This section describes how to connect storage systems for replication and shows cabling examples.

Topics:

- [Connecting two storage systems to replicate volumes](#)
- [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.

 **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 *ME5 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.

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.
- 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 systems are physically cabled, see the *Dell PowerVault ME5 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 ME5 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 41. Replication cabling — single controller enclosures, 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 ME5 Series controller enclosures that are equipped with dual-controller modules for replication.

Co-located replication with multiple servers, one switch, one network

In cases where you use a single switch connected to a single host server that is logically functioning as multiple servers, an example of optimal cabling dedicates a pair of cables connected to each controller for I/O traffic and the other pair connected to each controller for replication traffic. In the illustration, green cables show replication traffic and blue cables show I/O traffic.

Sample cabling for the first controller enclosure:

- Two SFP I/O cables connect Controller 0A, and two connect Controller 0B to the switch.
- Two SFP replication cables connect Controller 0A, and two connect Controller 0B to the switch.

Sample cabling for the second controller enclosure:

- Two SFP I/O cables connect Controller 0A, and two connect Controller 0B to the switch.
- Two SFP replication cables connect Controller 0A, and two connect Controller 0B to the switch.




Figure 42. Replication cabling — multiple servers, one switch, and one network

1. 2U controller enclosures
2. Switch (I/O, replication)
3. Connection to host servers

Co-located replication with multiple servers and switches

Ideally, use three separate switches to avoid a single point of failure and allow for physical isolation of I/O traffic from replication traffic. Two switches are dedicated to I/O traffic and act as the bridge to connect controller enclosures to multiple host servers. The third switch is the replication switch and acts as the bridge to connect controller enclosures to each other. In the illustration, green cables show replication traffic and blue cables show I/O traffic.

Sample cabling for the first controller enclosure:

- Two SFP I/O cables connect Controller OA, and two connect Controller OB to the left I/O switch
- Two SFP replication cables connect Controller OA, and two connect Controller OB to the center replication switch.

Sample cabling for the second controller enclosure:

- Two SFP I/O cables connect Controller OA, and two connect Controller OB to the right I/O switch.
- Two SFP replication cables connect Controller OA, and two connect Controller OB to the center replication switch.




Figure 43. Replication cabling — multiple servers, multiple switches, one network

1. Controller enclosure modules
2. Switches (I/O)

Remote replication

Volume replication can occur on the same physical network, or on different physical networks.

For situations where you need to replicate volumes in two physically remote locations, you must still isolate input and output(I/O) traffic from replication traffic. In such a case, host servers are on separate networks, each at a separate site, and use a shared wide area network (WAN). The Ethernet WAN might be located at one of the sites, or connect to the Cloud. Ideally, use two switches, one at each remote location, and an Ethernet WAN to avoid a single point of failure and allow for physical isolation of I/O traffic from replication traffic. Both switches are dedicated to I/O traffic. Each one acts as the bridge to connect the site controller enclosures to the site host server. The Ethernet WAN acts as the bridge to connect the controller enclosures to each other.

Sample cabling for the Site 1 controller enclosure and switch:

- Two SFP I/O cables connect Site 1 Controller OA, and two connect Controller OB to the Site 1 switch.
- Two SFP replication cables connect Controller OA, and two connect Controller OB to the Ethernet WAN.

Sample cabling for the Site 2 controller enclosure and switch:

- Two SFP I/O cables connect Site 2 Controller OA, and two connect Controller OB to the Site 2 switch.
- Two SFP replication cables connect Controller OA, and two connect Controller OB to the Ethernet WAN.
- 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.




Figure 44. Replication cabling — 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 | |

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 ME5 Series Storage System controller enclosures for replication supplements the troubleshooting procedures that follow.

- Familiarize yourself with replication content provided in the *Dell PowerVault ME5 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:
 - Find the port address on the secondary system:
Using the CLI, run the `show ports` command on the secondary system.
 - 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 go to **Settings > Peer Connection**.
 - Create a peer connection.
To create a peer connection, use the `create peer-connection` CLI command or in the PowerVault Manager go to **Settings > Peer Connection**.
 - Create a virtual replication set:
 - Use the `create replication-set` CLI command, or
 - In the PowerVault Manager go to **Provisioning > Volumes** and select a Volume. Then select **Add Data Protection** and follow the setup wizard to complete the replication configuration.
 - Initiate a replication:
 - Use the `replicate` CLI command, or
 - In the PowerVault Manager go to **Provisioning > Volumes**, select a Volume, and then select **Data Protection**. From there you can start, suspend, or remove a replication.
 - 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 PowerVault ME5 Series Series Storage System Administrator's Guide* for procedures to set up and manage replications.
 - See the *Dell PowerVault ME5 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 24. 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	Perform the following actions on each system used for virtual replication:

Table 24. Diagnostics for replication setup: Using the replication feature (continued)

Answer	Possible reasons	Action
	feature is not running on each system that is used for replication.	- In the PowerVault Manager Dashboard, select Maintenance > Firmware. The Firmware panel opens, showing the 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 PowerVault ME5 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: - Verify use of proper cables. - Verify proper cabling paths for host connections. - 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: Go to **Provisioning > Volumes**, select a volume and then click **Add Data Protection**. Follow the setup wizard to connect to another system and set up a replication schedule.

Table 25. 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 CHAP, see the topics about configuring CHAP and working in replications within the <i>Dell PowerVault ME5 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).	- 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: - A conflicting volume does not exist. - Available free space in the pool.
No	Communication link is down.	Review the Alerts and Activity in the PowerVault Manager Dashboard for indicators of a specific fault in a host or replication data path component.

Can you replicate a volume?

Table 26. Diagnostics for replication setup – Replicating a volume

Answer	Possible reasons	Action
Yes	System functioning properly.	No action required.
No	Nonexistent .	- Determine existence of primary or secondary volumes. - If a replication set has not been successfully created, go to Provisioning > Volumes, select a volume and then click Add Data Protection. Follow the setup wizard to connect to another system and set up a replication schedule.

Table 26. Diagnostics for replication setup – Replicating a volume (continued)

Answer	Possible reasons	Action
		Review the Alerts and Activity in the PowerVault Manager Dashboard for indicators of a specific fault in a replication data path component. Follow any Recommended Actions.
No	Network error occurred during in-progress replication.	- Review event logs for indicators of a specific fault in a replication data path component. Follow any Recommended Actions. - Go to Provisioning > Volumes and select the Data Protection table to display replications and associated metadata. - Replications that enter the suspended state can be resumed manually (see the <i>Dell PowerVault ME5 Series Storage System Administrator's Guide</i> for additional information).
No	Communication link is down.	Review the Alerts and Activity in the PowerVault Manager Dashboard for indicators of a specific fault in a host or replication data path component.

Has a replication run successfully?

Table 27. 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	- Go to Provisioning > Volumes, and select the volume that is a member of the replication set. - Select the Data Protection 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 PowerVault ME5 Series Storage System Administrator's Guide</i>.
No	Communication link is down	Review the Alerts and Activity in the PowerVault Manager Dashboard 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 that are shipped with ME5 Series FC or iSCSI controller enclosures.

Locate the SFP transceivers

Locate the SFP transceivers that shipped with the controller enclosure. For more information about whether a transceiver is required, see:

- [32Gb Fibre Channel host connection](#)
- [25 GbE iSCSI host connection](#)

 **NOTE:** See 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](#) 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.
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 the SFP installation procedure.

System Information Worksheet

Use the system information worksheet to record the information that is needed to install the ME5 Series Storage System.

ME5 Series Storage System information

Gather and record the following information about the ME5 Series Storage System network and the administrator user:

Table 28. ME5 Series Storage System management network

Item	Information
Service tag	
Management IPv4 address (ME5 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 29. ME5 Series Storage System administrator

Item	Information
Password for the default ME5 Series Storage System Admin user	
Email address of the default ME5 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.

NOTE: For a storage system deployed with two Ethernet switches, Dell recommends setting up separate subnets.

Table 30. iSCSI Subnet 1

Item	Information
Subnet mask	----- . ----- . ----- . -----

Table 30. 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 31. 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 ME5 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 wizard.

Table 32. 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 33. 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 34. 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

You can connect directly to the controller module using a micro USB port and set network addresses using the CLI.

Topics:

- [Set a network port IP address using the micro USB port](#)
- [Micro-USB device connection](#)

Set a network port IP address using the micro USB port

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.

About this task

If you did not use DHCP to set network port IP address, you can set them manually using the CLI port. You can use a generic micro-USB cable and the USB CLI port. If you plan on using a micro-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.

 **NOTE:** To connect to the micro USB port on a controller module, see [Micro-USB device connection](#).

 **NOTE:** If you are using a host computer running Linux, prepare the USB port as described in [Linux drivers](#).

Use the CLI commands described in the following steps to set the IP address for the network port on each controller module:

 **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.

Steps

1. From your network administrator, obtain an IP address, subnet mask, and gateway address for controller A and another for controller B.
2. Connect a micro-USB cable from a host computer to the USB CLI port on controller A.
3. Start a terminal emulator and configure it to use the display and connection settings shown in the following tables.

Table 35. Terminal emulator display settings

Parameter	Value
Terminal emulation mode	VT-100 or ANSI (for color support)

Table 35. Terminal emulator display settings (continued)

Parameter	Value
Font	Terminal
Translations	None
Columns	80

Table 36. 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 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.
 - c. Type **Y** at the prompt to continue.
 If you are connecting to a storage system 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. Set the network port using either DHCP or set a static address using IPv4.
 - a. To use DHCP to set network port IP addresses, type the following command at the prompt:

```
set network-parameters dhcp
```

- b. To use custom static IP IPv4 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
```

NOTE: See the CLI Reference Guide for information about IPv6 and the commands used to add IPv6 addresses and set IPv6 network parameters. The `ipv6` term is included within each relevant command name

7. Type the following CLI command to verify the new IP addresses:

For IPv4: `show network-parameters`

For IPv6: `show ipv6-network-parameters`

The network parameters, including the IP address, subnet mask, and gateway address are displayed for each controller module.

8. Use the CLI ping command to verify connectivity to the gateway address.
For example:

```
ping 192.168.0.1
```

9. 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
```

10. Record the IP address for the controller modules to use when connecting to the storage system using PowerVault Manager.
11. When you are done using the CLI, close the terminal emulator.

Micro-USB device connection

The following sections describe the connection to the micro-USB port:

Emulated serial port

When a computer is connected to a controller module using a micro-USB 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](#).

Supported host applications

The following terminal emulator applications can be used to communicate with an ME5 Series controller module:

Table 37. 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 micro-USB port is considered an out-of-band connection. The connection to the micro-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 38. USB identification code

USB identification code type	Code
USB Vendor ID	0x210C
USB Product ID	0xA4A7

Microsoft Windows drivers

Windows Server 2016 and later operating systems provide a native USB serial driver that supports the micro-USB port.

Linux drivers

Linux operating systems do not require the installation of an ME5 Series USB driver. However, certain parameters must be provided during driver loading to enable recognition of the micro-USB port on an ME5 Series controller module.

Type the following command to load the Linux device driver with the parameters that are required to recognize the micro-USB port:

```
# modprobe usbserial vendor=0x210c product=0xa4a7 use_acm=1
```

 **NOTE:** Optionally, this information can be incorporated into the `/etc/modules.conf` file.

Technical specifications

Enclosure dimensions

Table 39. 2U12 and 2U24 enclosure dimensions

Specification	mm	inches
Height	87.9 mm	3.46 in
Width	483 mm	19.01 in
Depth (2U12)	618.7 mm	24.36 in
Depth (2U24)	547.8 mm	21.56 in

NOTE:

- The 2U12 enclosure uses 3.5" LFF disks.
- The 2U24 enclosure uses 2.5" SFF disks.

Table 40. 5U84 enclosure dimensions

Specification	mm	inches
Height	222.3 mm	8.75 in
Width	483 mm	19.01 in
Depth	981 mm	38.62 in

NOTE: The 5U84 uses 3.5" LFF disks in the DDIC carrier. It can also use 2.5" SFF disks with 3.5" adapter in the DDIC.

Enclosure weights

Table 41. 2U12, 2U24, and 5U84 enclosure weights

CRU/component	2U12 (kg/lb)	2U24 (kg/lb)	5U84 (kg/lb)
Storage enclosure (empty)	4.8/10.56	4.8/10.56	64/141
Disk drive carrier	0.9/1.98	0.3/0.66	0.8/1.8
Blank disk drive carrier	0.05/0.11	0.05/0.11	—
Power Cooling Module (PCM)	3.5/7.7	3.5/7.7	—
Power Supply Unit (PSU)	—	—	2.7/6
Fan Cooling Module (FCM)	—	—	1.4/3
SBB controller module (maximum weight)	2.6/5.8	2.6/5.8	2.6/5.8
SBB expansion module	1.5/3.3	1.5/3.3	1.5/3.3
RBOD enclosure (fully populated with modules: maximum weight)	32/71	30/66	135/298

Table 41. 2U12, 2U24, and 5U84 enclosure weights (continued)

CRU/component	2U12 (kg/lb)	2U24 (kg/lb)	5U84 (kg/lb)
EBOD enclosure (fully populated with modules: maximum weight)	28/62	25/55	130/287

i NOTE:

- Weights shown are nominal, and subject to variances.
- Weights may vary due to different controller modules, IOMs, and power supplies; and differing calibrations between scales.
- Weights may also vary due to the number and type of disk drives (SAS or SSD) installed.

Environmental requirements

Table 42. Ambient temperature and humidity

Specification	Temperature range	Relative humidity	Max. Wet Bulb
Operating	- RBOD: 5°C to 35°C (41°F to 95°F) - EBOD: 5°C to 40°C (41°F to 104°F)	20% to 80% non-condensing	28°C
Non-operating (shipping)	-40°C to +70°C (-40°F to +158°F)	5% to 100% non-precipitating	29°C

Table 43. Additional environmental requirements

Specification	Measurement/description
Airflow	- System must be operated with low pressure rear exhaust installation. - Back pressure created by rack doors and obstacles not to exceed 5Pa (~0.5 mm H₂O)
Altitude, operating	2U enclosures: 0 to 3,000 meters (0 to 10,000 feet) - Maximum operating temperature is de-rated by 5°C above 2,133 meters (7,000 feet)
	5U84 enclosures: -100 to 3,000 meters (-330 to 10,000 feet) - Maximum operating temperature is de-rated by 1°C above 900 meters (3,000 feet)
Altitude, non-operating	-100 to 12,192 meters (-330 to 40,000 feet)
Shock, operating	5.0 g, 10 ms, ½ sine pulses, Y-axis
Shock, non-operating	2U enclosures: 30.0 g, 10 ms, ½ sine pulses
	5U84 enclosures: 30.0 g, 10 ms, ½ sine pulses (Z-axis); 20.0 g, 10 ms, ½ sine pulses (X- and Y-axes)
Vibration, operating	0.21 G _{rms} 5 Hz to 500 Hz random
Vibration, non-operating	1.04 G _{rms} 2 Hz to 200 Hz random
Vibration, relocation	0.3 G _{rms} 2 Hz to 200 Hz 0.4 decades per minute
Acoustics	Operating sound power 2U enclosures: ≤ L_{WAd} 6.6 Bels (re 1 pW) @ 23°C 5U84 enclosures: ≤ L_{WAd} 8.0 Bels (re 1 pW) @ 23°C
Orientation and mounting	19" rack mount (2 EIA units; 5 EIA units)

Power cooling module

Specifications for the PCM are provided in the following table.

Table 44. 2U Power cooling module specifications

Specification	Measurement/description	
Dimensions (size)	84.3 mm high x 104.5 mm wide x 340.8 mm long - X-axis length: 104.5 mm (4.11 in) - Y-axis length: 84.3 mm (3.32 in) - Z-axis length: 340.8 mm (37.03)	
Maximum output power	580 W	
Voltage range	100–200 VAC rated	
Frequency	50/60 Hz	
Voltage range selection	Auto-ranging: 90–264 VAC, 47–63 Hz	
Maximum inrush current	20A	
Power factor correction	≥ 95% @ nominal input voltage	
Efficiency	115 VAC/60 Hz	230 VAC/50 Hz
	> 80% @ 10% load	> 80% @ 10% load
	> 87% @ 20% load	> 88% @ 20% load
	> 90% @ 50% load	> 92% @ 50% load
	> 87% @ 100% load	> 88% @ 100% load
	> 85% @ surge	> 85% @ surge
Harmonics	Meets EN61000-3-2	
Output	+5 V @ 42A, +12 V @ 38A, +5 V standby voltage @ 2.7A	
Operating temperature	0°C to 57°C (32°F to +135°F)	
Hot pluggable	Yes	
Switches and LEDs	AC mains switch and four status indicator LEDs	
Enclosure cooling	Dual axial cooling fans with variable fan speed control	

Power supply unit

Table 45. 5U84 Power supply unit specifications

Specification	Measurement/description
Maximum output power	2,214 W maximum continuous output power at high line voltage
Voltage	+12 V at 183 A (2,196 W) +5 V standby voltage at 2.7 A
Voltage range	200–240 VAC rated
Frequency	50/60 Hz
Power factor correction	≥ 95% @ 100% load
Efficiency	> 82% @ 10% load > 90% @ 20% load > 94% @ 50% load

Table 45. 5U84 Power supply unit specifications (continued)

Specification	Measurement/description
	• > 91% @ 100% load
Holdup time	5 ms from ACOKn high to rails out of regulation (see SBB v2 specification)
Main inlet connector	IEC60320 C20 with cable retention
Weight	3 kg (6.6 lb)
Cooling fans	Two stacked fans: 80 mm x 80 mm x 38 mm (3.1 in. x 3.15 in. x 1.45 in.)