Dell PowerVault MD3600f and MD3620f Storage Arrays

Owner's Manual

Notes, Cautions, and Warnings

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

CAUTION: A CAUTION indicates potential damage to hardware or loss of data if instructions are not followed.

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

———————————————

Contents

11 Configuration: Premium Feature—Virtual Disk Copy 169

1

Introduction

📝 **NOTE:** Unless specified, MD3600f Series represents Dell PowerVault MD3600f and Dell PowerVault MD3620f storage arrays.

⚠️ **WARNING: See the Safety, Environmental, and Regulatory Information document for important safety information before following any procedures listed in this document.**

About This Document

This guide familiarizes you with the functions of the Dell PowerVault MD3600f Series storage array. The guide is organized according to the tasks that you must complete after receiving your MD3600f Series storage array. The tasks are:

Planning—Information about the storage array and its features.

Configuration—Tasks that must be completed to ensure that your storage array performs optimally.

Management—Tasks that may be performed to ensure that the storage array components are up to date and performing properly, including removal and installation of storage array components.

Troubleshooting—Tasks that you must complete to resolve problems that may occur with the storage array.

Additional information on these and other topics can be found in the *Dell PowerVault MD3600f and MD3620f Storage Array Deployment Guide* at **dell.com/support/manuals**.

Inside the Box of the Dell PowerVault MD3600f Series Storage Array

Your MD3600f Series product package includes:

- Dell PowerVault MD3600f Series storage array
- Small form-factor pluggable (SFP) transceiver (2 per controller)
- Fiber optic cables (optional)
- Power cables (2)
- Front bezel (optional)
- Mounting rails (2) (optional)
- MD3600f Series resource media
- *Rack Installation Instructions*
- *Getting Started With Your System* (provides an overview of enclosure features, setting up your enclosure, and technical specifications)

MD3600f Series Storage Array

The MD3600f Series is a 2U rack-mounted external redundant array of independent disks (RAID) storage array capable of accommodating up to twelve 3.5" or twenty four 2.5" 6.0-Gbps Serial-Attached SCSI (SAS) disks. The MD3600f Series storage arrays can be daisy-chained with MD1200 Series expansion enclosures, providing access to a maximum of 120 disks (or 192 disks with Premium Feature activation) in the entire storage array. Connectivity between the storage array and the host server is provided by a Fibre Channel (FC) host bus adapter (HBA).

Dell PowerVault Modular Disk Storage Manager

Dell PowerVault Modular Disk Storage Manager (MDSM) is a graphical user interface (GUI) application, used to configure and manage one or more MD3600f Series storage arrays. The MDSM software is located on the MD3600f Series resource media.

Other Information You May Need

⚠️ **WARNING: See the safety and regulatory information that shipped with your system. Warranty information may be included within this document or as a separate document.**

- The *Getting Started Guide* provides an overview of setting up and cabling your storage array.

- The *Deployment Guide* provides installation and configuration instructions for both software and hardware.

- The *Storage Manager CLI Guide* provides information about using the command line interface (CLI).

- The Resource media contains all system management tools.

- The *Configuring Fibre Channel with the Dell MD3600f Series Storage Array* provides information regarding configurations, host bus adapter (HBA) installation, and zoning. This document is available at **dell.com/support/manuals.**

- The *Systems Support Matrix* provides information on supported software and hardware for MD systems. The document is available at **dell.com/support/manuals**.

- The *Dell PowerEdge Cluster Documentation* is available at **dell.com/support/manuals**.

- *Release notes* or readme files are included to provide last-minute updates to the enclosure or documentation or advanced technical reference material intended for experienced users or technicians.

- This document as well as *Dell PowerVault MD 1200 Series Installation Guide* is available at **dell.com/support/manuals** for users who incorporate MD1200 expansion enclosures.

- The *Rack Installation Instructions* included with your rack solution describes how to install your enclosure into a rack.

📝 **NOTE:** Always check for updates on **dell.com/support/manuals** and read the updates first because they often supersede information in other documents.

2

Planning: About Your Storage Array

Overview

The MD3600f Series storage array is designed for high availability, offering redundant access to data storage. It's features support for both single and dual RAID controller configurations.

The Dell PowerVault MD3600f Series storage array provides Fibre Channel (FC) connectivity to the host server and enables access for up to eight non-redundant servers or four redundant servers.

The MD3600f Series storage array includes a number of components. These are:

- RAID controller module(s)
- PSU/fan modules
- Disk drives (also called physical disks/hard drives in this document)
- A front bezel (optional)
- A system enclosure, into which, the other components are plugged

Hardware Features

Front-Panel Features and Indicators

Figure 2-1.   Front-Panel Features and Indicators—Dell PowerVault MD3600f

Figure 2-2.   Front-Panel Features and Indicators—Dell PowerVault MD3620f

Figure 2-3. Front-Bezel Features and Indicators

Item	Indicator, Button, or Connector	Icon	Description
1	Enclosure status LED	⎍∿•	The enclosure status LED lights when the enclosure power is on.
			Lights blue during normal operation.
			Blinks blue when a host server is identifying the enclosure or when the system identification button is pressed.
			Lights amber as enclosure boots or is reset.
			Blinks amber when the enclosure is either in a fault state or the hosts are not using the preferred path to a virtual disk.
2	Power LED	①	The power LED lights green when at least one power supply is supplying power to the enclosure.

Item	Indicator, Button, or Connector	Icon	Description
3	Split mode LED		This LED must be unlit as the split mode function is not supported by the MD3600f Series Storage Arrays.
4	System identification button		The system identification button on the front control panel can be used to locate a particular enclosure within a rack. When the button is pushed, the system status indicators on the control panel and the RAID controller module(s) blink blue until the button is pushed again.
5	Hard drives		MD3600f—Up to twelve 3.5" SAS hot-swappable hard drives.
			MD3620f—Up to twenty four 2.5" SAS hot-swappable hard drives.
6	Enclosure mode switch		The function of this switch is not applicable to your storage array. However, if MD1200 Series expansion enclosures are daisy chained to the storage array, the enclosure mode switches of the MD1200 Series expansion enclosures must be set to the Unified-Mode position.
			NOTE: This switch must be set before turning on the MD1200 Series expansion enclosure. Changing the switch setting after the expansion enclosure is turned on has no effect on the enclosure configuration until the expansion enclosure goes through a complete power cycle.

Back Panel Features and Indicators

Figure 2-4.   Back-Panel Features and Indicators—Dell PowerVault MD3600f Series

1	600 W power supply/cooling fan module	2	RAID Controller Module 0
3	RAID Controller Module 1	4	600 W power supply/cooling fan module

Hard–Drive Indicator Patterns

Figure 2-5.   Hard Drive Indicators

1 hard-drive activity indicator (green)	2 hard-drive status indicator (green and amber)

Drive-Status Indicator Pattern	Condition
Off	The physical disk is: • not yet discovered by the host server • is spun down for removal • not supported for the RAID controller module or is not in the physical disk slot **NOTE:** The drive status indicator remains off until all hard drives are initialized after system power is turned on. Drives are not ready for insertion or removal during this time.
Steady green	Physical disk is online
Green flashing (On 250 ms, Off 250 ms)	Physical disk is being identified
Green flashing (On 400 ms, Off 100 ms)	Physical disk rebuilding
Amber flashing (On 150 ms, Off 150 ms)	Physical disk failed
Flashing green, amber, and Off (green On 500 ms, amber On 500 ms, Off 1000 ms)	Physical disk failure predicted (SMART)
Flashing green, amber, and Off (green 3 seconds, amber 3 seconds, and Off 3 seconds)	Physical disk rebuild aborted

Power Supply and Cooling Fan Features

The MD3600f Series storage array includes two integrated, hot-swappable power supply/cooling fan modules. Both modules must be installed to ensure proper cooling. The system requires at least one of the cooling fans to function to avoid overheating.

A power supply/cooling fan module can be replaced without powering down the system. For information on removing and installing the modules, see "Power Supply/Cooling Fan Module" on page 257.

⚠ **CAUTION: A power supply/cooling fan module can be removed from a powered-on enclosure for a maximum period of 5 minutes. Beyond that time, the system may automatically shut down to prevent damage.**

Power Indicator Codes and Features

Figure 2-6.   Power Indicator Codes and Features

Item	LED Type	Icon	Description
1	DC power	⏻	The LED lights green when the DC output voltage is within the limit.
			If this LED is off, it indicates that the DC output voltage are not within the limit.
2	Power supply/cooling fan fault	⚠	The LED lights amber when the DC output voltage is not within the limit or a fault with the fan is detected.
			If this LED is off, it indicates that no fault condition is present.
3	AC power		The LED lights green when the AC input voltage is within the limit.
			If this LED is off, it indicates either there is no power or the AC input voltage is not within the limit.
4	Power connector		Connect the external power supply to this connector.
5	Power switches (2)		The power switch controls the power supply output to the enclosure.

3

Planning: RAID Controller Modules

RAID Controller Modules

The RAID controller modules provide high-performance, advanced virtual disk configuration, and fault-tolerant disk subsystem management. Each RAID controller module contains 2 GB or 4 GB of mirrored cache for high availability and is protected by a battery powered cache offload mechanism.

NOTE: The 4 GB mirrored cache is an optional feature.

RAID controller modules provide the following data path and enclosure management functions:

- Monitoring and controlling enclosure environment elements (temperature, fans, power supplies, and enclosure LEDs)
- Controlling access to the physical disks
- Communicating enclosure attributes and states to the host server and management station

Each RAID controller module has multiple Fibre Channel (FC) ports (labeled FCIN) for host access. The ports provide redundant host connections and support a high availability storage environment. Each port must have a small form-factor pluggable (SFP) transceiver module installed for connectivity. Various configurations can be utilized, in both single controller (simplex) and dual controller (duplex) modes, to connect the storage enclosure to hosts depending on specific redundancy needs.

For information on cabling, see the *MD3600f and MD3620f Series Storage Array's Deployment Guide,* at **dell.com/support/manuals**.

RAID Controller Module Connectors and Features

Figure 3-1. MD3600f Series Fibre Channel RAID Controller Module

Item	Component	Function
1	SAS OUT port	Provides SAS connection for cabling to a downchain expansion enclosure.
2	FCIN port 0	Provides host-to-controller FC connection.
3	FCIN port 1	Provides host-to-controller FC connection.
4	FCIN port 2	Provides host-to-controller FC connection.
5	FCIN port 3	Provides host-to-controller FC connection.
6	MAC Address label	Provides MAC addresses of the management port.
7	Debug port	Dell support only.
8	Management port activity LED	Lights green when Ethernet connection is active. Off when Ethernet connection is not active.

Item	Component	Function
9	Management port Ethernet connector	Provides a 100/1000 Mbps Ethernet connection for out-of-band management of the enclosure.
10	Management port speed LED	Lights green when Ethernet connection is operating at 1000 Mbps.
		Lights amber when Ethernet connection is operating at 100 Mbps.
		Off when Ethernet connection is operating at 10 Mbps or is not active.
11	Password reset switch	Activating this switch resets the password.
12	Battery fault	Lights amber when battery backup unit or battery has failed.
		Off when battery backup unit is operating normally.
13	FCIN port 3 LED 1 / LED 0	Displays link rate and port status per the accompanying host channel LED link/rate table.
14	Cache active or cache offload LED	Lights green when on-board controller memory contains data.
		If AC power fails, this LED changes to indicate Cache Offload status.
		If the password reset function has successfully changed the password, this LED flashes on and off briefly.
15	FCIN port 2 LED 1 / LED 0	Displays link rate and port status per the accompanying host channel LED link/rate table.
16	System identification LED	Blinks blue when system identification switch push-button on enclosure front panel is pressed.
17	Controller fault LED	Lights amber when controller fault detected.
		Off when controller operating normally.
18	FCIN port 1 LED 1 / LED 0	Displays link rate and port status per the accompanying host channel LED link/rate table.
19	Controller power LED	Lights green when controller power is on.
		Off when controller is not powered.

Item	Component	Function
20	FCIN port 0 LED 1 / LED 0	Displays link rate and port status per the accompanying host channel LED link/rate table.
21	SAS OUT port link/fault LED	Lights green when all four links are connected. Lights amber when one to 3 links are disconnected. Off when all links in the port are disconnected or cable is disconnected.

Host Channel LED Link/Rate Indications

Below each FC port is a pair of LED indicators. The status of any FC port can be determined by applying Table 3-1 to the condition of each LED pair.

Table 3-1.   LED Status Indicators

LED 1	LED 0	CommRate
Off	Off	Link down
Off	On	2 Gbps
On	Off	4 Gbps
On	On	8 Gbps

RAID Controller Module—Additional Features

Battery Backup Unit

Each RAID controller contains a two-cell Lithium ion nanopolymer battery backup unit (BBU). It provides power to the RAID controller module in the event of a power outage, see "RAID Controller Module Backup Battery Unit" on page 255.

📝 **NOTE:** For virtual disks, the RAID controller firmware changes the data cache setting based on the state of the battery. If the battery is missing or does not have sufficient charge, the controller flushes the cache and sets the write cache attribute to **Write Through** for all virtual disks. When the battery is replaced, **Write Back** is re-enabled.

Storage Array Thermal Shutdown

The system automatically shuts down when system temperature exceeds the safe threshold. The battery backup unit protects against data loss by providing power to offload cache to non-volatile memory in the event of power loss. It is not necessary to shut down any MD1200 Series expansion enclosures attached to the storage array when thermal shutdown occurs.

Temperature threshold values determine the temperature at which shutdown occurs. These thresholds cannot be changed.

Table 3-2.   Shutdown Threshold Type

Threshold Temperature Exceeding	Event Description
Nominal failure threshold	A critical event is set.
Maximum failure threshold	Shutdown of the system power supplies occurs within 3 minutes.
Shutdown threshold	Shutdown of the system power supplies occurs within 5 seconds.

System Password Reset

The storage array password can be reset if it is forgotten. To reset the password, push and hold down the password reset switch for at least 5 seconds. The password is deleted. See Figure 3-1 to locate the password reset switch.

The RAID controller module allows you to change the password. For more about setting your password, see "Setting a Password" on page 75.

NOTE: The reset switch can be accessed by using a small object such as the tip of a pen.

Cache Functions and Features

Cache Mirroring

Cache mirroring function copies accepted host-write data from the primary controller to the partner controller. This action ensures that host-write data is safely mirrored to the partner controller before successful completion status is returned to the host. If a controller fails, the surviving controller safely retains all mirrored data. Cache mirroring is enabled by default.

Write–Back Cache

In write-back cache, write operations result in a completion signal being sent to the host operating system as soon as the cache receives the data to be written. The target physical disk receives the data at a more appropriate time in order to increase controller performance. In dual-active controller configurations with write-back caching enabled, the write data is always mirrored to the cache of the second controller before completion status is issued to the host initiator. Write-back cache is enabled by default unless cache mirroring is disabled.

Write–Through Cache

In write-through cache, data is written to the physical disk before completion status is returned to the host operating system. Write-through cache is considered more robust than write-back cache, since a power failure is less likely to cause loss of data. The RAID controller automatically switches to write-through if cache mirroring is disabled or if the battery is missing or has a fault condition.

SFP Transceivers, Fiber Optic, and SAS Cables

Figure 3-2 shows fiber optic cable and SFP transceivers.

NOTE: Your SFP transceivers and cables may look different from the ones shown below. The differences do not affect the performance of the SFP transceivers.

FC host connections may operate at 8 Gbps or at a lower data rate. Ports for 8 Gbps FC host connections require SFP transceivers designed for this data rate. SFP transceivers that support other data rates are incompatible.

WARNING: Do not disassemble or remove any part of a small form-factor pluggable (SFP) transceiver due to the possibility of being exposed to laser radiation.

Figure 3-2. Fiber Optic Cable Connection

| 1 | SFP transceiver | 2 | Fiber optic cable |

The RAID expansion ports support SAS expansion connections. Figure 3-3 shows a SAS cable and an SFF-8088 connector.

Figure 3-3. SAS Cable

| 1 | SFF-8088 connector | 2 | SAS cable |

Interoperability of 2 Gbps, 4 Gbps, and 8 Gbps Devices

The FC standard specifies a procedure for speedy auto-detection. If a 4 Gbps port on a switch or device is connected to a 2 Gbps port, it must negotiate down for the link to run at 2 Gbps. If there are two 8 Gbps ports on either end of a link, the negotiation runs the link at 8 Gbps if the link supports the required specifications. If a link connects at 4 Gbps in this scenario, possible causes could be excessive cable length, poor cable quality, loose cables or SFPs, and so on. Ensure that the cable lengths are not excessive and the quality of the cable and connector is good.

The same rules apply to 8 Gbps devices relative to 4 Gbps and 2 Gbps environments. The 8 Gbps devices have the ability to automatically negotiate down to either 4 Gbps or 2 Gbps, depending on the attached device and link quality. If the link negotiates down to a speed slower than expected, check the device.

4

Planning: MD3600f Series Storage Array Terms and Concepts

This chapter explains terms and concepts used for configuration and operation of MD3600f Series storage arrays.

Physical Disks, Virtual Disks, and Disk Groups

Physical disks in your storage array provide the physical storage capacity for your data. Before you can begin writing data to the storage array, you must configure the physical storage capacity into logical components, called disk groups and virtual disks.

A disk group is a set of physical disks upon which multiple virtual disks are created. The maximum number of physical disks supported in a disk group is 120 disks (or 192 disks with Premium Feature activation) for RAID 0, RAID 1, and RAID 10, and 30 drives for RAID 5 and RAID 6. You can create disk groups from unconfigured capacity on your storage array.

A virtual disk is a partition in a disk group that is made up of contiguous data segments of the physical disks in the disk group. A virtual disk consists of data segments from all physical disks in the disk group. Virtual disks and disk groups are set up according to how you plan to organize your data. For example, you may have one virtual disk for inventory, a second virtual disk for financial and tax information, and so on.

All virtual disks in a disk group support the same RAID level. The storage array supports up to 255 virtual disks (minimum size of 10 MB each) that can be assigned to host servers. Each virtual disk is assigned a Logical Unit Number (LUN) that is recognized by the host operating system.

Physical Disks

Only Dell supported 6.0 Gbps SAS physical disks are supported in the storage array. If the storage array detects unsupported physical disks, it marks the disk as unsupported and the physical disk becomes unavailable for all operations.

> **NOTE:** The MD3600f Series storage array must contain at least two physical disks for proper operation. This is necessary because the physical disks are used to store configuration information.

Physical Disk States

Table 4-1 describes the various states of the physical disk, which are recognized by the storage array and reported in the MDSM application.

Table 4-1. RAID Controller Physical Disk States

Status	Mode	Description	Physical Disk Status LED
Optimal	Assigned	The physical disk in the indicated slot is configured as part of a disk group.	Steady Green
Optimal	Unassigned	The physical disk in the indicated slot is unused and available to be configured.	Steady Green
Optimal	Hot Spare standby	The physical disk in the indicated slot is configured as a hot spare.	Steady Green
Optimal	Hot Spare in use	The physical disk in the indicated slot is in use as a hot spare within a disk group.	Steady Green
Failed	Assigned, Unassigned, Hot Spare in use, or Hot Spare standby	The physical disk in the indicated slot has failed because of an unrecoverable error, an incorrect drive type or drive size, or by its operational state being set to failed.	Amber flashing (150 ms)
Replaced	Assigned	The physical disk in the indicated slot is replaced and is ready to be, or is actively being configured into a disk group.	Green flashing (On 400 ms, Off 100 ms)

Table 4-1. RAID Controller Physical Disk States *(continued)*

Status	Mode	Description	Physical Disk Status LED
Pending Failure	Assigned, Unassigned, Hot Spare in use, or Hot Spare standby	A Self-Monitoring Analysis and Reporting Technology (SMART) error is detected on the physical disk in the indicated slot.	Green flashing (500 ms), Amber (500 ms), and Off (1000ms)
Offline	Not applicable	The physical disk has either been spun down or had a rebuild aborted by user request.	Green flashing (3000 ms), Amber (3000 ms), and Off (3000ms)
Identify	Assigned, Unassigned, Hot Spare in use, or Hot Spare standby	The physical disk is being identified.	Green flashing (250 ms)

If a disk drive rebuild fails because of a source drive failure or because the drive is too small, the MDSM reports a failure of the physical disk even though the LED state on the drive indicates the rebuild was aborted (green for 3 seconds, amber for 3 seconds, then off for 3 seconds).

Self-Monitoring Analysis and Reporting Technology

SMART monitors the internal performance of all physical disk components to detect faults indicating the potential for physical disk failure. SMART uses this information to report whether failure is imminent so that a physical disk can be replaced before failure occurs. The storage array monitors all attached drives and notifies you when a predicted failure is reported by a physical disk.

Virtual Disks and Disk Groups

1 Organize the physical disks into disk groups.

2 Create virtual disks within these disk groups.

3 Provide host server access.

4 Create mappings to associate the virtual disks with the host servers.

NOTE: Host server access must be created before mapping virtual disks.

Disk groups are always created in the unconfigured capacity of a storage array. Unconfigured capacity is the available physical disk space not already assigned in the storage array.

Virtual disks are created within the free capacity of a disk group. Free capacity is the space in a disk group that has not been assigned to a virtual disk.

Virtual Disk States

The storage array recognizes the following virtual disk states.

Table 4-2. RAID Controller Virtual Disk States

State	Description
Optimal	The virtual disk contains physical disks that are all online.
Degraded	The virtual disk with a redundant RAID level contains an inaccessible physical disk. The system can still work properly, but performance may be affected and additional disk failures may result in data loss.
Offline	A virtual disk with one or more member disks is in an inaccessible (failed, missing, or offline) state. Data on the virtual disk is no longer accessible.
Force online	The storage array forces a virtual disk that is in an **Offline** state to an **Optimal** state. If all the member physical disks are not available, the storage array forces the virtual disk to a **Degraded** state. The storage array can force a virtual disk to an **Online** state only when a sufficient number of physical disks are available to support the virtual disk.

RAID Levels

RAID levels determine the way in which data is written to physical disks. Different RAID levels provide different levels of accessibility, redundancy, and capacity.

Using multiple physical disks has the following advantages over using a single physical disk:

- Placing data on multiple physical disks (striping) allows input/output (I/O) operations to occur simultaneously and improve performance.

- Storing redundant data on multiple physical disks using mirroring or parity supports reconstruction of lost data if an error occurs, even if that error is the failure of a physical disk.

Each RAID level provides different performance and protection. You must select a RAID level based on the type of application, access, fault tolerance, and data you are storing.

The storage array supports RAID levels 0, 1, 5, 6, and 10. The maximum number of physical disks that can be used in a disk group depends on the RAID level:

- 192 for RAID levels 0, 1, and 10
- 30 for RAID levels 5 and 6.

RAID Level Usage

To ensure best performance, you must select an optimal RAID level when you create a system physical disk. The optimal RAID level for your disk array depends on:

- Number of physical disks in the disk array
- Capacity of the physical disks in the disk array
- Need for redundant access to the data (fault tolerance)
- Disk performance requirements

RAID 0

RAID 0 uses disk striping to provide high data throughput, especially for large files in an environment that requires no data redundancy. RAID 0 breaks the data down into segments and writes each segment to a separate physical disk. I/O performance is greatly improved by spreading the I/O load across many physical disks. Although it offers the best performance of any RAID level, RAID 0 lacks data redundancy. Select this option only for non-critical data, because failure of one physical disk results in the loss of all data. Examples of RAID 0 applications include video editing, image editing, prepress applications, or any application requiring high bandwidth.

RAID 1

RAID 1 uses disk mirroring so that data written to one physical disk is simultaneously written to another physical disk. This RAID level offers fast performance, the best data availability, and the highest disk overhead. RAID 1 is recommended for small databases or other applications that do not require large capacity. RAID 1 provides full data redundancy. For example accounting, payroll, or financial applications.

RAID 5

RAID 5 uses parity and striping data across all physical disks (distributed parity) to provide high data throughput and data redundancy, especially for small random access. This is a versatile RAID level and is suited for multi-user environments where typical I/O size is small and there is a high proportion of read activity such as file, application, database, web, e-mail, news, and intranet servers.

RAID 6

RAID 6 is similar to RAID 5 but provides an additional parity disk for better redundancy. This is the most versatile RAID level and is suited for multi-user environments where typical I/O size is small and there is a high proportion of read activity. RAID 6 is recommended when large size physical disks are used or large number of physical disks are used in a disk group.

RAID 10

RAID 10, a combination of RAID 1 and RAID 0, uses disk striping across mirrored disks. It provides high data throughput and complete data redundancy. Utilizing an even number of physical disks (four or more) creates a RAID level 10 disk group and/or virtual disk. Because RAID levels 1 and 10 use disk mirroring, half of the capacity of the physical disks is utilized for mirroring. This leaves the remaining half of the physical disk capacity for actual storage. RAID 10 is automatically used when a RAID level of 1 is chosen with four or more physical disks. RAID 10 works well for medium-sized databases or any environment that requires high performance and fault tolerance and moderate-to-medium capacity.

Segment Size

Disk striping enables data to be written across multiple physical disks. Disk striping enhances performance because striped disks are accessed simultaneously.

The segment size or stripe element size specifies the size of data in a stripe written to a single disk. The MD3600f Series array supports stripe element sizes of 8 KB, 16 KB, 32 KB, 64 KB, 128 KB, 256 KB, and 512 KB. The default stripe element size is 128 KB.

Stripe width, or depth, refers to the number of disks involved in an array where striping is implemented. For example, a four-disk group with disk striping has a stripe width of four.

NOTE: Although disk striping delivers excellent performance, striping alone does not provide data redundancy.

Virtual Disk Operations

Virtual Disk Initialization

Every virtual disk must be initialized. Initialization can be done in the foreground or the background. A maximum of four virtual disks can be initialized concurrently on each RAID controller module.

Background Initialization

The storage array executes a background initialization when the virtual disk is created to establish parity, while allowing full host server access to the virtual disks. Background initialization does not run on RAID 0 virtual disks. The background initialization rate is controlled by MDSM. To change the rate of background initialization, you must stop any existing background initialization. The rate change is implemented when the background initialization restarts automatically.

Foreground Initialization

The storage array supports foreground initialization for virtual disks. All access to the virtual disk is blocked during foreground initialization. During foreground initialization, zeros (0x00) are written to every sector of the virtual disk. The virtual disk is available after foreground initialization is completed.

Consistency Check

A consistency check verifies the correctness of data in a redundant array (RAID levels 1, 5, 6, and 10). For example, in a system with parity, checking consistency involves computing the data on one physical disk and comparing the results to the contents of the parity physical disk.

A consistency check is similar to a background initialization. The difference is that background initialization cannot be started or stopped manually, while consistency check can.

 NOTE: It is recommended that you run data consistency checks on a redundant array at least once a month. This allows detection and automatic replacement of unreadable sectors. Finding an unreadable sector during a rebuild of a failed physical disk is a serious problem, because the system does not have the redundancy to recover the data.

Media Verification

Another background task performed by the storage array is media verification of all configured physical disks in a disk group. The storage array uses the Read operation to perform verification on the space configured in virtual disks and the space reserved for the metadata.

Cycle Time

The media verification operation runs only on selected disk groups, independent of other disk groups. Cycle time is the time taken to complete verification of the metadata region of the disk group and all virtual disks in the disk group for which media verification is configured. The next cycle for a disk group starts automatically when the current cycle completes. You can set the cycle time for a media verification operation between 1 and 30 days. The storage controller throttles the media verification I/O accesses to disks based on the cycle time.

The storage array tracks the cycle for each disk group independent of other disk groups on the controller and creates a checkpoint. If the media verification operation on a disk group is preempted or blocked by another operation on the disk group, the storage array resumes after the current cycle. If the media verification process on a disk group is stopped due to a RAID controller module restart, the storage array resumes the process from the last checkpoint.

Virtual Disk Operations Limit

The maximum number of active, concurrent virtual disk processes per RAID controller module installed in the storage array is four. This limit is applied to the following virtual disk processes:

- Background initialization
- Foreground initialization
- Consistency check
- Rebuild
- Copy back

If a redundant RAID controller module fails with existing virtual disk processes, the processes on the failed controller are transferred to the peer controller. A transferred process is placed in a suspended state if there are four active processes on the peer controller. The suspended processes are resumed on the peer controller when the number of active processes falls below four.

Disk Group Operations

RAID Level Migration

You can migrate from one RAID level to another depending on your requirements. For example, fault-tolerant characteristics can be added to a stripe set (RAID 0) by converting it to a RAID 5 set. MDSM provides information about RAID attributes to assist you in selecting the appropriate RAID level. You can perform a RAID level migration while the system is still running and without rebooting, which maintains data availability.

Segment Size Migration

Segment size refers to the amount of data (in KB) that the storage array writes on a single physical disk in a virtual disk before writing data on the next physical disk. Valid values for the segment size are 8 KB, 16 KB, 32 KB, 64 KB, 128 KB, 256 KB, and 512 KB.

Dynamic segment size migration enables the segment size of a given virtual disk to be changed. A default segment size is set when the virtual disk is created, based on such factors as the RAID level and expected usage. You can change the default value (128 KB) if segment size usage does not match your needs.

When considering a segment-size change, two scenarios illustrate different approaches to the limitations:

- If I/O activity stretches beyond the segment size, you can increase it to reduce the number of disks required for a single I/O. Using a single physical disk for a single request frees disks to service other requests, especially when you have multiple users accessing a database or storage environment.

- If you use the virtual disk in a single-user, large I/O environment (such as for multimedia application storage), performance can be optimized when a single I/O request is serviced with a single data stripe (the segment size multiplied by the number of physical disks in the disk group used for data storage). In this case, multiple disks are used for the same request, but each disk is only accessed once.

Virtual Disk Capacity Expansion

When you configure a virtual disk, you select a capacity based on the amount of data you expect to store. However, you may need to increase the virtual disk capacity for a standard virtual disk by adding free capacity to the disk group. This creates more unused space for new virtual disks or to expand existing virtual disks.

Disk Group Expansion

Because the storage array supports hot pluggable physical disks, you can add two physical disks at a time for each disk group while the storage array remains online. Data remains accessible on virtual disk groups, virtual disks, and physical disks throughout the operation. The data and increased unused free space are dynamically redistributed across the disk group. RAID characteristics are also reapplied to the disk group as a whole.

Disk Group Defragmentation

Defragmenting consolidates the free capacity in the disk group into one contiguous area. Defragmentation does not change the way in which the data is stored on the virtual disks.

Disk Group Operations Limit

The maximum number of active, concurrent disk group processes per installed RAID controller module is one. This limit is applied to the following disk group processes:

- Virtual disk RAID level migration
- Segment size migration
- Virtual disk capacity expansion
- Disk group expansion
- Disk group defragmentation

If a redundant RAID controller module fails with an existing disk group process, the process on the failed controller is transferred to the peer controller. A transferred process is placed in a suspended state if there is an active disk group process on the peer controller. The suspended processes are resumed when the active process on the peer controller completes or is stopped.

NOTE: If you try to start a disk group process on a controller that does not have an existing active process, the start attempt fails if the first virtual disk in the disk group is owned by the other controller and there is an active process on the other controller.

RAID Background Operations Priority

The storage array supports a common configurable priority for the following RAID operations:

- Background initialization
- Rebuild
- Copy back
- Virtual disk capacity expansion
- Raid level migration
- Segment size migration
- Disk group expansion
- Disk group defragmentation

The priority of each of these operations can be changed to address performance requirements of the environment in which the operations are to be executed.

NOTE: Setting a high priority level impacts storage array performance. It is not advisable to set priority levels at the maximum level. Priority must also be assessed in terms of impact to host server access and time to complete an operation. For example, the longer a rebuild of a degraded virtual disk takes, the greater the risk for secondary disk failure.

Virtual Disk Migration and Disk Roaming

Virtual disk migration is moving a virtual disk or a hot spare from one array to another by detaching the physical disks and re-attaching them to the new array. Disk roaming is moving a physical disk from one slot to another on the same array.

Disk Migration

You can move virtual disks from one array to another without taking the target array offline. However, the disk group being migrated must be offline before your perform disk migration. If the disk group is not offline prior to migration, the source array holding the physical and virtual disks within the disk group marks them as missing. However, the disk groups themselves migrate to the target array.

An array can import a virtual disk only if it is in an optimal state. You can move virtual disks that are part of a disk group only if all members of the disk group are being migrated. The virtual disks automatically become available after the target array has finished importing all the disks in the disk group.

When you migrate a physical disk or a disk group from one MD3600f array to another, the MD3600f array you migrate to, recognizes any data structures and/or metadata you had in place on the migrating MD3600f array. However, if you are migrating from any device other than a MD3600f Series storage array, the MD3600f array does not recognize the migrating metadata and that data is lost. In this case, the MD3600f array initializes the physical disks and marks them as unconfigured capacity.

NOTE: Only disk groups and associated virtual disks with all member physical disks present can be migrated from one storage array to another. it is recommended that you only migrate disk groups that have all their associated member virtual disks in an optimal state.

NOTE: The number of physical disks and virtual disks that a storage array supports limits the scope of the migration.

Use either of the following methods to move disk groups and virtual disks:

- Hot virtual disk migration—Disk migration with the destination storage array power turned on.
- Cold virtual disk migration—Disk migration with the destination storage array power turned off.

✐ **NOTE:** To ensure that the migrating disk groups and virtual disks are correctly recognized when the target storage array has an existing physical disk, use hot virtual disk migration.

When attempting virtual disk migration, follow these recommendations:

- Moving physical disks to the destination array for migration—When inserting drives into the destination storage array during hot virtual disk migration, wait for the inserted physical disk to be displayed in MDSM before inserting the next physical disk.

⚠ **WARNING: Without the delay between drive insertions, the storage array can become unstable and manageability is temporarily lost.**

- Migrating virtual disks from multiple storage arrays into a single storage array—When migrating virtual disks from multiple or different storage arrays into a single destination storage array, move all of the physical disks from the same storage array as a set into the new destination storage array. Ensure that all of the physical disks from a storage array are migrated to the destination storage array before starting migration from the next storage array.

  ✐ **NOTE:** If the drive modules are not moved as a set to the destination storage array, the newly relocated disk groups may not be accessible.

- Migrating virtual disks to a storage array with no existing physical disks— Turn off the destination storage array, when migrating disk groups or a complete set of physical disks from a storage array to another storage array that has no existing physical disks. After the destination storage array is turned on and has successfully recognized the newly migrated physical disks, migration operations can continue.

  ✐ **NOTE:** Disk groups from multiple storage arrays must not be migrated at the same time to a storage array that has no existing physical disks.

- Enabling premium features before migration—Before migrating disk groups and virtual disks, enable the required premium features on the destination storage array. If a disk group is migrated from a MD3600f storage array that has a premium feature enabled and the destination array does not have this feature enabled, an **Out of Compliance** error message may be generated.

Disk Roaming

You can move physical disks within an array. The RAID controller module automatically recognizes the relocated physical disks and logically places them in the proper virtual disks that are part of the disk group. Disk roaming is permitted when the RAID controller module is either online or powered off.

✎ **NOTE:** The disk group must be exported before moving the physical disks.

Advanced Features

The RAID enclosure supports several advanced features:

- Virtual Disk Snapshots
- Virtual Disk Copy
- High Performance Tier
- Remote Replication (this premium feature is supported only in storage arrays with the Fibre Channel [FC] host ports)

✎ **NOTE:** Virtual Disk Snapshot, Virtual Disk Copy, High Performance Tier, and Remote Replication are premium features that must be activated separately. If you have purchased these features, an activation card is supplied that contains instructions for enabling this functionality.

Host Server–to–Virtual Disk Mapping

The host server attached to a storage array accesses various virtual disks on the storage array through its host ports. Specific virtual disk-to-LUN mappings to an individual host server can be defined. In addition, the host server can be part of a host group that shares access to one or more virtual disks.

You can manually configure a host server-to-virtual disk mapping. When you configure host server-to-virtual disk mapping, consider these guidelines:

- You can define one host server-to-virtual disk mapping for each virtual disk in the storage array.
- Host server-to-virtual disk mappings are shared between RAID controller modules in the storage array.
- A unique LUN must be used by a host group or host server to access a virtual disk.
- Not every operating system has the same number of LUNs available for use.

Host Types

A host server is a server that accesses a storage array. Host servers are mapped to the virtual disks. Host servers have the following attributes:

- Host name—A name that uniquely identifies the host server.
- Host group (used in Cluster solutions only)—Two or more host servers associated together to share access to the same virtual disks.

  This host group is a logical entity you can create in MDSM. All host servers in a host group must be running the same operating system.

- Host type—The operating system running on the host server.

Snapshot Virtual Disks

A snapshot is a point-in-time image of a virtual disk. The snapshot provides an image of the virtual disk at the time the snapshot was created. You create a snapshot so that an application (for example, a backup application) can access the snapshot and read the data while the source virtual disk remains online and user-accessible. When the backup is completed, the snapshot virtual disk is no longer needed. You can create up to four snapshots per virtual disk.

Snapshots are used to recover previous versions of files that have changed since the snapshot was taken. Snapshots are implemented using a copy-on-write algorithm, which makes a backup copy of data the instant an error occurs. Data on a virtual disk is copied to the snapshot repository before it is modified. Snapshots can be created instantaneously or can be scheduled and take up less overhead than a full physical copy process.

Snapshot Repository Virtual Disk

When you create a snapshot virtual disk, it automatically creates a snapshot repository virtual disk. A snapshot repository is a virtual disk created in the storage array as a resource for a snapshot virtual disk. A snapshot repository virtual disk contains snapshot virtual disk metadata and copy-on-write data for a particular snapshot virtual disk. The repository supports one snapshot only.

You cannot select a snapshot repository virtual disk as a source virtual disk or as a target virtual disk in a virtual disk copy. If you select a snapshot source virtual disk as the target virtual disk of a virtual disk copy, you must disable all snapshot virtual disks associated with the source virtual disk.

⚠ CAUTION: **Before using the Snapshot Virtual Disks Premium Feature in a Windows Clustered configuration, you must map the snapshot virtual disk to the cluster node that owns the source virtual disk. This ensures that the cluster nodes correctly recognize the snapshot virtual disk.**

⚠ CAUTION: **Mapping the snapshot virtual disk to the node that does not own the source virtual disk before the snapshot enabling process is completed can result in the operating system incorrectly identifying the snapshot virtual disk. This can result in data loss or an inaccessible snapshot.**

For more information on mapping the snapshot virtual disk to the secondary node, see the *Dell PowerVault MD3600f and MD3620f Storage Arrays With Microsoft Windows Server Failover Clusters* on **dell.com/support/manuals**.

Virtual Disk Copy

Virtual disk copy is a premium feature to:

- Back up data
- Copy data from disk groups that use smaller-capacity physical disks to disk groups using greater capacity physical disks
- Restore snapshot virtual disk data to the source virtual disk

Virtual disk copy generates a full copy of data from the source virtual disk to the target virtual disk in a storage array and can be performed either online or offline.

Source Virtual Disk

When you create a virtual disk copy, a copy pair consisting of a source virtual disk and a target virtual disk is created on the same storage array. When a virtual disk copy is started, data from the source virtual disk is copied completely to the target virtual disk.

Target Virtual Disk

When you start a virtual disk copy, the target virtual disk maintains a copy of the data from the source virtual disk. You can choose whether to use an existing virtual disk or create a new virtual disk as the target virtual disk. If you choose an existing virtual disk as the target, all data on the target is overwritten. A target virtual disk can be a standard virtual disk or the source virtual disk of a failed or disabled snapshot virtual disk.

NOTE: The target virtual disk capacity must be equal to or greater than the source virtual disk capacity.

When you begin the disk copy process, you must define the rate at which the copy is completed. Giving the copy process top priority slightly impacts I/O performance, while giving it lowest priority makes the copy process longer to complete. You can modify the copy priority while the disk copy is in progress.

For more information, see the *online help* topics.

Virtual Disk Recovery

You can use the **Edit host server-to-virtual disk mappings** feature to recover data from the backup virtual disk. This functionality enables you to unmap the original source virtual disk from its host server, then map the backup virtual disk to the same host server.

Ensure that you record the LUN used to provide access to the source virtual disk. You need this information when you define a host server-to-virtual disk mapping for the target (backup) virtual disk. Also, be sure to stop all I/O activity to the source virtual disk before beginning the virtual disk recovery procedure.

Using Snapshot and Disk Copy Together

You can use the **Snapshot Virtual Disk** and **Virtual Disk Copy** premium features together to back up data on the same storage array, or to restore the data on the snapshot virtual disk to its original source virtual disk.

You can copy data from a virtual disk by:

- Taking a point-in-time snapshot of the data (online)
- Copying the data to another virtual disk using a virtual disk copy (offline)

You can select a snapshot virtual disk as the source virtual disk for a virtual disk copy. This configuration is one of the best ways you can apply the snapshot virtual disk feature, since it enables complete backups without any impact on the storage array I/O.

You cannot use a snapshot repository virtual disk as a source virtual disk or as a target virtual disk in a virtual disk copy. If you select the source virtual disk as the target virtual disk of a virtual disk copy, you must disable all snapshot virtual disks associated with the source virtual disk.

Multi-Path Software

Multi-path software (also referred to as the failover driver) is a software resident on the host server that provides management of the redundant data path between the host server and the storage array.

The multi-path software identifies the existence of multiple paths to a virtual disk and establishes a preferred path to that disk. If any component in the preferred path fails, the multi-path software automatically re-routes I/O requests to the alternate path so that the storage array continues to operate without interruption.

✍ **NOTE:** Multi-path software is available on the MD3600f Series resource media.

Preferred and Alternate Controllers and Paths

A preferred controller is a RAID controller module designated as the owner of a virtual disk or disk group. The preferred controller is automatically selected by MDSM when a virtual disk is created. You can change the preferred RAID controller module owner of a virtual disk after it is created. If a host is only connected to only one RAID controller module, the preferred owner must manually be assigned to the RAID controller module the host can access.

Ownership of a virtual disk is moved from the preferred controller to the secondary controller (also called the alternate controller) when the preferred controller is:

- Physically removed
- Updating firmware
- Involved in an event that caused failover to the alternate controller

Paths used by the preferred RAID controller module to access either the disks or the host server are called the preferred paths; redundant paths are called the alternate paths. If a failure causes the preferred path to become inaccessible, the storage array automatically uses the alternate path to access data. When this occurs, the enclosure status LED blinks amber.

Virtual Disk Ownership

MDSM can be used to automatically build and view virtual disks. It uses optimal settings to stripe the disk group. Virtual disks are assigned to alternating RAID controller modules when they are created. This default assignation provides a simple means for load balancing the workload of the RAID controller modules.

Ownership can later be modified to balance workload according to actual usage. If virtual disk ownership is not manually balanced, it is possible for one controller to have the majority of the work, while the other controller is idle. Limit the number of virtual disks in a disk group. If multiple virtual disks are in a disk group, consider:

- The impact each virtual disk has on other virtual disks in the same disk group.
- The patterns of usage for each virtual disk.
- Different virtual disks have higher usage at different times of day.

Load Balancing

A load balance policy is used to determine which path is used to process I/O. Multiple options for setting the load balance policies lets you optimize I/O performance when mixed host interfaces are configured.

You can choose one of these load balance policies to optimize I/O performance:

- Round robin with subset—The round robin with subset I/O load balance policy routes I/O requests, in rotation, to each available data path to the RAID controller module that owns the virtual disks. This policy treats all paths to the RAID controller module that owns the virtual disk equally for I/O activity. Paths to the secondary RAID controller module are ignored until ownership changes. The basic assumption for the round-robin policy is that the data paths are equal. With mixed host support, the data paths may have different bandwidths or different data transfer speeds.

- Least queue depth with subset—The least queue depth with subset policy is also known as the least I/Os or least requests policy. This policy routes the next I/O request to a data path that has the least outstanding I/O requests queued. For this policy, an I/O request is simply a command in the queue. The type of command or the number of blocks that are associated with the command are not considered. The least queue depth with subset policy treats large block requests and small block requests equally. The data path selected is one of the paths in the path group of the RAID controller module that owns the virtual disk.

- Least path weight with subset (Windows operating systems only)—The least queue depth with subset policy is also known as the least I/Os or least requests policy. This policy routes the next I/O request to a data path that has the least outstanding I/O requests queued. For this policy, an I/O request is simply a command in the queue. The type of command or the number of blocks that are associated with the command are not considered. The least queue depth with subset policy treats large block requests and small block requests equally. The data path selected is one of the paths in the path group of the RAID controller module that owns the virtual disk.

Monitoring MD3600f Series System Performance

You can use the Performance Monitor to select virtual disks and RAID controller modules to monitor or to change the polling interval. Keep these guidelines in mind when using the Performance Monitor:

- The Performance Monitor does not dynamically update its display if any configuration changes occur while the window is open. You must close the **Performance Monitor** window and reopen it for the changes to be displayed.

- Using the Performance Monitor to retrieve performance data can affect the normal storage array performance depending on the polling interval that you set.

- If the storage array you are monitoring begins in or transitions to an unresponsive state, an informational dialog is displayed. The dialog informs you that the Performance Monitor cannot poll the storage array for performance data.

To monitor the performance of the arrays:

1 Open MDSM and select the appropriate storage array.

2 Open the **Array Management Window** (AMW) for the selected storage array.

3 In the AMW, select **Storage Array→ Monitor Performance**.

4 Click **Settings**.

    a Select the components that you want to monitor. You can monitor:

       - RAID controller modules

       - Virtual disks

       - Storage array totals

    NOTE: By default, all items are selected.

    b In **Polling interval**, select how often you want to update the performance statistics.

    NOTE: For an accurate elapsed time, do not use the **Set RAID Controller Module Clocks** option while using the Performance Monitor.

Each time the polling interval elapses, the Performance Monitor queries the storage array again and updates the statistics in the table.

5 Click **Start**.

Values are displayed for the selected storage arrays in the Performance Monitor data table. The table is updated at the interval specified in the Polling Interval setting.

6 Click **Update** to force an immediate poll of the storage array.

7 Click **Stop** to stop monitoring the storage array.

8 Click **Save As** on the **Performance Monitor** main dialog to save the currently displayed performance statistics.

9 Select an appropriate directory.

10 Type a file name in the **File name** text box.

   ![note icon] **NOTE:** The .perf extension is the default.

11 Select a file type from the **Files of type** list.

   • Use the **Report format** (ASCII text) file type to save the data to a report form for viewing or printing.

   • Use the **Comma Delimited Format** file type to save the data in a form that can be imported into a commercial spreadsheet application for further analysis. Most leading commercial spreadsheet applications recognize a comma delimiter. These applications use the delimiter to import the data into spreadsheet cells.

12 Click **Save**.

The Performance Monitor data provides information about how your storage array is performing. The data is presented in eight columns, which are described in this table. Use this data to make performance tuning decisions for your storage array.

Table 4-3. Performance Monitor Table Description

Column Headings	Description
Devices	Controller, virtual disk or storage array total.
Total IOs	Cumulative I/O's per second from last start time.
Read Percentage	Percentage of cumulative I/O's that are reads.
Cache Hit Percentage	Percentage of cumulative I/O's that are in-cache.
Current KB/second	Snapshot of throughput value per second (1KB = 1024 bytes).
Maximum KB/second	Maximum recorded throughput value from last start time.
Current IO/second	Snapshot of I/O's per second (IOP = Input/output per second or one completed I/O transaction).
Maximum IO/second	Maximum recorded IOP from last start time.

For more information, see the *MD PowerVault Support Matrix* at **dell.com/support/manuals**.

5

Configuration: Overview

Dell PowerVault Modular Disk Storage Manager (MDSM) online help contains information on how to use the MDSM application to perform the configuration and management tasks described in this document. You can access online help by selecting **Help** located at the top right corner of MDSM interface. For information on installing the MDSM, see the *MD3600f and MD3620f Storage Array's Deployment Guide* at **dell.com/support/manuals**.

> **NOTE:** MDSM supports the MD3000 series, MD3200 series, and MD3600 series storage arrays and can automatically detect all of these storage arrays.

User Interface

The Storage Manager screen is divided into two primary windows:

- Enterprise Management Window (EMW)—The EMW provides high-level management of the storage arrays. You can launch the Array Management Window from the EMW.

- Array Management Window (AMW)—The AMW provides management functions for a single storage array. You can launch more than one AMW at the same time to manage different storage arrays.

The EMW and the AMW consist of the following:

- The title bar at the top of the window shows the name of the application.

- Beneath the title bar, is the menu bar. You can select menu options from the menu bar to perform tasks on a storage array.

- Beneath the menu bar, is the toolbar. You can select options in the toolbar to perform tasks on a storage array.

- Beneath the toolbar, are the tabs. Tabs are used to group the tasks that you can perform on a storage array.

- Beneath the tabs, is the status bar. The status bar shows status messages and status icons related to the storage array.

> **NOTE:** By default, the toolbar and status bar is not displayed. To view the toolbar or the status bar, select **View**→ Toolbar or **View**→ Status Bar, respectively.

Enterprise Management Window

The EMW provides high-level management of storage arrays. When you start MDSM, the EMW is displayed. The EMW has these tabs:

- **Devices** tab—Provides information about the storage arrays.
- **Setup** tab—Presents the initial setup tasks that guide you through adding storage arrays and configuring alerts.

The **Devices** tab has a Tree view on the left side of the window that shows discovered storage arrays, unidentified storage arrays, and the status conditions for the storage arrays. Discovered storage arrays are managed by MDSM. Unidentified storage arrays are available to MDSM but not configured for management. The right side of the **Devices** tab has a Table view that shows detailed information for each storage array.

In the EMW, you can:

- Discover hosts and managed storage arrays on the local sub-network.
- Manually add and remove hosts and storage arrays.
- Locate the storage arrays.
- Name or rename discovered storage arrays.
- Add storage array comments to the Table view.
- Sort rows in the Table view according to different criteria.
- Store your EMW view preferences and configuration data in local configuration files. The next time you open the EMW, data from the local configuration files is used to show customized view and preferences.
- Monitor the status of managed storage arrays and indicate status using appropriate icons.
- Add or remove management connections.
- Configure alert notifications for all selected storage arrays through e-mail or SNMP traps.
- Report critical events to the configured alert destinations.
- Launch the AMW for a selected storage array.
- Run a script to perform batch management tasks on specific storage arrays.
- Import the operating system theme settings into the MDSM.
- Upgrade firmware on multiple storage arrays concurrently.

Inheriting the System Settings

Use the **Inherit System Settings** option to import the operating system theme settings into the MDSM. Importing system theme settings affects features like font type, font size, color, and contrast in the MDSM.

1   Open the **Inherit System Settings** window in one of these ways:

  •   Select **Tools→ Inherit System Settings**.

  •   Select the **Setup** tab and click **Inherit System Settings**.

2   Select **Inherit system settings for color and font**.

3   Click **OK**.

Array Management Window

You can launch the AMW from the EMW. The AMW provides management functions for a single storage array. You can have multiple AMWs open simultaneously to manage different storage arrays.

To launch the AMW:

1   In the **EMW**, on the **Devices** tab, double click on the relevant storage array.

  The context menu for the selected storage is displayed.

2   In the context menu, select **Manage Storage Array.**

  The AMW for the selected storage is displayed.

The AMW has the following tabs:

  •   **Summary** tab—You can view the following information about the storage array:

    –   Status

    –   Hardware components

    –   Capacity

    –   Hosts and Mappings

    –   Storage partitions

    –   Disk groups and virtual disks

  •   **Logical** tab—You can view the organization of the storage array by virtual disks, disk groups, free capacity nodes, and any unconfigured capacity for the storage array.

- **Physical** tab—You can view the organization of the storage array by RAID controller modules, physical disks, and other hardware components.
- **Mappings** tab—You can define the hosts, host groups, and host ports. You can change the mappings to grant virtual disk access to host groups and hosts and create storage partitions.
- **Setup** tab—You can complete the initial setup tasks to configure the storage array.
- **Support**—You can complete common support tasks like downloading RAID controller module firmware, viewing the online help, and so on.

In the AMW, you can:

- Provide storage array options, for example, renaming a storage array, changing a password, or enabling a background media scan.
- Provide the ability to configure virtual disks from the storage array capacity, define hosts and host groups, and grant host or host group access to sets of virtual disks called storage partitions.
- Monitor the health of storage array components and report detailed status using applicable icons.
- Provide applicable recovery procedures for a failed logical component or a failed hardware component.
- Present a view of the Event Log for the storage array.
- Present profile information about hardware components, such as RAID controller modules and physical disks.
- Provide RAID controller module management options, such as changing ownership of virtual disks or placing a RAID controller module online or offline.
- Provide physical disk management options, such as assignment of hot spares and locating the physical disk.
- Monitor storage array performance.

6

Configuration: About Your Storage Array

You can manage a storage array in two ways:

- Out-of-band management
- In-band management

Out-of-Band Management

Using the out-of-band management method, data is separated from commands and events. Data travels through the host-to-controller interface, while commands and events travel through the management port Ethernet cables.

This management method lets you configure the maximum number of virtual disks that are supported by your operating system and host adapters. A maximum of eight storage management stations can concurrently monitor an out-of-band managed storage array. This limit does not apply to systems that manage the storage array through the in-band management method.

When you use out-of-band management, you must set the network configuration for each RAID controller module's management Ethernet port. This includes the Internet Protocol (IP) address, subnetwork mask (subnet mask), and gateway. If you are using a Dynamic Host Configuration Protocol (DHCP) server, you can enable automatic network configuration, but if you are not using a DHCP server, you must enter the network configuration manually.

NOTE: RAID controller module network configurations can be assigned using a DHCP server (the default setting). However, if a DHCP server is not available for 150 seconds, the RAID controller modules assign static IP addresses. The addresses assigned are 192.168.128.101 for controller 0 and 192.168.128.102 for controller 1.

In-Band Management

Using in-band management, commands, events, and data travel through the host-to-controller interface. Unlike out-of-band management, commands and events are mixed with data.

NOTE: For detailed information on setting up in-band and out-of-band management see the Deployment Guide.

When you add storage arrays by using this management method, you need to specify only the host name or IP address of the host. After you add the specific host name or IP address, the host-agent software automatically detects any storage arrays that are connected to that host.

⚠ CAUTION: Some operating systems can be used only as storage management stations. For more information about the operating system that you are using, see the *MD PowerVault Support Matrix* at dell.com/support/manuals.

For more information, see the *PowerVault Modular Disk Storage Manager online help* topics.

Access Virtual Disk

Each RAID controller module in an MD3600f Series storage array maintains a special virtual disk, called the access virtual disk. The host-agent software uses the access virtual disk to communicate management requests and event information between the storage management station and the RAID controller module in an in-band–managed storage array. The access virtual disk is not available for application data storage. The default LUN is 31.

Storage Arrays

You must add the storage arrays to MDSM before you can setup the storage array for optimal use.

Adding Storage Arrays

You can add storage arrays only in the EMW. You can:

- Automatically discover storage arrays
- Manually add storage arrays

✎ NOTE: Verify that your host or management station network configuration—including station IP address, subnet mask, and default gateway—is correct before adding a new storage array using the Automatic option.

✎ NOTE: For Linux, set the default gateway so that broadcast packets are sent to 255.255.255.0. For Red Hat Enterprise Linux, if no gateway exists on the network, set the default gateway to the IP address of the NIC.

✎ NOTE: MDSM uses TCP/UDP port 2463 for communication to the MD Storage Array.

Automatic Discovery of Storage Arrays

The Automatic Discovery process sends out a broadcast message across the local subnetwork (subnet) and adds any storage array that responds to the message. The Automatic Discovery process finds both in-band and out-of-band storage arrays.

NOTE: The Automatic Discovery option and the Re-scan Hosts option in the Enterprise Management Window provide automatic methods to discover managed storage arrays.

Manual Addition of a Storage Array

Use Manual Addition if the storage array resides outside of the local subnet. This process requires specific identification information to manually add a storage array.

To add a storage array that uses out-of-band management, specify the host name or management port IP address of each controller in the storage array.

Before using this option, verify that the applicable network configuration tasks are performed.

To add an in-band storage array, add the host through which the storage array is attached to the network.

NOTE: It can take several minutes for MDSM to connect to the specified storage array.

To add a storage array manually:

1 Select **Edit→ Add Storage Array**.

2 Select the relevant management method:

- **Out-of-band management**—Enter a host name or an IP address for the **RAID controller Modules** in the storage array.

- **In-band management**—Enter a name or an IP address for the **Host** through which the storage array is attached to the network.

  NOTE: The host agent must be restarted before in-band management communication can be established. See "Starting or Restarting the Host Context Agent Software" on page 293.

3 Click **Add**.

4 Use one of these methods to name a storage array:

- In the EMW, select the **Setup** tab, and select **Name/Rename Storage Arrays**.
- In the AMW, select the **Setup** tab, and select **Rename Storage Array**.
- In the EMW, right-click the icon corresponding to the array and select **Rename**.

Setting Up Your Storage Array

A list of initial setup tasks is displayed on the **Setup** tab in the AMW. The list of initial setup tasks shows you how to set up a storage array. Using the steps outlined in the Initial Setup Tasks area, ensures that the basic setup steps are completed properly.

Use the Initial Setup Tasks list the first time that you set up a storage array to perform these tasks:

- Locate the storage array—Find the physical location of the storage array on your network by turning on the unit identify LEDs. The storage array can be identified with a label.
- Give a new name to the storage array—Use a unique name that identifies each storage array.
- Set a storage array password—Configure the storage array with a password to protect it from unauthorized access. MDSM prompts for the password when you attempt to change the storage array configuration. For example, when a virtual disk is created or deleted.
- Configure the storage array—Create disk groups, virtual disks, and hot spare physical disks by using the Automatic configuration method or the Manual configuration method. For more information, see the *PowerVault Modular Disk Storage Manager online help* topics.
- Map Virtual Disks—Map virtual disks to hosts or host groups.
- Save Configuration—Save the configuration parameters in a file that you can use to restore the configuration, or reuse the configuration on another storage array. For more information, see the *PowerVault Modular Disk Storage Manager online help* topics.

After you complete the basic steps for configuring the storage array, you can perform these optional tasks:

- Manually define hosts—Define the hosts and the host port identifiers that are connected to the storage array. Use this option only if the host is not automatically recognized and shown in the **Mappings** tab.

- Configure ethernet management ports—Configure the network parameters for the Ethernet management ports on the RAID controller modules if you are managing the storage array by using the out-of-band management connections.

- View and enable premium features—Your MDSM may include premium features. View the premium features that are available and the premium features that are already started. You can start available premium features that are currently stopped.

Locating Storage Arrays

You can use the Blink option to physically locate and identify a storage array.

To locate the storage array:

1 Select the relevant storage array and:

- In the EMW, right-click the appropriate storage array, and select **Blink Storage Array**.

- In the AMW, select the **Setup** tab, click **Blink Storage Array**.

- In the AMW, select **Storage Array**→ **Blink**→ **Storage Array**.

The LEDs blink on the physical disks in the storage array.

2 After locating the storage array, click **OK**.

The LEDs stop blinking.

3 If the LEDs do not stop blinking, select **Storage Array**→ **Blink**→ **Stop All Indications**.

A confirmation message is displayed.

4 Click **OK**.

Naming or Renaming Storage Arrays

Each storage array must be assigned a unique up to 30-character alphanumeric name. A name can consist of letters, numbers, and the special characters underscore (_), dash (–), and pound sign (#). No other special characters are allowed.

To rename a selected storage array:

1 Perform one of these actions:

  • In the AMW **Setup** tab, select **Rename Storage Array**.

  • In the EMW **Devices** tab Tree view, select **Edit→ Rename.**

  • In the EMW **Devices** tab Table view, select **Edit→ Rename.**

  • In the EMW **Devices** tab Tree view, right-click on the desired array icon and select **Rename.**

  The **Name/Rename Storage Arrays** dialog is displayed.

2 Select the relevant storage array from the **Select storage array** table.

  If you do not know the name or physical location of the storage array, click **Blink**. After locating the storage array, click **OK** to turn off the LEDs.

  The name of the storage array is displayed in the **Storage array name**.

3 In **Storage array name**, type the new name of the storage array. If applicable, add a comment for the storage array in **Additional comment**.

4 Click **Apply**.

  A message is displayed warning you about the implications of changing the storage array name.

5 Click **Yes**.

  The new storage array name is displayed in the **Select storage array** table.

6 Repeat step 2 through step 4 to name or rename additional storage arrays.

  📝 **NOTE:** Avoid arbitrary names or names that may lose meaning in the future.

Setting a Password

You can configure each storage array with a password to protect it from unauthorized access. MDSM prompts for the password when an attempt is made to change the storage array configuration, such as, when a virtual disk is created or deleted. View operations do not change the storage array configuration and do not require a password. You can create a new password or change an existing password.

NOTE: It is recommended that you use a long password with at least 15 alphanumeric characters to increase security.

To set a new password or change an existing password:

1 Select the relevant storage array and navigate to the AMW for that storage array. See "Array Management Window" on page 67.

The AMW for the selected storage array is displayed.

2 In the AMW, perform one of these actions:

- Select the storage array in the **Logical** pane, and then select **Storage Array→ Set Password**.

- Select the **Setup** tab, and then click **Set a Storage Array Password**.

- Select the **Logical** tab, right-click and select **Set Password.**

The **Set Password** dialog is displayed.

3 If you are resetting the password, type the **Current password**.

NOTE: If you are setting the password for the first time, leave the **Current password** blank.

4 Type the **New password**.

5 Re-type the new password in **Confirm new password**.

NOTE: The password in **Confirm new password** and **New password** must be exactly the same.

6 Click **OK**.

NOTE: You are not prompted for a password when you attempt to change the storage array configuration in the current management session.

Password Guidelines

Follow these guidelines when you create a password:

- Use secure passwords for your storage array. A password must be easy for you to remember but difficult for others to determine. Consider using numbers or special characters in the place of letters, such as a 1 in the place of the letter I, or the at sign (@) in the place of the letter a.

- For increased protection, use a long password with at least 15 alphanumeric characters. The maximum password length is 30 characters.

- Passwords are case sensitive.

NOTE: You can attempt to enter a password up to ten times before the storage array enters a lockout state. Before you can try to enter a password again, you must wait 10 minutes for the storage array to reset. To reset the password, press the password reset switch on your RAID controller module, see Figure 3-1.

Viewing Storage Array Connections

You can use the **View Connections** option to view the expansion enclosures connected to the RAID controller module.

To view the storage array connections:

1   From the toolbar in AMW, select **Storage Array**→ **View**→ **Connections**.

    The <**Storage Array**>**:Connections** dialog is displayed.

2   Click the column name to sort the connections according to your preference.

3   Click **Close**.

If you receive an error message for a port, you can use this dialog to identify the components on the port that may have caused the error. By isolating these components, you prevent accidentally disconnecting components that are still in operation, which could cause an interruption in data flow.

Adding/Editing a Comment to an Existing Storage Array

A descriptive comment, with an applicable storage array name, is a helpful identification tool. You can add or edit a comment for a storage array in the EMW only.

To add or edit a comment:

1   In the EMW, select the **Devices** tab and select the relevant managed storage array.

2   Select **Edit**→ **Comment**.

    The **Edit Comment** dialog is displayed.

3   Type a 60-character comment.

4   Click **OK**.

This option updates the comment in the table view and saves it in your local storage management station file system. The comment is not displayed to administrators who are using other storage management stations.

Removing Storage Arrays

You can remove a storage array from the list of managed arrays if you no longer want to manage it from a specific storage management station. Removing a storage array does not affect the storage array or its data in any way. Removing a storage array simply removes it from the list of storage arrays that are displayed in the drop-down list in the Array Selector. If a storage array is accidentally removed, it can be added again. See "Adding Storage Arrays" on page 70.

You can remove the storage array only from the EMW.

To remove the storage array:

1   In the EMW, select the **Devices** tab and select the relevant managed storage array.

2   Select **Edit**→ **Remove**→ **Storage Array**.

    A message prompts you for a confirmation for the removal of the selected storage array.

3   To remove the storage array, click **Yes**.

Enabling Premium Features

You can enable premium features on the storage array. To enable the premium features, you must obtain a feature key file specific to the premium feature that you want to enable from your storage supplier.

To enable premium features:

1 From the toolbar in AMW, select **Storage Array**→ **Premium Features**.

   The **Premium Features and Feature Pack Information** window is displayed.

2 Select the relevant premium feature, and click **Enable**.

   The **Select Feature Key File** dialog is displayed.

3 Navigate to the relevant folder, select the appropriate key file, and click **OK**.

4 Click **Close**.

For more information, see the *PowerVault Modular Disk Storage Manager online help* topics.

Failover Alert Display

You can change the failover alert delay for a storage array. The failover alert delay lets you delay the logging of a critical event if the multi-path driver transfers virtual disks to the non-preferred controller. If the multi-path driver transfers the virtual disks back to the preferred controller within the specified delay period, a critical event is not logged. If the transfer exceeds this delay period, then a virtual disk-not-on-preferred-path alert is issued as a critical event. You can also use this option to minimize multiple alerts when more than one virtual disk fails over because of a system error, such as a failed host adapter.

For more information, see the *PowerVault Modular Disk Storage Manager online help* topics.

Changing the Cache Settings on the Storage Array

1. In the AMW, select **Storage Array→ Change→ Cache Settings**.

   The **Change Cache Settings** window is displayed.

2. Select or enter the percentage of unwritten data in the cache to trigger a cache flush in **Start flushing**.

3. Select or enter the percentage of unwritten data in the cache to stop a cache flush in progress in **Stop flushing**.

4. Select the appropriate **Cache block size**.

   A smaller cache size is a good choice for file-system use or database-application use. A larger cache size is a good choice for applications that generate sequential I/O, such as multimedia.

5. In the **Enter Password** dialog, type the current password for the storage array, and click **OK**.

Changing Expansion Enclosure ID Number

When an MD1200 series expansion enclosure is connected to an MD3600f series storage array for the first time, an enclosure ID number is assigned and maintained by the expansion enclosure. This enclosure ID number is also displayed in the MDSM.

To change the enclosure ID numbers:

1. In the AMW, select the storage array, and select **Storage Array→ Change→ Enclosure ID**.

2. Select a new enclosure ID number from the **Change Enclosure ID** list.

   The enclosure ID must be between 0 and 99 (inclusive).

3. To save the changed enclosure ID, click **Change**.

Changing the Enclosure Order in the Physical Pane

You can change the order of the RAID controller modules and the expansion enclosures in the **Physical** pane to match the hardware configuration in your storage array. The **Physical** pane that initially is displayed is a default view that may not match your storage array. The enclosure order change remains in effect until it is modified again.

To change the enclosure order in the **Physical** pane:

1 In the AMW, select **Storage Array**→ **Change**→ **Enclosure Order**.

2 From the **Enclosures list,** select the enclosure you want to move and click either **Up** or **Down** to move the enclosure to the new position.

3 Click **OK**.

 If you have set a password for the selected storage array, the **Enter Password** dialog is displayed.

4 Type the current password for the storage array.

5 Click **OK**.

Configuring Alert Notifications

MDSM can send an alert for any condition on the storage array that requires your attention. Alerts can be sent as e-mail messages or as Simple Network Management Protocol (SNMP) trap messages.

You can configure alert notifications either for all the storage arrays or a single storage array.

To configure alert notifications for all storage arrays:

1 In the EMW, select the **Setup** tab.

2 Select **Configure Alerts**.

 The **Configure Alerts** dialog is displayed.

3 Select **All storage arrays**.

4 Click **OK**.

 The **Configure Alerts** dialog is displayed. To configure e-mail alerts, see "Configuring E–mail Alerts" on page 81. To configure SNMP alerts, see "Configuring SNMP Alerts" on page 83.

To configure alert notifications for a single storage array:

1 In the EMW, select the **Devices** tab.

2 Select the relevant storage array, then select **Edit→ Configure→ Alerts**.

The **Configure Alerts** dialog is displayed. To configure e-mail alerts, see "Configuring E–mail Alerts" on page 81. To configure SNMP alerts, see "Configuring SNMP Alerts" on page 83.

Configuring E–mail Alerts

For more information on configuring alert notifications, see "Configuring Alert Notifications" on page 80.

To configure e-mail alerts:

1 Open the **Configure Alerts** dialog by performing one of these actions:

- In the Tree view or the Table view on the **Devices** tab in the EMW, select a node, and then select **Edit→ Configure Alerts**. Go to step 3.

- In the **Setup** tab in the EMW, select **Configure Alerts**. Go to step 2.

2 Select one of the following radio buttons to specify an alert level:

- **All storage arrays**—Select this option to send an alert e-mail about events on all storage arrays.

- **An individual storage array**—Select this option to send an alert e-mail about events that occur on only a specified storage array.

These results occur, depending on your selection:

- If you select all storage arrays, the **Configure Alerts** dialog is displayed.

- If you select an individual storage array, the **Select Storage Array** dialog is displayed. Select the storage array for which you want to receive e-mail alerts and click **OK**. The **Configure Alerts** dialog is displayed.

- If you do not know which storage array to select, click **Blink** to turn on the LEDs of the storage array.

3 In the **Configure Alerts** dialog, select the **Mail Server** tab.

4 In **Mail server**, type the name of the Simple Mail Transfer Protocol (SMTP) mail server.

The SMTP mail server is the name of the mail server that forwards the alert e-mails to the configured e-mail addresses.

5   In **E-mail sender address**, type the valid sender e-mail address.

The e-mail address of the sender (the network administrator) is displayed on each e-mail alert sent to the destination.

6   To include the contact information of the sender in the e-mail alert, select **Include contact information with the alerts**, and type the contact information.

    📝 **NOTE:** Including the contact information in the e-mail alert is optional.

7   Select the **e-mail** tab to configure the e-mail destinations.

    • Adding an e-mail address—In **e-mail address**, type the e-mail address, and click **Add**.

    • Replacing an e-mail address—In the **Configured E-mail addresses** area, select the e-mail address to be replaced, type the replacement e-mail address in **E-mail address**, and click **Replace**.

    • Deleting an e-mail address—In the **Configured E-mail addresses** area, select the e-mail address, and click **Delete**.

    • Validating an e-mail address—Type the e-mail address in **E-mail address** or select the e-mail address in the **Configured E-mail addresses** area, and click **Test**. A test e-mail is sent to the selected e-mail address. A dialog with the results of the test and any error is displayed.

8   For the selected e-mail address, in **Information To Send**, select:

    • **Event Only**—The alert e-mail contains only the event information. This alert type is the default.

    • **Event + Profile**—The alert e-mail contains the event information and the storage array profile.

    • **Event + Support**—The alert e-mail contains the event information and a compressed file that contains complete support information for the storage array that has generated the alert.

9 For the selected e-mail address, in **Frequency**, select:

- **Every event**—Sends an alert e-mail whenever an event occurs. This is the default option.

- **Every x hours**—Sends an alert e-mail after the specified time interval if an event occurred during that time interval. You can select this option only if you have selected either **Event + Profile** or **Event + Support** in the **Information To Send** drop down list.

10 Click **OK**.

An alert icon is displayed next to each node in the Tree view where an alert is set.

To ensure that the e-mail is sent successfully:

- You must provide an SMTP mail server name and an e-mail sender address for the e-mail addresses to work.

- The e-mail addresses that you had previously configured are displayed in the Configured e-mail addresses area.

- You must use fully qualified e-mail addresses, for example, name@mycompany.com.

- You can configure multiple e-mail addresses before you click **OK**.

Configuring SNMP Alerts

To add a management console to the list of addresses configured to receive SNMP alerts:

1 Open the **Configure Alerts** dialog by performing one of these actions:

- In the Tree view or the Table view on the **Devices** tab in the EMW, select a node, and select **Edit→ Configure Alerts**. Go to step 3.

- In the **Setup** tab in the EMW, select **Configure Alerts.** Go to step 2.

2 Select one of the following radio buttons to specify an alert level:

- **All storage arrays**—Select this option to send an alert notification about events on all storage arrays.

- **An individual storage array**—Select this option to send an alert notification about events that occur in only a specified storage array.

These results occur, depending on your selection:

- If you select all storage arrays, the **Configure Alerts** dialog is displayed.

- If you select an individual storage array, the **Select Storage Array** dialog is displayed. Select the storage array for which you want to receive alert notifications and click **OK**. The **Configure Alerts** dialog is displayed.

📝 **NOTE:** If you do not know which storage array to select, click **Blink** to turn on the LEDs of the storage array.

3 Select the **SNMP** tab to configure the SNMP alert destinations.

- Adding an SNMP address—In **Community name**, type the community name. In **Trap destination**, type the trap destination, and click **Add**.

📝 **NOTE:** The community name is an American Standard Code for Information Interchange (ASCII) string that identifies a known set of network management stations and is set by the network administrator. The default community name is the string "public".The trap destination is the IP address or the host name of a computer running an SNMP management application. An example of a SNMP enabled management application is the Dell Management Console. For more information on Dell Management Console, see **dell.com**.

- Replacing an SNMP address—Select the SNMP address in the **Configured SNMP addresses** area, type the replacement community name in **Community name** and the trap destination in **Trap destination**, and click **Replace**.

- Deleting an SNMP address—Select the SNMP address in the **Configured SNMP addresses** area, and click **Delete**.

- Validating an SNMP address—Select the SNMP address in the **Configured SNMP addresses** area, and click **Test**. A test message is sent to the SNMP address. A message box with the results of the validation and any error information is displayed.

4 Click **OK**.

An alert icon is displayed next to each node in the Tree view for which an alert is set.

Follow these guideline for SNMP alerts:

- Any SNMP addresses that you had previously configured are displayed in the Configured SNMP addresses area.
- The SNMP Community Name is determined by the system administrator and configured within the management application, such as the Dell Management Console. More information about the Dell Management Console is available at **dell.com**.
- You can configure multiple SNMP addresses before you click **OK**.

Battery Settings

A smart battery backup unit (BBU) can perform a learn cycle. The smart BBU module includes the battery, a battery gas gauge, and a battery charger. The learn cycle calibrates the smart battery gas gauge so that it provides a measurement of the charge of the battery module. A learn cycle can only start when the battery is fully charged.

The learn cycle completes the following operations:

- Discharges the battery to a predetermined threshold.
- Charges the battery back to full capacity.

A learn cycle starts automatically when you install a new battery module. Learn cycles for batteries in both RAID controller modules in a duplex system occur simultaneously.

Learn cycles are scheduled to start automatically at regular intervals, at the same time and on the same day of the week. The interval between cycles is described in weeks.

Use the following guidelines to adjust the interval:

- You can use the default interval.
- You can run a learn cycle at any time.
- You can set the learn cycle earlier than the currently scheduled time.
- You cannot set the learn cycle to start more than seven days later than the currently scheduled time.

To change the battery settings perform these steps:

1   In the AMW, select **Storage Array**→ **Change**→ **Battery Settings**.

    The **Battery Settings** dialog is displayed.

2   In **Battery location**, select a battery.

3   Check these details about the battery:

    •   Battery status

    •   Battery age

    •   Days until replacement

For more information, see the *PowerVault Modular Disk Storage Manager online help* topics.

Setting the Storage Array RAID Controller Module Clocks

You can use the **Synchronize RAID Controller Module Clocks** option to synchronize the storage array RAID controller module clocks with the storage management station. This option makes sure that the event timestamps written by the RAID controller modules to the Event Log match the event timestamps written to host log files. The RAID controller modules remain available during synchronization.

To synchronize the RAID controller module clocks with the storage management station:

1   In the AMW, select **Storage Array**→ **Synchronize RAID Controller Module Clocks**.

2   If a password is set, in the **Enter Password** dialog, type the current password for the storage array, and click **Synchronize**.

    The RAID controller module clocks are synchronized with the storage management station.

7

Configuration: Event Monitor

An event monitor is provided with Dell PowerVault Modular Disk Storage Manager (MDSM). The event monitor runs continuously in the background and monitors activity on the managed storage arrays. If the event monitor detects any critical problems, it can notify a host or remote system using e-mail, Simple Network Management Protocol (SNMP) trap messages, or both.

For the most timely and continuous notification of events, enable the event monitor on a management station that runs 24 hours a day. Enabling the event monitor on multiple systems or having a combination of an event monitor and MDSM active can result in duplicate events, but this does not indicate multiple failures on the array.

To use the Event Monitor:

* Set up alert destinations for the managed device that you want to monitor. A possible alert destination would be the Dell Management Console. More information about the Dell Management Console can be found at **dell.com**.

* Replicate the alert settings from a particular managed device by copying the **emwdata.bin** file to every storage management station from which you want to receive alerts.

Each managed device shows a check mark that indicates that alerts are set.

Enabling or Disabling the Event Monitor

You can enable or disable the event monitor at any time.

Disable the event monitor if you do not want the system to send alert notifications. If you are running the event monitor on multiple systems, disabling the event monitor on all but one system prevents the sending of duplicate messages.

✎ **NOTE:** It is recommended that you configure the event monitor to start by default on a management station that runs 24 hours a day.

Windows

To enable or disable the event monitor:

1   Click **Start**→ **Administrative Tools**→ **Services**.

    or

    Click **Start**→ **Settings**→ **Control Panel**→ **Administrative Tools**→ **Services**.

2   From the list of services, select **Modular Disk Storage Manager Event Monitor**.

3   Select **Action**→ **Properties**.

4   To enable the event monitor, in the **Service Status** area, click **Start**.

5   To disable the event monitor, in the **Service Status** area, click **Stop**.

Linux

To enable the event monitor, at the command prompt, type `SMmonitor start` and press < Enter>. When the program startup begins, the system displays the following message:

`SMmonitor started.`

To disable the event monitor, start terminal emulation application (console ox xterm) and at the command prompt, type `SMmonitor stop` and press < Enter>. When the program shutdown is complete, the following message is displayed:

`Stopping Monitor process.`

8

Configuration: About Your Host

Configuring Host Access

Dell PowerVault Modular Disk Storage Manager (MDSM) software is comprised of multiple modules. One of these modules is the Host Context Agent. The host context agent is installed as part of the MDSM installation and runs continuously in the background.

If the host context agent is running on a host, the host and the host ports connected from it to the storage array are automatically detected by MDSM and are displayed on the **Mappings** tab in the **Array Management Window** (AMW). For more information on the **Mappings** tab, see "Using the Mappings Tab" on page 90.

If the hosts are not detected automatically, you can use the **Define Host Wizard** to define the hosts that access the virtual disks in the storage array. Defining a host is one of the steps required to let the storage array know which hosts are attached to it and to allow access to the virtual disks. For more information on defining the hosts, see "Defining a Host" on page 90.

To enable the host to write to the storage array, you must map the host to the virtual disk. This mapping grants a host or a host group access to a particular virtual disk or to a number of virtual disks in a storage array. You can define the mappings on the **Mappings** tab in the AMW.

On the **Summary** tab in the AMW, the **Hosts & Mappings** area indicates how many hosts are configured to access the storage array. You can click **Configured Hosts** in the **Hosts & Mappings** area to see the names of the hosts.

A collection of elements, such as default host groups, hosts, and host ports, are displayed as nodes in the **Topology** pane of the **Mappings** tab in the AMW.

The host topology is reconfigurable. You can perform the following tasks:

- Create a host and assign an alias or user label.
- Add or associate a new host port identifier to a particular host.
- Change the host port identifier alias or user label.
- Move or associate a host port identifier to a different host.
- Replace a host port identifier with a new host port identifier.
- Manually activate an inactive host port so that the port can gain access to host specific or host group specific LUN mappings.
- Change the host port type to another type.
- Move a host from one host group to another host group.
- Remove a host group, a host, or a host port identifier.
- Rename a host group or a host.

Using the Mappings Tab

In the **Mappings** tab, you can:

- Define hosts and hosts groups
- Add mappings to the selected host groups

For more information, see the *PowerVault Modular Disk Storage Manager online help* topics.

Defining a Host

You can use the Define Host Wizard in the AMW to define a host for a storage array. Either a known unassociated host port identifier or a new host port identifier can be added.

> **NOTE:** A user label must be specified before the host port identifier may be added (the add button is disabled until one is entered). This user label cannot be the same as the host name. It is recommended that you use the host name with HBA port number. For example: Server1-HBA-0.

To define a host:

1 In the AMW, select the **Mappings** tab and select the appropriate storage array.

2 Perform one of the actions:

- Select **Mappings**→ **Define**→ **Host**.

- Select the **Setup** tab, and click **Manually Define Hosts**.

- Select the **Mappings** tab. Right-click the root node (storage array name), **Default Group** node, or **Host Group** node in the **Topology** pane to which you want to add the host, and select **Define**→ **Host** from the pop-up menu.

The **Specify Host Name** window is displayed.

3 In **Host name**, enter an up to 30-character alphanumeric name.

4 Select the relevant option in **Do you plan to use the storage partitions in the this storage array**? and click **Next**.

The **Specify Host Port Identifiers** window is displayed.

5 Select the relevant option to add a host port identifier to the host, you can select:

- **Add by selecting a known unassociated host port identifier**—In **Known unassociated host port identifiers**, select the relevant host port identifier.

- **Add by creating a new host port identifier**—In **New host port identifier**, enter a 16 hexadecimal character name and an up to 30 character **Alias** for the host port identifier and click **Add**.

> ✐ **NOTE:** The host port identifier name is in hexadecimal and must contain the letters A through F and numbers 0 through 9.

6 Click **Next**.

The **Specify Host Type** window is displayed.

7 In **Host** type, select the relevant operating system for the host.

The **Host Group Question** window is displayed.

8 In this window, you can select:

- **Yes**—this host shares access to the same virtual disks with other hosts.

- **No**—this host does NOT share access to the same virtual disks with other hosts.

9   Click **Next**.

   If you select **Yes**, the **Specify Host Group** window is displayed. If you select **No**, see step 11

10   Enter the name of the host group or select an existing host group and click **Next**.

   The **Preview** window is displayed.

11   Click **Finish**.

Removing Host Access

1   In the AMW, select the **Mappings** tab, select the host node in the **Topology** pane.

2   Perform one of these actions:
   - Select **Mappings**→ **Remove**.
   - Right-click the host node and select **Remove** from the pop-up menu.

   The **Remove confirmation** dialog is displayed.

3   Type **yes**.

4   Click **OK**.

For more information, see the *PowerVault Modular Disk Storage Manager online help* topics.

Managing Host Groups

A host group is a logical entity of two or more hosts that share access to specific virtual disks on the storage array. You create host groups with MDSM.

All hosts in a host group must have the same host type (operating system). In addition, all hosts in the host group must have special software, such as clustering software, to manage virtual disk sharing and accessibility.

If a host is part of a cluster, every host in the cluster must be connected to the storage array, and every host in the cluster must be added to the host group.

Creating a Host Group

1   In the AMW, select the **Mappings** tab.

2   In the **Topology** pane, select the storage array or the **Default Group**.

3   Perform one of the following actions:

- •   Select **Mappings→ Define→ Host Group**
- •   Right-click the storage array or the **Default Group,** and select **Define→ Host Group** from the pop-up menu.

4   Type the name of the new host group in **Enter new host group name**.

5   Select the appropriate hosts in the **Select hosts to add** area.

6   Click **Add**.

> **NOTE:** To remove hosts, select the hosts in the **Hosts in group** area, and click **Remove**.

7   Click **OK**.

The host group is added to the storage array.

Adding a Host to a Host Group

You can add a host to an existing host group or a new host group using the **Define Host Wizard**. For more information, see "Defining a Host" on page 90.

You can also move a host to a different host group. For more information, see "Moving a Host to a Different Host Group" on page 94.

Removing a Host From a Host Group

You can remove a host from the **Topology** pane on the **Mappings** tab of the **Array Management Window**. For more information, see "Removing a Host Group" on page 94.

Moving a Host to a Different Host Group

1   In the AMW, select the **Mappings** tab, select the host node in the
    Topology pane.

2   Perform one of these actions:

    •   Select **Mappings→ Move**.

    •   Right-click the host node, and select **Move** from the pop-up menu.

    The **Move Host** dialog is displayed.

3   In the **Select host group**, select the host group to which you want to move
    the host.

    The **Move Host Confirmation** dialog is displayed.

4   Click **Yes**.

    The host is moved to the selected host group with the following mappings:

    •   The host retains the specific virtual disk mappings assigned to it.

    •   The host inherits the virtual disk mappings assigned to the host group
        to which it is moved.

    •   The host loses the virtual disk mappings assigned to the host group
        from which it was moved.

Removing a Host Group

1   In the AMW, select the **Mappings** tab, select the host node in the
    Topology pane.

2   Perform one of these actions:

    •   Select **Mappings→ Remove**.

    •   Right-click the host node, and select **Remove** from the pop-up menu.

    The **Remove** dialog is displayed.

3   Click **Yes**.

    The selected host group is removed.

For more information, see the *PowerVault Modular Disk Storage Manager
online help* topics.

Host Topology

Host topology is the organization of hosts, host groups, and host interfaces configured for a storage array. You can view the host topology in the **Mappings** tab of the AMW. For more information, see "Using the Mappings Tab" on page 90.

The following tasks change the host topology:

- Moving a host or a host connection
- Renaming a host group, a host, or a host connection
- Adding a host connection
- Replacing a host connection
- Changing a host type

MDSM automatically detects these changes for any host running the host agent software.

Starting or Stopping the Host Context Agent

The host context agent discovers the host topology and starts and stops with the host. The topology discovered by the host context agent can be viewed by clicking **Configure Host Access (Automatic)** in the **Configure** tab in the MDSM.

You must stop and restart the host context agent to see the changes to the host topology if:

- A new storage array is attached to the host server.
- A host is added while turning on power to the RAID controller modules.

Linux

To start or stop the host context agent, enter the following command at the prompt:

```
SMagent start
SMagent stop
```

You stop and then restart SMagent after:

- Moving a controller offline or replacing a controller.
- Removing host-to-array connections from or attaching host-to-array connections to a Linux host server.

Windows

To start or stop the host context agent:

1 Click **Start**→ **Settings**→ **Control Panel**→ **Administrative Tools**→ **Services**.

 or

 Click **Start**→ **Administrative Tools**→ **Services**.

2 From the list of services, select **Modular Disk Storage Manager Agent**.

3 If the host context agent is running, click **Action**→ **Stop**, then wait approximately 5 seconds.

4 Click **Action**→ **Start**.

I/O Data Path Protection

You can have multiple host-to-array connections for a host. Ensure that you select all the connections to the array when configuring host access to the storage array.

📝 **NOTE:** See the **Deployment Guide** for more information on cabling configurations, at **dell.com/support/manuals**.

📝 **NOTE:** For maximum redundancy, you must select all host connections to the array when manually defining host topology. For example, a host may have two host connections listed when manually configuring host access. For this host, you would select the two host connections listed in the **Available hosts** section and add them to the **Selected hosts** section using the **Add** button.

📝 **NOTE:** For more information on configuring hosts see "Configuration: About Your Host" on page 89.

If a component such as a RAID controller module or a cable fails, or an error occurs on the data path to the preferred RAID controller module, the virtual disk ownership is moved to the alternate non-preferred RAID controller module for processing. This process is called failover.

Drivers for multi-path frameworks such as Microsoft Multi-Path IO (MPIO) and Linux Device Mapper (DM) are installed on host systems that access the storage array and provide I/O path failover.

For more information on Linux DM, see "Configuration: Device Mapper Multipath for Linux" on page 219. For more information on MPIO, see **microsoft.com**.

NOTE: You must have the multi-path driver installed on the hosts at all times, even in a configuration where there is only one path to the storage system, such as a single port cluster configuration.

During a failover, the virtual disk transfer is logged as a critical event, and an alert notification is sent automatically if you have configured alert destinations for the storage array.

Managing Host Port Identifiers

You can manage the host port identifiers that are added to the storage array. You can:

- Add—Add or associate a new host port identifier to a particular host.
- Edit—Change the host port identifier alias or user label. You can move (associate) the host port identifier to a new host.
- Replace—Replace a particular host port identifier with another host port identifier.
- Remove—Remove the association between a particular host port identifier and the associated host.

NOTE: User label cannot be the same as the host name. It is recommended that you use the host name with HBA port number. For example: Server1-HBA-0.

To manage a host port identifier:

1   Perform one of these actions:

    •   Right-click the host in the **Topology** pane, and select **Manage Host Port Identifiers** in the pop-up menu.

    •   From the menu bar, select **Mappings**→ **Manage Host Port Identifiers**.

    The **Manage Host Port Identifiers** dialog is displayed. You can choose to manage the host port identifiers for a specific host or all of the host port identifiers for all of the hosts in **Show host port identifiers associated with**.

2   If you want to manage the host port identifiers for a specific host, select the host from the list of hosts that are associated with the storage array. If you want to manage the host port identifiers for all hosts, select **All hosts** from the list of hosts that are associated with the storage array.

3   If you are adding a new host port identifier, go to step 4. If you are managing an existing host port identifier, go to step 8.

4   Click **Add**.

    The **Add Host Port Identifier** dialog is displayed.

5   Select the method to add a host port identifier to the host. You can select:

    •   **Add by selecting a know unassociated host port identifier**—Select the appropriate host port identifier from the existing list of **Known unassociated host port identifiers**.

    •   **Add by creating a new host port identifier**—In **New host port identifier**, enter the name of the new host port identifier.

6   In **User label**, enter up to 30–character alphanumeric name.

7   In **Associated with host**, select the appropriate host or host group.

8 Select the host port identifier that you would like to manage from the list of host port identifiers in the **Host port identifier information** area.

9 Perform one of these actions for the selected host port identifier:

- To edit the host port identifier—Select the appropriate host port identifier and click **Edit**, the **Edit Host Port Identifier** dialog is displayed, update **User label** and **Associated with host** and click **Save**.

- To replace the host port identifier—Select the appropriate host port identifier and click **Replace**, the **Replace Host Port Identifier** dialog is displayed, replace the current host port identifier with a known unassociated host port identifier or create a new host port identifier, update **User label** and click **Replace**.

- To remove the host port identifier—Select the appropriate host port identifier and click **Edit**, the **Remove Host Port Identifier** dialog is displayed, type **yes** and click **OK**.

For more information, see the *PowerVault Modular Disk Storage Manager online help* topics.

9

Configuration: Disk Groups and Virtual Disks

Creating Disk Groups and Virtual Disks

Disk groups are created in the unconfigured capacity of a storage array, and virtual disks are created in the free capacity of a disk group. The maximum number of physical disks supported in a disk group is 30. The hosts attached to the storage array read and write data to the virtual disks.

📝 **NOTE:** Before you can create virtual disks, you must first organize the physical disks into disk groups and configure host access. Then you can create virtual disks within a disk group.

To create a virtual disk, use one of the following methods:

- Create a new disk group from unconfigured capacity. You first define the RAID level and free capacity (available storage space) for the disk group, and then you define the parameters for the first virtual disk in the new disk group.

- Create a new virtual disk in the free capacity of an existing disk group. You only need to specify the parameters for the new virtual disk.

A disk group has a set amount of free capacity that is configured when the disk group is created. You can use that free capacity to subdivide the disk group into one or more virtual disks.

You can create disk groups and virtual disks using:

- Automatic configuration—Provides the fastest method, but with limited configuration options.

- Manual configuration—Provides more configuration options.

When creating a virtual disk, consider the uses for that virtual disk, and select an appropriate capacity for those uses. For example, if a disk group has a virtual disk that stores multimedia files (which tend to be large) and another virtual disk that stores text files (which tend to be small), the multimedia file virtual disk requires more capacity than the text file virtual disk.

A disk group must be organized according to its related tasks and subtasks. For example, if you create a disk group for the Accounting Department, you can create virtual disks that match the different types of accounting transactions performed in the department: Accounts Receivable (AR), Accounts Payable (AP), internal billing, and so forth. In this scenario, the AR and AP virtual disks probably need more capacity than the internal billing virtual disk.

NOTE: In Linux, the host must be rebooted after deleting virtual disks to reset the /dev entries.

NOTE: Before you can use a virtual disk, you must register the disk with the host systems. See "Host-to-Virtual Disk Mapping" on page 125.

Creating Disk Groups

You can create disk groups using either Automatic configuration or Manual configuration.

To create disk groups using automatic configuration:

1 To start the **Create Disk Group Wizard**, perform one of these actions:

   • To create a disk group from unconfigured capacity in the storage array—On the **Logical** tab, select an **Unconfigured Capacity** node, and select **Disk Group**→ **Create**. Alternatively, you can right-click the **Unconfigured Capacity** node, and select **Create Disk Group** from the pop-up menu.

   • To create a disk group from unassigned physical disks in the storage array—On the **Physical** tab, select one or more unassigned physical disks of the same physical disk type, and select **Disk Group**→ **Create**. Alternatively, you can right-click the unassigned physical disks, and select **Create Disk Group** from the pop-up menu.

   • To create a secure disk group—On the **Physical** tab, select one or more unassigned security capable physical disks of the same physical disk type, and select **Disk Group**→ **Create**. Alternatively, you can right-click the unassigned security capable physical disks, and select **Create Disk Group** from the pop-up menu.

   The I**ntroduction (Create Disk Group)** window is displayed.

2 Click **Next**.

   The **Disk Group Name and Physical Disk Selection** window is displayed.

3 Type the name of the disk group (up to 30 characters) in **Disk group name**.

4 Select the appropriate **Physical Disk selection choices**, you can select:

- **Automatic**, see step 6
- **Manual**, see step 9

5 Click **Next**.

6 For Automatic configuration, the **RAID Level and Capacity** window is displayed.

7 Select the appropriate RAID level in **Select RAID level**. You can select RAID levels 0, 1/10, 6, and 5.

Depending on your RAID level selection, the physical disks available for the selected RAID level is displayed in **Select capacity** table.

8 In the **Select Capacity** table, select the relevant disk group capacity, and click **Finish**.

9 For Manual configuration, the **Manual Physical Disk Selection** window is displayed.

10 Select the appropriate RAID level in **Select RAID level**. You can select RAID levels 0, 1/10, 6, and 5.

Depending on your RAID level selection, the physical disks available for the selected RAID level is displayed in **Unselected physical disks** table.

11 In the **Unselected physical disks** table, select the appropriate physical disks and click **Add**.

> **NOTE:** You can select multiple physical disks at the same time by holding <Ctrl> or <Shift> and selecting additional physical disks.

12 To view the capacity of the new disk group, click **Calculate Capacity**.

13 Click **Finish**.

A message prompts you that the disk group is successfully created and that you must create at least one virtual disk before you can use the capacity of the new disk group. For more information on creating virtual disks, see "Creating Virtual Disks" on page 104.

Locating a Disk Group

You can physically locate and identify all of the physical disks that comprise a selected disk group. An LED blinks on each physical disk in the disk group.

To locate a disk group:

1 In the AMW, select the **Logical** tab.

2 Select the appropriate disk group and from the toolbar select **Disk Group→ Blink**.

   The LEDs for the selected disk group blink.

3 After locating the disk group, click **OK**.

   The LEDs stop blinking.

4 If the LEDs for the disk group do not stop blinking, from the toolbar in AMW, select **Storage Array→ Blink→ Stop All Indications**.

   If the LEDs successfully stop blinking, a confirmation message is displayed.

5 Click **OK**.

Creating Virtual Disks

Keep these important guidelines in mind when you create a virtual disk:

- Many hosts can have 256 logical unit numbers (LUNs) mapped per storage partition, but the number varies per operating system.

- After you create one or more virtual disks and assign a mapping, you must register the virtual disk with the operating system. In addition, you must ensure that the host recognizes the mapping between the physical storage array name and the virtual disk name. Depending on the operating system, run the host-based utilities, **hot_add** and **SMdevices**.

- If the storage array contains physical disks with different media types or different interface types, multiple Unconfigured Capacity nodes may be displayed in the **Logical** pane of the **Logical** tab. Each physical disk type has an associated Unconfigured Capacity node if unassigned physical disks are available in the expansion enclosure.

- You cannot create a disk group and subsequent virtual disk from different physical disk technology types. Each physical disk that comprises the disk group must be of the same physical disk type.

NOTE: Ensure that you create disk groups before creating virtual disks.

To create virtual disks:

1 Choose one of these methods to start the **Create Virtual Disk Wizard**:

- To create a virtual disk from unconfigured capacity in the storage array—On the **Logical** tab, select an **Unconfigured Capacity** node, and select **Virtual Disk→ Create**. Alternatively, you can right-click the **Unconfigured Capacity** node, and select **Create Virtual Disk** from the pop-up menu.

- To create a virtual disk from free capacity on a disk group—On the **Logical** tab, select a **Free Capacity** node, and select **Virtual Disk→ Create**. Alternatively, you can right-click the **Free Capacity** node, and select **Create Virtual Disk** from the pop-up menu.

- To create a virtual disk from unassigned physical disks in the storage array—On the **Physical** tab, select one or more unassigned physical disks of the same physical disk type, and select **Virtual Disk→ Create**. Alternatively, you can right-click the unassigned physical disks, and select **Create Virtual Disk** from the pop-up menu.

- To create a secure virtual disk—On the **Physical** tab, select one or more unassigned security capable physical disks of the same physical disk type, and select **Virtual Disk→ Create**. Alternatively, you can right-click the unassigned security capable physical disks, and select **Create Virtual Disk** from the pop-up menu.

If you selected an **Unconfigured Capacity** node or unassigned physical disks to create a virtual disk, the **Disk Group Required** dialog is displayed. Click **Yes** and create a disk group by using the **Create Disk Group Wizard**. The **Create Virtual Disk Wizard** is displayed after you create the disk group.

If you selected a **Free Capacity** node, the The **Introduction (Create Virtual Disk)** window is displayed.

2 Click **Next**.

The **Specify Capacity/Name** window is displayed.

3 Select the appropriate unit for memory in **Units** and enter the capacity of the virtual disk in **New virtual disk capacity**.

4 Enter an up to 30–character name for the virtual disk in **Virtual disk name**.

5 In **Advanced virtual disk parameters**, you can select:

- **Use recommended settings**.
- **Customize settings**.

6 Click **Next**.

7 In the **Customize Advanced Virtual Disk Parameters** window, select the appropriate **Virtual Disk I/O characteristics** type. You can select:

- **File system (typical)**
- **Database**
- **Multimedia**
- **Custom**

   NOTE: If you select **Custom**, you must select an appropriate segment size.

8 Select the appropriate **Preferred RAID controller module ownership** and click **Next**.

   The **Specify Virtual Disk-to-LUN Mapping** window is displayed.

9 In the **Specify Virtual Disk-to-LUN Mapping** window, you can select:

- **Default Mapping**
- The appropriate **Host type**
- **Map later Using Mappings View**.

10 Click **Finish**.

   The virtual disks are created.

Changing the Virtual Disk Modification Priority

You can specify the modification priority setting for a single virtual disk or multiple virtual disks on a storage array.

Guidelines to change the modification priority of a virtual disk:

- If more than one virtual disk is selected, the modification priority defaults to the lowest priority. The current priority is shown only if a single virtual disk is selected.

- Changing the modification priority by using this option modifies the priority for the selected virtual disks.

To change the virtual disk modification priority:

1 In the AMW, select the **Logical** tab.

2 Select a virtual disk.

3 In the toolbar, select **Virtual Disk**→ **Change**→ **Modification Priority**.

   The **Change Modification Priority** window is displayed.

4 Select one or more virtual disks.

   To select nonadjacent virtual disks, press < Ctrl> click. To select adjacent virtual disks, press < Shift> click. To select all of the available virtual disks, click **Select All**. Move the **Select modification priority** slider bar to the desired priority.

5 Click **OK**.

   A message prompts you to confirm the change in the virtual disk modification priority.

6 Click **Yes**.

7 Click **OK**.

Changing the Virtual Disk Cache Settings

You can specify the cache memory settings for a single virtual disk or for multiple virtual disks on a storage array.

Guidelines to change cache settings for a virtual disk:

- After opening the **Change Cache Settings** dialog, the system may prompt you that the RAID controller module has temporarily suspended caching operations. This action may occur when a new battery is charging, when a RAID controller module is removed, or if a mismatch in cache sizes is detected by the RAID controller module. After the condition is cleared, the cache properties selected in the dialog become active. If the selected cache properties do not become active, contact your Technical Support representative.

- If you select more than one virtual disk, the cache settings default to no settings selected. The current cache settings are displayed only if you select a single virtual disk.

- If you change the cache settings by using this option, the priority of all of the virtual disks that you selected is modified.

To change the virtual disk cache settings:

1   In the AMW, select the **Logical** tab and select a virtual disk.

2   In the toolbar, select **Virtual Disk→ Change→ Cache Settings**.

    The **Change Cache Settings** window is displayed.

3   Select one or more virtual disks.

    To select nonadjacent virtual disks, press < Ctrl> click. To select adjacent
    virtual disks, press < Shift> click. To select all of the available virtual disks,
    click **Select All**.

4   In the **Select cache properties** area, you can select:

    •   **Enable read caching**—to enable read caching.

    •   **Enable dynamic cache read prefetch**—to enable dynamic cache
        read prefetch.

    •   **Enable write caching**—to enable write caching.

        •   **Enable write caching with mirroring**—to mirror cached data
            across two redundant RAID controller modules that have the
            same cache size.

        •   **Enable write caching without batteries**—to permit write caching
            to continue even if the RAID controller module batteries are
            discharged completely, not fully charged, or are not present.

⚠ CAUTION: Possible loss of data – Selecting the Enable write caching without
batteries option allows write caching to continue even when the batteries are
discharged completely or are not fully charged. Typically, write caching is turned off
temporarily by the RAID controller module until the batteries are charged. If you
select this option and do not have a universal power supply for protection, you could
lose data. In addition, you could lose data if you do not have RAID controller module
batteries and you select the Enable write caching without batteries option.

📝 NOTE: When the Optional RAID controller module batteries option is enabled,
the Enable write caching is not displayed. The Enable write caching without
batteries is still available, but it is not checked by default.

📝 NOTE: Cache is automatically flushed after the Enable write caching check
box is disabled.

5 Click **OK**.

A message prompts you to confirm the change in the virtual disk modification priority.

6 Click **Yes**.

7 Click **OK**.

Changing the Segment Size of a Virtual Disk

You can change the segment size on a selected virtual disk. During this operation, I/O performance is affected, but your data remains available.

Guidelines to proceed with changing the segment size:

- You cannot cancel this operation after it starts.
- Do not start this operation unless the disk group is in Optimal status.
- MDSM determines the segment size transitions that are allowed. Segment sizes that are inappropriate transitions from the current segment size are unavailable on the menu. Allowed transitions usually are double or half of current segment size. For example, if the current virtual disk segment size is 32 KB, a new virtual disk segment size of either 16 KB or 64 KB is allowed.

🖉 **NOTE:** The operation to change the segment size is slower than other modification operations (for example, changing RAID levels or adding free capacity to a disk group). This slowness is the result of how the data is reorganized and the temporary internal backup procedures that occur during the operation.

The amount of time that a change segment size operation takes depends on:

- The I/O load from the host
- The modification priority of the virtual disk
- The number of physical disks in the disk group
- The number of physical disk ports
- The processing power of the storage array RAID controller modules

If you want this operation to complete faster, you can change the modification priority, although this may decrease system I/O performance.

To change the segment size of a virtual disk:

1 In the AMW, select the **Logical** tab and select a virtual disk.

2 Select **Virtual Disk**→ **Change**→ **Segment Size**.

3 Select the required segment size.

A message prompts you to confirm the selected segment size.

4 Click **Yes**.

The segment size modification operation begins. The virtual disk icon in the **Logical** pane shows an **Operation in Progress** status while the operation is taking place.

✍ **NOTE:** To view the progress or change the priority of the modification operation, select a virtual disk in the disk group, and select **Virtual Disk**→ Change→ Modification Priority.

Changing the I/O Type

You can specify the virtual disk I/O characteristics for the virtual disks that you are defining as part of the storage array configuration. The expected I/O characteristics of the virtual disk is used by the system to indicate an applicable default virtual disk segment size and dynamic cache read prefetch setting. For more information about the Automatic Configuration Wizard, see the *PowerVault Modular Disk Storage Manager online help* topics.

✍ **NOTE:** The dynamic cache read prefetch setting can be changed later by selecting **Virtual Disk**→ Change→ Cache Settings. You can change the segment size later by selecting **Virtual Disk**→ Change→ Segment Size.

The I/O characteristic types shown below are only presented during the create virtual disk process.

When you choose one of the virtual disk I/O characteristics, the corresponding dynamic cache prefetch setting and segment size that are typically well suited for expected I/O patterns are populated in the **Dynamic cache read prefetch** field and the **Segment size** field.

To change the I/O type:

1   Select from these virtual disk I/O characteristic types, based on your application needs:

   • **File system (typical)**

   • **Database**

   • **Multimedia**

   • **Custom**

   The corresponding dynamic cache read prefetch setting and segment size values that are typically well suited for the selected virtual disk I/O characteristic type are populated in the **Dynamic cache read prefetch** and **Segment size** fields.

   > **NOTE:** If you selected the Custom option, select your preferred dynamic cache read prefetch setting (enabled/disabled) and segment size (8 KB to 512 KB).

2   Click **OK**.

Choosing an Appropriate Physical Disk Type

You can create disk groups and virtual disks in the storage array. You must select the capacity that you want to allocate for the virtual disk from either unconfigured capacity or free capacity available in the storage array. Then you define basic and optional advanced parameters for the virtual disk.

With the advent of different physical disk technologies, it is now possible to mix physical disks with different media types and different interface types within a single storage array. In this release of MDSM, the following media types are supported:

   • Hard physical disk

   • Solid State Disk (SSD)

Physical Disk Security With Self Encrypting Disk

Self encrypting disk (SED) technology prevents unauthorized access to the data on a physical disk that is physically removed from the storage array. The storage array has a security key. Self encrypting disks provide access to data only through an array that has the correct security key.

The self encrypting disk or a security capable physical disk encrypts data during writes and decrypts data during reads. For more information, see the *PowerVault Modular Disk Storage Manager online help* topics.

You can create a secure disk group from security capable physical disks. When you create a secure disk group from security capable physical disks, the physical disks in that disk group become security enabled. When a security capable physical disk is security enabled, the physical disk requires the correct security key from a RAID controller module to read or write the data. All of the physical disks and RAID controller modules in a storage array share the same security key. The shared security key provides read and write access to the physical disks, while the physical disk encryption key on each physical disk is used to encrypt the data. A security capable physical disk works like any other physical disk until it is security enabled.

Whenever the power is turned off and turned on again, all of the security-enabled physical disks change to a security locked state. In this state, the data is inaccessible until the correct security key is provided by a RAID controller module.

You can view the self encrypting disk status of any physical disk in the storage array from the **Physical Disk Properties** dialog. The status information reports whether the physical disk is:

- Security Capable
- Secure—Security enabled or disabled
- Read/Write Accessible—Security locked or unlocked

You can view the self encrypting disk status of any disk group in the storage array. The status information reports whether the storage array is:

- Security Capable
- Secure

Table 9-1 shows how to interpret the security status of a disk group.

Table 9-1.   Interpreting Security Status of a Disk Group

Secure	Security Capable–Yes	Security Capable–No
Yes	The disk group is composed of all SED physical disks and is in a Secure state.	Not applicable. Only SED physical disks can be in a Secure state.
No	The disk group is composed of all SED physical disks and is in a Non-Secure state.	The disk group is not entirely composed of SED physical disks.

The **Physical Disk Security** menu is displayed in the **Storage Array** menu. The P**hysical Disk Security** menu has these options:

- **Create Security Key**
- **Change Security Key**
- **Save Security Key File**
- **Validate Security Key**
- **Unlock Drives**

📝 **NOTE:** If you have not created a security key for the storage array, the **Create Security Key** option is active. If you have created a security key for the storage array, the **Create Security Key** option is inactive with a check mark to the left. The **Change Security Key** option, the **Save Security Key** option, and the **Validate Security Key** option are now active.

The **Secure Physical Disks** option is displayed in the **Disk Group** menu. The **Secure Physical Disks** option is active if these conditions are true:

- The selected storage array is not security enabled but is comprised entirely of security capable physical disks.
- The storage array contains no snapshot source virtual disks or snapshot repository virtual disks.
- The disk group is in an **Optimal** state.
- A security key is set up for the storage array.

📝 **NOTE:** The **Secure Physical Disks** option is inactive if these conditions are not true.

The **Secure Physical Disks** option is inactive with a check mark on the left if the disk group is already security enabled.

The **Create a secure disk group** option is displayed in the **Create Disk Group Wizard** - **Disk Group Name** and **Physical Disk Selection** dialog. The **Create a secure disk group** option is active only when these conditions are met:

- A security key is installed in the storage array.
- At least one security capable physical disk is installed in the storage array.
- All of the physical disks that you selected on the **Physical** tab are security capable physical disks.

You can erase security-enabled physical disks so that you can reuse the drives in another disk group or in another storage array. When you erase security-enabled physical disks, ensure that the data cannot be read. When all of the physical disks that you have selected in the **Physical** tab are security enabled, and none of the selected physical disks is part of a disk group, the **Secure Erase** option is displayed in the **Physical Disk** menu.

The storage array password protects a storage array from potentially destructive operations by unauthorized users. The storage array password is independent from self encrypting disk and must not be confused with the pass phrase that is used to protect copies of a security key. It is recommended that you set a storage array password.

Creating a Security Key

When you create a security key, it is generated by and securely stored by the array. You cannot read or view the security key.

✎ **NOTE:** A copy of the security key must be kept on some other storage medium for backup in case of system failure or for transfer to another storage array.

A pass phrase that you provide is used to encrypt and decrypt the security key for storage on other media. When you create a security key, you also provide information to create a security key identifier. Unlike the security key, you can read or view the security key identifier. The security key identifier is also stored on a physical disk or transportable media. The security key identifier is used to identify which key the storage array is using.

To create a security key:

1 In the AMW toolbar, select **Storage Array**→ **Physical Disk Security**→ **Create Security Key**.

The **Create Security Key** window is displayed.

2 In **Security key identifier**, enter a string that becomes part of the secure key identifier.

You can enter up to 189 alphanumeric characters without spaces, punctuation, or symbols. Additional characters is generated automatically and is appended to the end of the string that you enter. The generated characters help to ensure that the secure key identifier is unique.

3 Edit the default path by adding a file name to the end of the path or click **Browse,** navigate to the required folder and enter the name of the file.

4 In **Pass phrase**, enter a string for the pass phrase.

The pass phrase must:

- be between 8 and 32 characters long.
- contain at least 1 uppercase letter.
- contain at least 1 lowercase letter.
- contain at least 1 number.
- contain at least 1 non-alphanumeric character, for example, $< > @ +$.

The pass phrase that you enter is masked.

NOTE: Create Key is active only if the pass phrase meets the above mentioned criterion.

5 In **Confirm pass phrase**, re-enter the exact string that you entered in **Pass phrase**.

Make a record of the pass phrase that you entered and the security key identifier that is associated with the pass phrase. You need this information for later secure operations.

6 Click **Create Key**.

7 Make a record of the security key identifier and the file name from the **Create Security Key Complete** dialog, and click **OK**.

After you have created a security key, you can create secure disk groups from security capable physical disks. Creating a secure disk group makes the physical disks in the disk group security-enabled. Security-enabled physical disks enter **Security Locked** status whenever power is re-applied. They can be unlocked only by a RAID controller module that supplies the correct key during physical disk initialization. Otherwise, the physical disks remain locked, and the data is inaccessible. The **Security Locked** status prevents any unauthorized person from accessing data on a security-enabled physical disk by physically removing the physical disk and installing the physical disk in another computer or storage array.

Changing a Security Key

When you change a security key, a new security key is generated by the system. The new key replaces the previous key. You cannot view or read the key. However, a copy of the security key must be kept on some other storage medium for backup in case of system failure or for transfer to another storage array. A pass phrase that you provide encrypts and decrypts the security key for storage on other media.

When you change a security key, you also provide information to create a security key identifier.

Changing the security key does not destroy any data. You can change the security key at any time.

Before you change the security key, ensure that:

- All virtual disks in the storage array are in **Optimal** status.
- In storage arrays with two RAID controller modules, both are present and working normally.

To change the security key:

1   In the AMW toolbar, select **Storage Array**→ **Physical Disk Security**→ **Change Security Key**.

    The **Confirm Change Security Key** window is displayed.

2   Type **yes** in the text field, and click **OK**.

    The **Change Security Key** window is displayed.

3 In **Secure key identifier**, enter a string that becomes part of the secure key identifier.

You may leave the text box blank, or enter up to 189 alphanumeric characters without white space, punctuation, or symbols. Additional characters are generated automatically.

4 Edit the default path by adding a file name to the end of the path or click **Browse,** navigate to the required folder and enter the name of the file.

5 In **Pass phrase**, enter a string for the pass phrase.

The pass phrase must:

- be between 8 and 32 characters long.

- contain at least 1 uppercase letter.

- contain at least 1 lowercase letter.

- contain at least 1 number.

- contain at least 1 non-alphanumeric character, for example, < > @ + .

The pass phrase that you enter is masked.

6 In **Confirm pass phrase**, re-enter the exact string you entered in **Pass phrase**.

Make a record of the pass phrase you entered and the security key identifier it is associated with. You need this information for later secure operations.

7 Click **Change Key.**

8 Make a record of the security key identifier and the file name from the **Change Security Key Complete** dialog, and click **OK**.

Saving a Security Key

You save an externally storable copy of the security key when the security key is first created and each time it is changed. You can create additional storable copies at any time. To save a new copy of the security key, you must provide a pass phrase. The pass phrase you choose does not need to match the pass phrase used when the security key was created or last changed. The pass phrase is applied to the particular copy of the security key you are saving.

To save the security key for the storage array:

1   In the AMW toolbar, select **Storage Array**→ **Physical Disk Security**→ **Save Security Key File**.

The **Save Security Key File** - **Enter Pass Phrase** window is displayed.

2   Edit the default path by adding a file name to the end of the path or click **Browse,** navigate to the required folder and enter the name of the file.

3   In **Pass phrase**, enter a string for the pass phrase.

The pass phrase must:

- be between 8 and 32 characters long.
- contain at least 1 uppercase letter.
- contain at least 1 lowercase letter.
- contain at least 1 number.
- contain at least 1 non-alphanumeric character, for example, $< > @ +$.

The pass phrase that you enter is masked.

4   In **Confirm pass phrase**, re-enter the exact string you entered in **Pass phrase**.

Make a record of the pass phrase you entered. You need it for later secure operations.

5   Click **Save**.

6   Make a record of the security key identifier and the file name from the **Save Security Key Complete** dialog and click **OK**.

Validate Security Key

A file in which a security key is stored is validated through the **Validate Security Key** dialog. To transfer, archive, or back up the security key, the RAID controller module firmware encrypts (or wraps) the security key and stores it in a file. You must provide a pass phrase and identify the corresponding file to decrypt the file and recover the security key.

Data can be read from a security enabled physical disk only if a RAID controller module in the storage array provides the correct security key. If security enabled physical disks are moved from one storage array to another,

the appropriate security key must also be imported to the new storage array. Otherwise, the data on the security enabled physical disks that were moved is inaccessible.

For more information on validating the security key, see the *PowerVault Modular Disk Storage Manager* online help topics.

Unlocking Secure Physical Disks

You can export a security-enabled disk group to move the associated physical disks to a different storage array. After you install those physical disks in the new storage array, you must unlock the physical disks before data can be read from or written to the physical disks. To unlock the physical disks, you must supply the security key from the original storage array. The security key on the new storage array is different and cannot unlock the physical disks.

You must supply the security key from a security key file that was saved on the original storage array. You must provide the pass phrase that was used to encrypt the security key file to extract the security key from this file.

For more information, see the *PowerVault Modular Disk Storage Manager online help* topics.

Erasing Secure Physical Disks

In the AMW, when you select a security enabled physical disk that is not part of a disk group, the **Secure Erase** menu item is enabled on the Physical Disk menu. You can use the secure erase procedure to re-provision a physical disk. You can use the **Secure Erase** option if you want to remove all of the data on the physical disk and reset the physical disk security attributes.

⚠ **CAUTION: Possible loss of data access – The Secure Erase option removes all of the data that is currently on the physical disk. This action cannot be undone.**

Before you complete this option, ensure that the physical disk that you have selected is the correct physical disk. You cannot recover any of the data that is currently on the physical disk.

After you complete the secure erase procedure, the physical disk is available for use in another disk group or in another storage array. For more information on the secure erase procedure, see the *PowerVault Modular Disk Storage Manager online help* topics.

Configuring Hot Spare Physical Disks

Guidelines to configure host spare physical disks:

- You can use only unassigned physical disks with Optimal status as hot spare physical disks.

- You can unassign only hot spare physical disks with Optimal, or Standby status. You cannot unassign a hot spare physical disk that has the In Use status. A hot spare physical disk has the In Use status when it is in the process of taking over for a failed physical disk.

- If a hot spare physical disk does not have Optimal status, follow the Recovery Guru procedures displayed by the MDSM application to correct any problem before trying to unassign the physical disk.

- Hot spare physical disks must be of the same media type and interface type as the physical disks that they are protecting.

- If there are secure disk groups and security capable disk groups in the storage array, the hot spare physical disk must match the security capability of the disk group.

- Hot spare physical disks must have capacities equal to or larger than the used capacity on the physical disks that they are protecting.

- The availability of enclosure loss protection for a disk group depends on the location of the physical disks that comprise the disk group. To ensure that enclosure loss protection is not affected, you must replace a failed physical disk to initiate the copyback process. See "Enclosure Loss Protection" on page 123.

⚠ **CAUTION: If a hot spare physical disk does not have Optimal status, follow the Recovery Guru procedures to correct the problem before you try to unassign the physical disk. You cannot assign a hot spare physical disk if it is in use (taking over for a failed physical disk).**

To assign or unassign hot spare physical disks:

1 In the AMW, select the **Physical** tab.

2 Select one or more physical disks.

3 Perform one of these actions:

- Select **Physical disk**→ **Hot Spare Coverage**.

- Right-click the physical disk and select **Hot Spare Coverage** from the pop-up menu.

  The **Hot Spare Physical Disk Options** window is displayed.

4   Select the appropriate option, you can select:

- **View/change current hot spare coverage**—to review hot spare coverage and to assign or unassign hot spare physical disks, if necessary. See step 5.

- **Automatically assign physical disks**— to create hot spare physical disks automatically for the best hot spare coverage using available physical disks.

- **Manually assign individual physical disks**—to create hot spare physical disks out of the selected physical disks on the **Physical** tab.

- **Manually unassign individual physical disks**—to unassign the selected hot spare physical disks on the **Physical** tab. See step 12.

5   To assign hot spares, in the **Hot Spare Coverage** window, select a disk group in the **Hot spare coverage** area.

6   Review the information about the hot spare coverage in the **Details** area.

7   Click **Assign**.

  The **Assign Hot Spare** window is displayed.

8   Select the relevant physical disks in the **Unassigned physical disks** area, as hot spares for the selected disk and click **OK**.

9   To unassign hot spares, in the **Hot Spare Coverage** window, select the physical disks in the **Hot spare physical disks** area.

10   Review the information about the hot spare coverage in the **Details** area.

11   Click **Unassign**.

  A message prompts you to confirm the operation.

12   Type **yes** and click **OK**.

Hot Spares and Rebuild

A valuable strategy to protect data is to assign available physical disks in the storage array as hot spares. A hot spare adds another level of fault tolerance to the storage array.

A hot spare is an idle, powered-on, stand-by physical disk ready for immediate use in case of disk failure. If a hot spare is defined in an enclosure in which a redundant virtual disk experiences a physical disk failure, a rebuild of the degraded virtual disk is automatically initiated by the RAID controller modules. If no hot spares are defined, the rebuild process is initiated by the RAID controller modules when a replacement physical disk is inserted into the storage array.

Global Hot Spares

The MD3600f Series supports global hot spares. A global hot spare can replace a failed physical disk in any virtual disk with a redundant RAID level as long as the capacity of the hot spare is equal to or larger than the size of the configured capacity on the physical disk it replaces, including its metadata.

Hot Spare Operation

When a physical disk fails, the virtual disk automatically rebuilds using an available hot spare. When a replacement physical disk is installed, data from the hot spare is copied back to the replacement physical disk. This function is called copy back. By default, the RAID controller module automatically configures the number and type of hot spares based on the number and capacity of the physical disks in your system.

A hot spare may have the following states:

- Standby hot spare—is a physical disk that is assigned as a hot spare and is available to take over for any failed physical disk.

- In-use hot spare—is a physical disk that is assigned as a hot spare and is currently replacing a failed physical disk.

Hot Spare Drive Protection

You can use a hot spare physical disk for additional data protection from physical disk failures that occur in a RAID level 1, or RAID level 5 disk group. If the hot spare physical disk is available when a physical disk fails, the RAID controller module uses redundancy data to reconstruct the data from the failed physical disk to the hot spare physical disk. When you have physically replaced the failed physical disk, a copyback operation occurs from the hot spare physical disk to the replaced physical disk.

If there are secure disk groups and security capable disk groups in the storage array, the hot spare physical disk must match the security capability of the disk group. For example, a non-security capable physical disk cannot be used as a hot spare for a secure disk group.

NOTE: For a security capable disk group, security capable hot spare physical disks are preferred. If security capable physical disks are not available, non-security capable physical disks may be used as hot spare physical disks. To ensure that the disk group is retained as security capable, the non-security capable hot spare physical disk must be replaced with a security capable physical disk.

If you select a security capable physical disk as hot spare for a non-secure disk group, a dialog box is displayed indicating that a security capable physical disk is being used as a hot spare for a non-secure disk group.

The availability of enclosure loss protection for a disk group depends on the location of the physical disks that comprise the disk group. The enclosure loss protection may be lost because of a failed physical disk and location of the hot spare physical disk. To ensure that enclosure loss protection is not affected, you must replace a failed physical disk to initiate the copyback process.

The virtual disk remains online and accessible while you are replacing the failed physical disk, because the hot spare physical disk is automatically substituted for the failed physical disk.

Enclosure Loss Protection

Enclosure loss protection is an attribute of a disk group. Enclosure loss protection guarantees accessibility to the data on the virtual disks in a disk group if a total loss of communication occurs with a single expansion enclosure. An example of total loss of communication may be loss of power to the expansion enclosure or failure of both RAID controller modules.

⚠ **CAUTION: Enclosure loss protection is not guaranteed if a physical disk has already failed in the disk group. In this situation, losing access to an expansion enclosure and consequently another physical disk in the disk group causes a double physical disk failure and loss of data.**

Enclosure loss protection is achieved when you create a disk group where all of the physical disks that comprise the disk group are located in different expansion enclosures. This distinction depends on the RAID level. If you choose to create a disk group by using the **Automatic** method, the software attempts to choose physical disks that provide enclosure loss protection. If you choose to create a disk group by using the **Manual** method, you must use the criteria specified in Table 9-2.

Table 9-2.   Criteria for Enclosure Loss Protection

RAID Level	Criteria for Enclosure Loss Protection
RAID level 5 or RAID level 6	Ensure that all the physical disks in the disk group are located in different expansion enclosures.
	Because a RAID level 5 requires a minimum of 3 physical disks, enclosure loss protection cannot be achieved if your storage array has less than 3 expansion enclosures.
	Because a RAID level 6 requires a minimum of 5 physical disks, enclosure loss protections cannot be achieved if your storage array has less than 5 expansion enclosures.

Table 9-2. Criteria for Enclosure Loss Protection *(continued)*

RAID Level	Criteria for Enclosure Loss Protection
RAID level 1	Ensure that each physical disk in a mirrored pair is located in a different expansion enclosure. This enables you to have more than two physical disks in the disk group within the same expansion enclosure.
	For example, if you are creating a six physical disk, disk group (3-mirrored pairs), you can achieve enclosure loss protection with only two expansion enclosures by specifying that the physical disk in each mirrored pair are located in separate expansion enclosures. For example:
	• Mirror pair 1—Physical disk in enclosure 1, slot 1 and physical disk in enclosure 2, slot 1.
	• Mirror pair 2—Physical disk in enclosure 1, slot 2 and physical disk in enclosure 2, slot 2.
	• Mirror pair 3—Physical disk in enclosure 1, slot 3 and physical disk in enclosure 2,slot 3.
	Because a RAID level 1 disk group requires a minimum of two physical disks, enclosure loss protections cannot be achieved if your storage array has less than two expansion enclosures.
RAID level 0	Because RAID level 0 does not have consistency, you cannot achieve enclosure loss protection.

Host-to-Virtual Disk Mapping

After you create virtual disks, you must map them to the host(s) connected to the array.

Guidelines to configure host-to-virtual disk mapping:

- Each virtual disk in the storage array can be mapped to only one host or host group.

- Host-to-virtual disk mappings are shared between controllers in the storage array.

- A unique LUN must be used by a host group or host to access a virtual disk.

- Each host has its own LUN address space. MDSM permits the same LUN to be used by different hosts or host groups to access virtual disks in a storage array.

- Not every operating system has the same number of LUNs available.
- You can define the mappings on the **Mappings** tab in the AMW. See "Using the Mappings Tab" on page 90.

Creating Host-to-Virtual Disk Mappings

Guidelines to define the mappings:

- An access virtual disk mapping is not required for an out-of-band storage array. If your storage array is managed using an out-of-band connection, and an access virtual disk mapping is assigned to the Default Group, an access virtual disk mapping is assigned to every host created from the Default Group. To prevent this action from occurring, remove the access virtual disk mapping from the Default Group.
- Most hosts have 256 LUNs mapped per storage partition. The LUN numbering is from 0 through 255. If your operating system restricts LUNs to 127, and you try to map a virtual disk to a LUN that is greater than or equal to 127, the host cannot access it.
- An initial mapping of the host group or host must be created using the Storage Partitioning Wizard before defining additional mappings. See "Storage Partitioning" on page 137.

To create host to virtual disk mappings:

1 In the AMW, select the **Mappings** tab.
2 In the **Topology** pane, select:
    - **Default Group**
    - **Undefined mappings node**
    - **Individual defined mapping**
    - **Host group**
    - **Host**
3 In the toolbar, select **Mappings→ Define→ Additional Mapping**.

   The **Define Additional Mapping** window is displayed.
4 In **Host group or host**, select the appropriate host group or host.

   All defined hosts, host groups, and the default group is displayed in the list.

5 In **Logical unit number**, select a LUN. The supported LUNs are 0 through 255.

6 Select the virtual disk to be mapped in the **Virtual Disk** area.

The **Virtual Disk** area lists the names and capacity of the virtual disks that are available for mapping based on the selected host group or selected host.

7 Click **Add**.

> ✍ **NOTE:** The **Add** button is inactive until a host group or host, LUN, and virtual disk are selected.

8 To define additional mappings, repeat step 4 through step 7.

> ✍ **NOTE:** After a virtual disk is mapped, it is no longer available in the Virtual Disk area.

9 Click **Close**.

The mappings are saved. The **Topology** pane and the **Defined Mappings** pane in the **Mappings** tab are updated to reflect the mappings.

Modifying and Removing Host-to-Virtual Disk Mapping

You can modify or remove a host-to-virtual disk mapping for several reasons, such as an incorrect mapping or reconfiguration of the storage array. Modifying or removing a host-to-virtual disk mapping applies to both hosts and host groups.

To modify or remove host to virtual disk mapping:

> ✍ **NOTE:** Before you modify or remove a host-to-virtual disk mapping, you must stop any data access (I/O) to the virtual disks to prevent data loss.

1 In the AMW, select the **Mappings** tab.

2 In the **Defined Mappings** pane, perform one of these actions:

- Select a single virtual disk, and select **Mappings→ Change→ Mapping.**

- Right-click the virtual disk, and select **Change Mapping** from the pop-up menu.

3 In **Host group or host,** select the appropriate host group or host.

By default, the drop-down list shows the current host group or the host associated with the selected virtual disk.

4 In **Logical unit number**, select the appropriate LUN.

The drop down list shows only the currently available LUNs that are associated with the selected virtual disk.

5 Click **OK**.

Stop any host applications associated with this virtual disk, and unmount the virtual disk, if applicable, from your operating system.

6 In the **Change Mapping** dialog, click **Yes** to confirm the changes.

The mapping is checked for validity and is saved. The **Defined Mappings** pane is updated to reflect the new mapping. The **Topology** pane is also updated to reflect any movement of host groups or hosts.

✍ **NOTE:** If a password is set on the storage array, the **Enter Password** dialog is displayed. Type the current password for the storage array, and click **OK**.

7 If configuring a Linux host, run the **rescan_dm_devs** utility on the host, and remount the virtual disk if required. This utility is installed on the host as part of the MDSM install process.

8 Restart the host applications.

Changing Controller Ownership of the Virtual Disk

If the host has a single data-path to the MD storage array, the virtual disk must be owned by the controller to which the host is connected. You must configure this storage array before you start I/O operations and after the virtual disk is created.

You can change the RAID controller module ownership of a standard virtual disk or a snapshot repository virtual disk. You cannot directly change the RAID controller module ownership of a snapshot virtual disk because the snapshot virtual disk inherits the RAID controller module owner of its associated source virtual disk. Changing the RAID controller module ownership of a virtual disk changes the preferred RAID controller module ownership of the virtual disk.

During a virtual disk copy, the same RAID controller module must own both the source virtual disk and the target virtual disk. Sometimes both virtual disks do not have the same preferred RAID controller module when the virtual disk copy starts. Therefore, the ownership of the target virtual disk is automatically transferred to the preferred RAID controller module of the source virtual disk. When the virtual disk copy is completed or is stopped,

ownership of the target virtual disk is restored to its preferred RAID controller module. If ownership of the source virtual disk is changed during the virtual disk copy, ownership of the target virtual disk is also changed. Under certain operating system environments, it may be necessary to reconfigure the multi-path driver before an I/O path can be used.

To change the ownership of the virtual disk to the connected controller:

1 In the AMW, select the **Logical** tab and select a virtual disk.

2 Select **Virtual Disk**→ **Change**→ **Ownership/Preferred Path**.

3 Select the appropriate RAID controller module slot and click **Yes** to confirm the selection.

Removing Host-to-Virtual Disk Mapping

1 In the AMW, select the **Mapping** tab.

2 Select a virtual disk from the **Defined Mappings** pane.

3 Perform one of these actions:

   • Select **Mappings**→ **Remove**.

   • Right-click the virtual disk, and select **Remove Mapping** from the pop-up menu.

4 Click **Yes** to remove the mapping.

Changing the RAID Controller Module Ownership of a Disk Group

You can change the RAID controller module ownership of a disk group.

You can change the RAID controller module ownership of a standard virtual disk or a snapshot repository virtual disk. You cannot directly change the RAID controller module ownership of a snapshot virtual disk because the snapshot virtual disk inherits the RAID controller module owner of its associated source virtual disk. Changing the RAID controller module ownership of a virtual disk changes the preferred RAID controller module ownership of the virtual disk.

During a virtual disk copy, the same RAID controller module must own both the source virtual disk and the target virtual disk. Sometimes both virtual disks do not have the same preferred RAID controller module when the virtual disk copy starts. Therefore, the ownership of the target virtual disk is automatically transferred to the preferred RAID controller module of the source virtual disk. When the virtual disk copy is completed or is stopped,

ownership of the target virtual disk is restored to its preferred RAID controller module. If ownership of the source virtual disk is changed during the virtual disk copy, ownership of the target virtual disk is also changed. Under certain operating system environments, it may be necessary to reconfigure the multi-path driver before an I/O path can be used.

To change the RAID controller module ownership of a disk group:

1 In the AMW, select the **Logical** tab and select a disk group.

2 Select **Disk Group**→ **Change**→ **Ownership/Preferred Path**.

3 Select the appropriate RAID controller module slot and click **Yes** to confirm the selection.

⚠ **CAUTION: Possible loss of data access– Changing ownership at the disk group level causes every virtual disk in that disk group to transfer to the other RAID controller module and use the new I/O path. If you do not want to set every virtual disk to the new path, change ownership at the virtual disk level instead.**

The ownership of the disk group is changed. I/O to the disk group is now directed through this I/O path.

✍ **NOTE:** The disk group may not use the new I/O path until the multi-path driver reconfigures and recognizes the new path. This action usually takes less than 5 minutes.

Changing the RAID Level of a Disk Group

Changing the RAID level of a disk group changes the RAID levels of every virtual disk that comprises the disk group. Performance may be slightly affected during the operation.

Guidelines to change the RAID level of a disk group:

• You cannot cancel this operation after it begins.

• The disk group must be in **Optimal** status before you can perform this operation.

• Your data remains available during this operation.

• If you do not have enough capacity in the disk group to convert to the new RAID level, an error message is displayed, and the operation does not continue. If you have unassigned physical disks, use the **Disk Group**→ **Add Free Capacity** (Physical Disks) option to add additional capacity to the disk group. Then retry the operation.

To change the RAID level of a disk group:

1 In the AMW, select the **Logical** tab and select a disk group.

2 Select **Disk Group**→ **Change**→ **RAID Level**.

3 Select the appropriate RAID level and click **Yes** to confirm the selection.

The RAID level operation begins.

Removing a Host-to-Virtual Disk Mapping Using Linux DMMP

To remove a host-to-virtual disk mapping using Linux DMMP, follow these steps:

1 Unmount the filesystem containing the virtual disk:

```
# umount filesystemDirectory
```

2 Run the following command to display multi-pathing topology:

```
# multipath -ll
```

Note the virtual disk that you want to delete from the mapping. For example, the following information may be displayed:

```
mpath6 (3600a0b80000fb6e50000000e487b02f5) dm-10
DELL, MD32xx

    [size=1.6T][features=3 queue_if_no_path
pg_init_retries 50][hwhandler=1 rdac]

    \_ round-robin 0 [prio=6][active]

        \_ 1:0:0:2 sdf 8:80 [active][ready]

     \_ round-robin 0 [prio=1][enabled]

        \_ 0:0:0:2 sde 8:64 [active][ghost]
```

In this example, the `mpath6` device contains two paths:

```
-- /dev/sdf at Host 1, Channel 0, Target 0, LUN 2
--/dev/sde at Host 0, Channel 0, Target 0, LUN 2
```

3 Flush the multi-pathing device mapping using the following command:

```
# multipath -f /dev/mapper/mapth_x
```

where `mapth_x` is the device you want to delete.

4 Delete the paths related with this device using the following command:

```
# echo 1 > /sys/block/sd_x/device/delete
```

where `sd_x` is the SD node (disk device) returned by the multipath command. Repeat this command for all paths related to this device.

For example:

```
#echo 1 > /sys/block/sdf/device/delete
```

```
#echo 1 > /sys/block/sde/device/delete
```

5 Remove mapping from MDSM, or delete the LUN if necessary.

6 If you want to map another LUN or increase volume capacity, perform this action from MDSM.

    📝 **NOTE:** If you are only testing LUN removal, you can stop at this step.

7 If a new LUN is mapped or volume capacity is changed, run the following command:

```
# rescan_dm_devs
```

8 Use the `multipath -ll` command to verify that:

   • If a new LUN is mapped, the new LUN is detected and given a multi-pathing device node

   • If you increased volume capacity, the new capacity is displayed.

Restricted Mappings

Many hosts are able to map up to 256 logical unit numbers (LUNs) (0 to 255) per storage partition. However, the maximum number of mappings differs because of operating system variables, failover driver issues, and potential data problems. The hosts listed in the Table 9-3 have these mapping restrictions.

If you try to map a virtual disk to a LUN that exceeds the restriction on these operating systems, the host is unable to access the virtual disk.

Table 9-3. Mapping Restrictions

Operating System	Highest LUN
Windows Server 2003 and Windows server 2008	255
Linux	255

Guidelines when you work with host types with LUN mapping restrictions:

- You cannot change a host adapter port to a restricted host type if there are already mappings in the storage partition that would exceed the limit imposed by the restricted host type.

- Consider the case of the Default Group that has access to LUNs up to 256 (0 to 255) and a restricted host type is added to the Default Group. In this case, the host that is associated with the restricted host type is able to access virtual disks in the Default Group with LUNs within its limits. For example, if the Default Group had two virtual disks mapped to LUNs 254 and 255, the host with the restricted host type would not be able to access those two virtual disks.

- If the Default Group has a restricted host type assigned and the storage partitions are disabled, you can map only a total of 32 LUNs. Any additional virtual disks that are created are put in the Unid**entified Mappings** area. If additional mappings are defined for one of these **Unidentified Mappings**, the **Define Additional Mapping** dialog shows the LUN list, and the **Add** button is unavailable.

- Do not configure dual mappings on a Windows host.

- If there is a host with a restricted host type that is part of a specific storage partition, all of the hosts in that storage partition are limited to the maximum number of LUNs allowed by the restricted host type.

- You cannot move a host with a restricted host type into a storage partition that already has LUNs mapped that are greater than what is allowed by the restricted host type. For example, if you have a restricted host type that allows only LUNs up to 31, you cannot move that restricted host type into a storage partition that has LUNs greater than 31 already mapped.

The Default Group on the Mappings tab has a default host type. You can change this type by selecting **Storage Array**→ **Change**→ **Default Host Type**. If you set the default host type to a host type that is restricted, the maximum number of LUNs that are allowed in the Default Group for any host are restricted to the limit imposed by the restricted host type. If a particular host with a non-restricted host type becomes part of a specific storage partition, you can change the mapping to a higher LUN.

Changing the RAID Controller Module Ownership of a Virtual Disk or a Disk Group

You can change the RAID controller module ownership of a virtual disk or a disk group.

You can change the RAID controller module ownership of a standard virtual disk or a snapshot repository virtual disk. You cannot directly change the RAID controller module ownership of a snapshot virtual disk because the snapshot virtual disk inherits the RAID controller module owner of its associated source virtual disk. Changing the RAID controller module ownership of a virtual disk changes the preferred RAID controller module ownership of the virtual disk.

During a virtual disk copy, the same RAID controller module must own both the source virtual disk and the target virtual disk. Sometimes both virtual disks do not have the same preferred RAID controller module when the virtual disk copy starts. Therefore, the ownership of the target virtual disk is automatically transferred to the preferred RAID controller module of the source virtual disk. When the virtual disk copy is completed or is stopped, ownership of the target virtual disk is restored to its preferred RAID controller module. If ownership of the source virtual disk is changed during the virtual disk copy, ownership of the target virtual disk is also changed. Under certain operating system environments, it may be necessary to reconfigure the multi-path driver before an I/O path can be used.

1 To change:

   a The RAID controller module ownership of a virtual disk—Go to step 2.

   b The RAID controller module ownership of a disk group—Go to step 3.

2 To change the RAID controller module ownership of a virtual disk, perform these steps:

   a Select the **Logical** tab.

   b Select the virtual disk.

   c Select **Virtual Disk→ Change→ Ownership/Preferred Path**. Alternatively, you can right-click the virtual disk and select **Change→ Ownership/Preferred Path** from the pop-up menu.

   d Select the RAID controller module.

⚠ **CAUTION: Possible loss of data access—If you do not use a multi-path driver, shut down any host applications that are currently using the virtual disk. This action prevents application errors when the I/O path changes.**

    e   Click **Yes**.

The ownership of the virtual disk is changed. I/O to the virtual disk is now directed through this I/O path. You are finished with this procedure.

📝 **NOTE:** The virtual disk may not use the new I/O path until the multi-path driver reconfigures and recognizes the new path. This action usually takes less than 5 minutes.

  3 To change the RAID controller module ownership of a disk group, perform these steps:

    a   Select the **Logical** tab.

    b   Select the disk group.

    c   Select **Disk Group→ Change→ Ownership/Preferred Path**. Alternatively, you can right-click the disk group and select **Change→ Ownership/Preferred Path** from the pop-up menu.

    d   Select the RAID controller module.

⚠ **CAUTION: Possible loss of data access—Changing ownership at the disk group level causes every virtual disk in that disk group to transfer to the other RAID controller module and use the new I/O path. If you do not want to set every virtual disk to the new path, change ownership at the virtual disk level instead.**

Click **Yes**.The ownership of the disk group is changed. I/O to the disk group is now directed through this I/O path.

📝 **NOTE:** The disk group may not use the new I/O path until the multi-path driver reconfigures and recognizes the new path. This action usually takes less than 5 minutes.

Changing the RAID Level of a Disk Group

Use the **Change→ RAID Level** option to change the RAID level on a selected disk group. Using this option changes the RAID levels of every virtual disk that comprises the disk group. Performance may be slightly affected during the operation.

Keep these guidelines in mind when you change the RAID level of a disk group:

- You cannot cancel this operation after it begins.
- The disk group must be in **Optimal** status before you can perform this operation.
- Your data remains available during this operation.
- If you do not have enough capacity in the disk group to convert to the new RAID level, an error message is displayed, and the operation does not continue. If you have unassigned physical disks, use the **Disk Group→ Add Free Capacity** (Physical Disks) option to add additional capacity to the disk group. Then retry the operation.

To change the RAID level of a disk group:

1 Select the **Logical** tab.

2 Select the disk group.

3 Select **Disk Group→ Change→ RAID Level**.

4 Select the RAID level (RAID level 0, RAID level 1, RAID level 5, or RAID level 6). The currently selected option is designated with a dot.

5 Click **Yes**.

The RAID level operation begins.

Storage Partitioning

A storage partition is a logical entity consisting of one or more virtual disks that can be accessed by a single host or shared among hosts that are part of a host group. The first time you map a virtual disk to a specific host or host group, a storage partition is created. Subsequent virtual disk mappings to that host or host group do not create another storage partition.

One storage partition is sufficient if:

- Only one attached host accesses all of the virtual disks in the storage array.

- All attached hosts share access to all of the virtual disks in the storage array.

When you choose this type of configuration, all of the hosts must have the same operating system and special software (such as clustering software) to manage virtual disk sharing and accessibility.

More than one storage partition is required if:

- Specific hosts must access specific virtual disks in the storage array.

- Hosts with different operating systems are attached to the same storage array. In this case, a storage partition is created for each host type.

You can use the **Storage Partitioning** wizard to define a single storage partition. The **Storage Partitioning** wizard guides you through the major steps required to specify which host groups, hosts, virtual disks, and associated logical unit numbers (LUNs) are to be included in the storage partition.

Storage partitioning fails when:

- All mappings are defined.

- You create a mapping for a host group that conflicts with an established mapping for a host in the host group.

- You create a mapping for a host in a host group that conflicts with an established mapping for the host group.

Storage partitioning is unavailable when:

- No valid host groups or hosts exist in the **Topology** pane on the **Mappings** tab.
- No host ports are defined for the host being included in the storage partition.
- All mappings are defined.

    🖉 NOTE: You can include a secondary virtual disk in a storage partition. However, any hosts that are mapped to the secondary virtual disk has read-only access until the virtual disk is promoted to a primary virtual disk, or the mirror relationship is removed.

Storage partitioning topology is the collection of elements, such as Default Group, host groups, hosts, and host ports shown as nodes in the **Topology** pane of the **Mappings** tab in the AMW. For more information, see "Using the Mappings Tab" on page 90.

If a storage partitioning topology is not defined, an informational dialog is displayed each time you select the **Mappings** tab. You must define the storage partitioning topology before you define the actual storage partition.

Disk Group and Virtual Disk Expansion

Adding free capacity to a disk group is achieved by adding unconfigured capacity on the array to the disk group. Data is accessible on disk groups, virtual disks, and physical disks throughout the entire modification operation. The additional free capacity can then be used to perform a virtual disk expansion on a standard or snapshot repository virtual disk.

Disk Group Expansion

1 In the AMW, select the **Logical** tab.

2 Select a disk group.

3 Select **Disk Group→ Add Free Capacity (Physical Disks)**.

    The **Add Free Capacity** window is displayed. Based on the RAID level, and the enclosure loss protection of the current disk group, a list of unassigned physical disks is displayed.

> ✎ **NOTE:** If the RAID level of the disk group is RAID level 5, or RAID level 6, and the expansion enclosure has enclosure loss protection, **Display only physical disks that ensures enclosure loss protection** is displayed and is selected by default.

4 in the **Available physical disks** area, select physical disks up to the allowed maximum number of physical disks.

> ✎ **NOTE:** You cannot mix different media types or different interface types within a single disk group or virtual disk.

5 Click **Add**.

A message prompts you to confirm your selection.

6 To add the capacity to the disk group, click **Yes**.

You can also use the Command Line Interface (CLI) on both Windows and Linux hosts to add free capacity to a disk group.

After the capacity expansion is completed, additional free capacity is available in the disk group for creation of new virtual disks or expansion of existing virtual disks.

Virtual Disk Expansion

Virtual disk expansion is a dynamic modification operation that increases the capacity of standard virtual disks.

> ✎ **NOTE:** Snapshot repository virtual disks can be expanded from the CLI or from MDSM. All other virtual disk types are expandable only from the CLI.

If you receive a warning that the snapshot repository virtual disk is becoming full, you may expand the snapshot repository virtual disk from MDSM. See "Snapshot Repository Capacity" on page 159 for step-by-step instructions.

Using Free Capacity

You can increase the capacity of a virtual disk using the free capacity on the disk group of the standard virtual disk or the snapshot repository virtual disk.

The **Free Capacity** node, shown in the **Logical** pane, is a contiguous region of unassigned capacity on a defined disk group. When increasing virtual disk capacity, some or all of the free capacity may be used to achieve the required final capacity. Data on the selected virtual disk remains accessible while the process for increasing virtual disk capacity is in progress.

Using Unconfigured Capacity

You can increase the capacity of a standard virtual disk or a snapshot repository virtual disk using the unconfigured capacity when no free capacity exists on a disk group. An increase is achieved by adding unconfigured capacity, in the form of unassigned physical disks to the disk group of the standard virtual disk or the snapshot repository virtual disk. See "Disk Group Expansion" on page 138.

For more information, see the *PowerVault Modular Disk Storage Manager online help* topics.

Disk Group Migration

Disk group migration allows to you export a disk group so that you can import the disk group to a different storage array. You can also export a disk group so that you can store the data offline.

✍ **NOTE:** During the export process (before the disk group is imported) you lose access to the data on the exported disk group.

✍ **NOTE:** You must export a disk group before you move the disk group or import the disk group.

Export Disk Group

The export disk group operation prepares the physical disks in the disk group for removal. You can remove the physical disks for offline storage, or you can import the disk group to a different storage array. After you complete the export disk group operation, all of the physical disks are offline. Any associated virtual disks or free capacity nodes are no longer shown in MDSM.

Non-Exportable Components

You must remove or clear any non-exportable settings before you can complete the export disk group procedure. Remove or clear the following settings:

- Persistent reservations
- Host-to-virtual disk mappings
- Virtual disk copy pairs
- Snapshot virtual disks and snapshot repository virtual disks
- Remote replication pairs
- Replication repositories

Exporting a Disk Group

On the source storage array:

1 Save the storage array configuration.

2 Stop all I/O and unmount or disconnect the file systems on the virtual disks in the disk group.

3 Back up the data on the virtual disks in the disk group.

4 Locate the disk group and label the physical disks.

5 Place the disk group offline.

6 Obtain blank physical disk modules or new physical disks.

On the target storage array:

1 Verify that the target storage array has available physical disk slots.

2 Verify that the target storage array supports the physical disks that you import.

3 Verify that the target storage array can support the new virtual disks.

4 Verify that the latest version of firmware is installed on the RAID controller module.

Import Disk Group

The import disk group operation adds the imported disk group to the target storage array. After you complete the import disk group operation, all of the physical disks have **Optimal** status. Any associated virtual disks or free capacity nodes are now shown in MDSM installed on the target storage array.

NOTE: You lose access to your data during the export/import process.

NOTE: You must export a disk group before you move the disk group or import the disk group.

Importing a Disk Group

✎ **NOTE:** You must insert all of the physical disks that are part of the disk group into the enclosure before the disk group can be imported.

On the target storage array:

1 Insert the exported physical disks into the available physical disk slots.

2 Review the Import Report for an overview of the disk group that you are importing.

3 Check for non-importable components.

4 Confirm that you want to proceed with the import procedure.

✎ **NOTE:** Some settings cannot be imported during the import disk group procedure.

The following settings are removed/cleared during the procedure:

- Persistent reservations
- Host-to-virtual disk mappings
- Virtual disk copy pairs
- Snapshot virtual disks and snapshot repository virtual disks
- Remote replication pairs
- Replication repositories

Non-Importable Components

Some components cannot be imported during the import disk group procedure. These components are removed during the procedure:

- Persistent reservations
- Host-to-virtual disk mappings
- Virtual disk copy pairs
- Snapshot virtual disks and snapshot repository virtual disks
- Remote replication pairs
- Replication repositories

Storage Array Media Scan

The media scan is a background operation that examines virtual disks to verify that data is accessible. The process finds media errors before normal read and write activity is disrupted and reports errors to the event log.

NOTE: You cannot enable background media scans on a virtual disk comprised of Solid State Disks (SSDs).

Errors discovered by the media scan include:

- Unrecovered media error—Data could not be read on the first attempt or on any subsequent attempts. For virtual disks with redundancy protection, data is reconstructed, rewritten to the physical disk, and verified and the error is reported to the event log. For virtual disks without redundancy protection (RAID level 1, RAID level 5, and RAID level 6 virtual disks), the error is not corrected but is reported to the event log.

- Recovered media error—Data could not be read by the physical disk on the first attempt but was successfully read on a subsequent attempt. Data is rewritten to the physical disk and verified and the error is reported to the event log.

- Redundancy mismatches error—The first 10 redundancy mismatches that are found on the virtual disk are reported to the event log.

- Unfixable error—Data could not be read and parity or redundancy information could not be used to regenerate the data. For example, redundancy information cannot be used to reconstruct the data on a degraded virtual disk. The error is reported to the event log.

For more information, see the *PowerVault Modular Disk Storage Manager online help* topics.

Changing Media Scan Settings

1 In the AMW, select the **Logical** tab and select any virtual disk.

2 Select **Virtual Disk**→ **Change**→ **Media Scan Settings**.

   The **Change Media Scan Settings** window is displayed.

3 Deselect **Suspend media scan** if selected.

4 In **Scan duration**, enter or select the duration (in days) for the media scan.

The media scan duration specifies the number of days for which the media scan runs on the selected virtual disks.

5 To disable media scans on an individual virtual disk, select the virtual disk in the **Select virtual disks to scan** area, and deselect **Scan selected virtual disks**.

6 To enable media scans on an individual virtual disk, select the virtual disk in the **Select virtual disks to scan** area, and select S**can selected virtual disks**.

7 To enable or disable the consistency check, select either **With consistency check** or **Without consistency check**.

   NOTE: A consistency check scans the data blocks in a RAID level 5 virtual disk, or a RAID level 6 virtual disk and checks the consistency information for each block. A consistency check compares data blocks on RAID level 1 replicated physical disks. RAID level 0 virtual disks have no data consistency.

8 Click **OK**.

Suspending the Media Scan

You cannot perform a media scan while performing another long-running operation on the disk drive such as reconstruction, copy-back, reconfiguration, virtual disk initialization, or immediate availability formatting. If you want to perform another long-running operation, you must suspend the media scan.

   NOTE: A background media scan is the lowest priority of the long-running operations.

To suspend a media scan:

1 In the AMW, select the **Logical** tab and select any virtual disk.

2 Select **Virtual Disk**→ **Change**→ **Media Scan Settings**.

   The **Change Media Scan** Settings window is displayed.

3 Select **Suspend media scan**.

   NOTE: This applies to all the virtual disks on the disk group.

4 Click **OK**.

10

Configuration: Premium Feature—Snapshot Virtual Disks

📝 **NOTE:** If you ordered this feature, you received a Premium Feature Activation card shipped in the same box as your Dell PowerVault MD storage array. Follow the directions on the card to obtain a key file and to enable the feature.

📝 **NOTE:** The snapshot feature allows up to 16 snapshots per LUN and 256 per array to be present at the same time.

A snapshot virtual disk is a point-in-time image of a virtual disk in a storage array. It is not an actual virtual disk containing a copy of the original data; it is a reference to the data that was contained on a virtual disk at a specific time. A snapshot virtual disk is the logical equivalent of a complete physical copy. However, you can create a snapshot virtual disk much faster than a physical copy, using less disk space.

The virtual disk on which the snapshot is based, called the source virtual disk, must be a standard virtual disk in your storage array. Typically, you create a snapshot so that an application, such as a backup application, can access the snapshot and read the data while the source virtual disk remains online and accessible.

📝 **NOTE:** No I/O requests are permitted on the source virtual disk while the virtual disk snapshot is being created.

A snapshot repository virtual disk containing metadata and copy-on-write data is automatically created when a snapshot virtual disk is created. The only data stored in the snapshot repository virtual disk is that which has changed since the time of the snapshot.

After the snapshot repository virtual disk is created, I/O write requests to the source virtual disk resume. Before a data block on the source virtual disk is modified, the contents of the block to be modified are copied to the snapshot repository virtual disk for safekeeping. Because the snapshot repository virtual disk stores copies of the original data in those data blocks, further changes to those data blocks write only to the source virtual disk. The snapshot

repository uses less disk space than a full physical copy, because the only data blocks that are stored in the snapshot repository virtual disk are those that have changed since the time of the snapshot.

When you create a snapshot virtual disk, specify its location, capacity, schedule, and other parameters. You can disable or delete the snapshot virtual disk when it is not required. If you disable a snapshot virtual disk, you can re-create and reuse it the next time you perform a backup. For more information, see "Re-creating Snapshot Virtual Disks" on page 164. If you delete a snapshot virtual disk, you also delete the associated snapshot repository virtual disk.

NOTE: Deleting a snapshot does not affect data on the source virtual disk.

NOTE: The following host preparation sections also apply when using the snapshot feature through the CLI interface.

Scheduling a Snapshot Virtual Disk

When you create a snapshot virtual disk, you can choose whether the snapshot is created immediately or is created according to a schedule that you determine. This schedule can be a one-time snapshot creation or an ongoing snapshot creation that occurs at regularly occurring intervals. If a schedule is not specified, the snapshot virtual disk creation happens immediately upon execution of the command.

A schedule can be specified when a snapshot virtual disk is first created, or it can be added to an existing snapshot virtual disk at any time. One schedule per snapshot virtual disk is supported.

Common Reasons for Scheduling a Snapshot Virtual Disk

Scheduling a snapshot virtual disk can serve multiple purposes across a data storage environment. Most common uses of a snapshot scheduler are:

- Data backups
- Rapid recovery from a data loss event

A scheduled data backup can protect against data loss on a regular, unmonitored basis. For example, if an application stores business-critical data on two virtual disks in the storage array, you may choose to perform an automatic backup every day. To implement this backup, select the first virtual disk and create a backup schedule that runs once a day, Monday through

Friday, at a time between the end of the work day and 11PM. Do not select an end date. Apply the same schedule to the second virtual disk, then map the two snapshot virtual disks to your backup host server and perform your regular backup procedures. Remember to unmap the two resulting snapshot virtual disks before the next scheduled snapshot begins. If the snapshot virtual disks are not unmapped, the storage array does not perform the next scheduled snapshot operation in order to avoid data corruption.

Scheduled snapshots are also valuable in the event of a data loss. For example, if you back up your data at the end of every work day and keep hourly snapshots from 8AM to 5PM, data can be recovered from the snapshots in windows smaller than one hour. To accomplish this type of rapid recovery, create a schedule that contains a start time of 8AM and an end time of 5PM, then select 10 snapshots per day on Monday through Friday with no end date.

For more information on creating snapshot virtual disk schedules, see the following sections on creating snapshots.

Guidelines for Creating Snapshot Schedules

Certain guidelines apply when creating snapshot virtual disk schedules:

- Scheduled virtual disk snapshot operations do not occur if:
  - The snapshot virtual disk is mapped
  - The storage array is offline or powered off
  - The snapshot virtual disk is in use as a source virtual disk during a Virtual Disk Copy operation
  - A copy operation is Pending or In progress
- Deleting a snapshot virtual disk that contains a schedule also deletes the schedule
- Snapshot schedules are stored in the configuration database on the storage array. The Management Station does not need to be running for scheduled snapshot operations to occur.
- Snapshot schedules can be created when the snapshot virtual disk is initially created or can be added to existing snapshot virtual disks.

Enabling and Disabling Snapshot Schedules

A scheduled snapshot operation can be temporarily suspended by disabling the schedule. When a schedule is disabled, the schedule timer continues to run but any scheduled snapshot operation do not occur.

Scheduled Snapshot Icons

Scheduled snapshots are displayed in the AMW using the following icons.

Icon	Description
	The schedule is enabled. Scheduled snapshots occurs.
	The schedule is disabled. Scheduled snapshots do not occur.

For more information on scheduling snapshots virtual disks, see the *PowerVault Modular Disk Storage Manager online help* topics and the *CLI Guide*.

Creating a Snapshot Virtual Disk Using the Simple Path

You can choose the simple path to create a snapshot virtual disk if the disk group of the source virtual disk has the required amount of free space. A snapshot repository virtual disk requires a minimum 8 MB of free capacity. The destination of a snapshot repository virtual disk is determined based on the free capacity available in the disk group.

If 8 MB of free capacity is not available in the disk group of the source virtual disk, the Create Snapshot Virtual Disks feature defaults to the advanced path. For more information, see "Creating a Snapshot Virtual Disk Using the Advanced Path" on page 152.

In the advanced path option, you can choose to place the snapshot repository virtual disk in another disk group or you can use unconfigured capacity on the storage array to create a new disk group.

About the Simple Path

Using the simple path, you can specify:

- Snapshot Virtual Disk Name—A user-specified name that helps you associate the snapshot virtual disk to its corresponding snapshot repository virtual disk and source virtual disk.

- Snapshot Repository Virtual Disk Name—A user-specified name that helps you associate the snapshot repository virtual disk to its corresponding snapshot virtual disk and source virtual disk.

- Snapshot Repository Virtual Disk Capacity—The snapshot repository virtual disk capacity is expressed as a percentage of the source virtual disk capacity (maximum 220 percent).

- Schedule—Creates the snapshot virtual disk at a specified time, or according to a regularly occurring interval. If no schedule is specified, the snapshot operation begins immediately. This parameter can also be used to apply a schedule to an existing snapshot virtual disk.

Using the simple path, the following defaults are used for the other parameters of a snapshot virtual disk:

- Capacity Allocation—The snapshot repository virtual disk is created using free capacity on the same disk group where the source virtual disk resides.

- Host-to-Virtual Disk Mapping—The default setting is **Map now**.

- Percent Full—When the snapshot repository virtual disk reaches the specified repository full percentage level, the event is logged in the Major Event Log (MEL). The default snapshot repository full percentage level is 50% of the source virtual disk.

- Snapshot Repository Virtual Disk Full Conditions—When the snapshot repository virtual disk is full, you are given a choice of failing write activity to the source virtual disk or failing the snapshot virtual disk.

Preparing Host Servers to Create the Snapshot Using the Simple Path

NOTE: Before using the Snapshot Virtual Disks Premium Feature in a Microsoft Windows clustered configuration, you must first map the snapshot virtual disk to the cluster node that owns the source virtual disk. This ensures that the cluster nodes correctly recognize the snapshot virtual disk.

NOTE: Mapping the snapshot virtual disk to the node that does not own the source virtual disk before the Snapshot enabling process is completed can result in the operating system mis-identifying the snapshot virtual disk. This, in turn, can result in data loss on the source virtual disk or an inaccessible snapshot.

NOTE: For more information on mapping the snapshot virtual disk to the secondary node, see the *Dell PowerVault MD3600f and MD3620f Storage Arrays With Microsoft Windows Server Failover Clusters* at **dell.com/support/manuals**.

NOTE: You can create concurrent snapshots of a source virtual disk on both the source disk group and on another disk group.

Before creating a Snapshot Virtual Disk:

- The following types of virtual disks are not valid source virtual disks:
  - Snapshot repository virtual disks
  - Snapshot virtual disks
  - Replication repository virtual disks
  - Target virtual disks that are participating in a virtual disk copy

  NOTE: Virtual Disk Copy is an Advanced (Premium) feature.

- You cannot create a snapshot of a virtual disk that contains unreadable sectors.

- You must satisfy the requirements of your host operating system for creating snapshot virtual disks. Failure to meet the requirements of your host operating system results in an inaccurate snapshot of the source virtual disk or the target virtual disk in a virtual disk copy.

  NOTE: Before you create a new snapshot of a source virtual disk, stop any data access (I/O) activity or suspend data transfer to the source virtual disk to ensure that you capture an accurate snapshot of the source virtual disk. Close all applications, including Windows Internet Explorer, to ensure all I/O activity has stopped.

NOTE: Removing the drive letter of the associated virtual disk(s) in Windows or unmounting the virtual drive in Linux helps to guarantee a stable copy of the drive for the Snapshot.

Before creating a snapshot virtual disk, the host server has to be in the proper state. To ensure that the host server is properly prepared to create a snapshot virtual disk, you can either use an application to carry out this task, or you can perform the following steps:

1 Stop all I/O activity to the source.

2 In the AMW, select the **Logical** tab and select a valid source virtual disk.

3 Select **Virtual Disk**→ **Snapshot**→ **Create.**

   or

   Right-click the source virtual disk and select **Create Snapshot Virtual Disk** from the pop-up menu.

   The **Create Snapshot Virtual Disk Wizard** - **Introduction** dialog is displayed.

4 Select **Simple (Recommended)** and click **Next**.

   The **Specify Snapshot Schedule** window is displayed.

5 Choose **Yes** to set up a schedule for the new snapshot virtual disk creation. To skip this option and create the snapshot immediately, choose **No**.

6 If you specified a snapshot schedule, define the schedule details in the **Create Snapshot Schedule** window and click **Next**.

7 Enter the **Snapshot virtual disk name** and the **Snapshot repository virtual disk name** and click **Next.**

   The **Specify Snapshot Repository Capacity** window is displayed.

8 Enter the snapshot repository virtual disks capacity as a percentage of the source virtual disks capacity and click **Next**.

   The **Preview** window containing the summary of the snapshot virtual disk is displayed.

9 Click **Finish**.

   The **Completed** window is displayed.

10 Click **OK**.

After creating one or more snapshot virtual disks, mount the source virtual disk, and restart the host application using that source virtual disk.

11 In the AMW, select the **Mappings** tab, assign mappings between the snapshot virtual disk and the host that accesses the snapshot virtual disk.

> ✏️ **NOTE:** In some cases, conflicts may result from mapping the same host to both a source virtual disk and its associated snapshot virtual disk. This conflict depends on the host operating system and any virtual disk manager software in use.

12 To register the snapshot virtual disk with the host operating system, run the host-based hot_add utility.

13 To associate the mapping between the storage array name and the virtual disk name, run the host-based SMdevices utility.

> ✏️ **NOTE:** If your operating system requires additional instructions, you can find those instructions in your operating system documentation.

Creating a Snapshot Virtual Disk Using the Advanced Path

About the Advanced Path

Use the advanced path to choose whether to place the snapshot repository virtual disk on free capacity or unconfigured capacity and to change the snapshot repository virtual disk parameters. You can select the advanced path regardless of whether you use free capacity or unconfigured capacity for the snapshot virtual disk.

Using the advanced path, you can specify the following parameters for your snapshot virtual disk:

- Snapshot Virtual Disk Name—A user-specified name that helps you associate the snapshot virtual disk to its corresponding snapshot repository virtual disk and source virtual disk.

- Snapshot Repository Virtual Disk Name—A user-specified name that helps you associate the snapshot repository virtual disk to its corresponding snapshot virtual disk and source virtual disk.

- Capacity Allocation—This parameter allows you to choose where to create the snapshot repository virtual disk. You can allocate capacity by using one of the following methods:
  - Use free capacity on the same disk group where the source virtual disk resides.
  - Use free capacity on another disk group.
  - Use unconfigured capacity and create a new disk group for the snapshot repository virtual disk.
  - Dell recommends placing the snapshot repository virtual disk within the disk group of the source virtual disk. This ensures that if drives associated with the disk group are moved to another storage array, all the virtual disks associated with the snapshot virtual disk remain in the same group.
- Snapshot Repository Virtual Disk Capacity—The snapshot repository virtual disk capacity is expressed as a percentage of the source virtual disk capacity (maximum 220 percent).
- Percent Full—When the snapshot repository virtual disk reaches the user-specified repository full percentage level, the event is logged in the Major Event Log (MEL). The default snapshot repository full percentage level is 50% of the source virtual disk.
- Snapshot Repository Virtual Disk Full Conditions—Choose whether to fail writes to the source virtual disk or fail the snapshot virtual disk when the snapshot repository virtual disk becomes full.
- Host-to-Virtual Disk Mapping—Choose whether to map the snapshot virtual disk to a host or host group now or to map the snapshot virtual disk later. The default setting is Map later.
- Schedule—Creates the snapshot virtual disk at a specified time, or according to a regularly occurring interval. If no schedule is specified, the snapshot operation begins immediately. This parameter can also be used to apply a schedule to an existing snapshot virtual disk.

Preparing Host Servers to Create the Snapshot Using the Advanced Path

NOTE: Before using the Snapshot Virtual Disks Premium Feature in a Microsoft Windows clustered configuration, you must first map the snapshot virtual disk to the cluster node that owns the source virtual disk. This ensures that the cluster nodes correctly recognize the snapshot virtual disk.

NOTE: Mapping the snapshot virtual disk to the node that does not own the source virtual disk before the Snapshot enabling process is completed can result in the operating system mis-identifying the snapshot virtual disk. This, in turn, can result in data loss on the source virtual disk or an inaccessible snapshot.

NOTE: For more information on mapping the snapshot virtual disk to the secondary node, see the *Dell PowerVault MD3200 and MD3220 Storage Arrays With Microsoft Windows Server Failover Clusters* at **dell.com/support/manuals**.

The destination of a snapshot repository virtual disk is determined based on the free capacity available in the disk group. A snapshot repository virtual disk requires a minimum 8 MB of free capacity. You can choose your preferred creation path—simple or advanced—if the disk group of the source virtual disk has the required amount of free space.

If 8 MB of free capacity is not available in the disk group of the source virtual disk, the **Create Snapshot Virtual Disks** feature defaults to the advanced path see "Creating a Snapshot Virtual Disk Using the Advanced Path" on page 152. In the advanced path option, you can choose to place the snapshot repository virtual disk in another disk group or you can use unconfigured capacity on the storage array to create a new disk group.

NOTE: You can create concurrent snapshots of a source virtual disk on both the source disk group and on another disk group.

Remember the following guidelines before creating a Snapshot Virtual Disk:

- The following types of virtual disks are not valid source virtual disks:
  - Snapshot repository virtual disks.
  - Snapshot virtual disks.
  - Replication repository virtual disks.
  - Target virtual disks that are participating in a virtual disk copy.

  NOTE: Virtual Disk Copy is an Advanced (Premium) feature.

- You cannot create a snapshot of a virtual disk that contains unreadable sectors.

- You must satisfy the requirements of your host operating system for creating snapshot virtual disks. Failure to meet the requirements of your host operating system results in an inaccurate snapshot of the source virtual disk or the target virtual disk in a virtual disk copy.

📝 **NOTE:** Before you create a new snapshot of a source virtual disk, stop any data access (I/O) activity or suspend data transfer to the source virtual disk to ensure that you capture an accurate snapshot of the source virtual disk. Close all applications, including Windows Internet Explorer, to ensure all I/O activity has stopped.

📝 **NOTE:** Removing the drive letter of the associated virtual disk(s) in Windows or unmounting the virtual drive in Linux helps to guarantee a stable copy of the drive for the Snapshot.

Before creating a snapshot virtual disk, the host server must be in the proper state. To prepare your host server:

1 Stop all I/O activity to the source.

2 Using your Windows system, flush the cache to the source. At the host prompt, type `SMrepassist -f <filename-identifier>` and press <Enter>. For more information, see "SMrepassist Utility" on page 289.

3 Remove the drive letter(s) of the source in Windows or unmount the virtual drive(s) in Linux to help guarantee a stable copy of the drive for the Snapshot. If this is not done, the snapshot operation reports that it has completed successfully, but the snapshot data is not updated properly.

📝 **NOTE:** Verify that the virtual disk has a status of Optimal or Disabled by clicking the **Summary** tab and then clicking **Disk Groups & Virtual Disks**.

4 Follow any additional instructions for your operating system. Failure to follow these additional instructions can create unusable snapshot virtual disks.

📝 **NOTE:** If your operating system requires additional instructions, you can find those instructions in your operating system documentation.

After your host server is prepared, see "Creating a Snapshot Virtual Disk Using the Advanced Path" on page 152 to create the snapshot using the advanced path.

If you want to use a snapshot regularly, such as for backups, use the **Disable Snapshot** and **Re-create Snapshot** options to reuse the snapshot. Disabling and re-creating snapshots preserves the existing virtual disk-to-host mappings to the snapshot virtual disk.

Creating the Snapshot Using the Advanced Path

NOTE: Removing the drive letter of the associated virtual disk in Windows or unmounting the virtual drive in Linux helps to guarantee a stable copy of the drive for the Snapshot.

After first preparing the host server(s) as specified in the preceding procedure, complete the following steps to create a virtual disk snapshot using the advanced path:

1 Stop the host application accessing the source virtual disk, and unmount the source virtual disk.

2 In the AMW, select the **Logical** tab, select a valid source virtual disk.

3 Select **Virtual Disk**→ **Snapshot**→ **Create**.

    or

4 Right-click the source virtual disk and select **Create Snapshot Virtual Disk** from the pop-up menu.

    The **Create Snapshot Virtual Disk Wizard** - **Introduction** dialog is displayed.

5 Select **Advanced**, and click **Next**.

    The **Specify Names** window is displayed.

6 Enter the **Snapshot visual disk name** and the **Snapshot repository virtual disk name** and click **Next.**

    The **Allocate Capacity** window is displayed.

7 In the **Capacity allocation** area, select:

    • Free capacity on same disk group as base (recommended)

    • Free capacity on different disk group

    • Unconfigured capacity (create new disk group)

8 Enter the snapshot repository virtual disks capacity as a percentage of the source virtual disks capacity and click **Next**.

The **Specify Virtual Disk Parameters** window is displayed.

9   In the **Snapshot virtual disk parameters** area, select the relevant mapping option, you can select:

   •   **Automatic**

   •   **Map later**

10 In the **Snapshot repository virtual disk parameters** area, enter the system behavior when:

   •   The snapshot repository virtual disk is full to the selected percentage level.

   •   The snapshot repository virtual disk is full.

11 Click **Next**.

   The **Preview** window containing the summary of the snapshot virtual disk is displayed.

12 Click **Finish**.

   The **Completed** window is displayed.

13 Click **OK**.

14 In the **Mappings** tab, assign mappings between the snapshot virtual disk and the host that accesses the snapshot virtual disk.

15 To register the snapshot virtual disk with the host operating system, run the host-based hot_add utility.

16 To associate the mapping between the storage array name and the virtual disk name, run the host-based SMdevices utility.

Specifying Snapshot Virtual Disk Names

Choose a name that helps you associate the snapshot virtual disk and snapshot repository virtual disk with its corresponding source virtual disk. The following information is useful when naming virtual disks:

By default, the snapshot name is shown in the **Snapshot virtual disk name** field as:

```
<source-virtual disk-name>-<sequence-number>
```

where `sequence-number` is the chronological number of the snapshot relative to the source virtual disk.

The default name for the associated snapshot repository virtual disk that is shown in the **Snapshot repository virtual disk** field is:

`<source-virtual disk-name>–R<sequence-number>`

For example, if you are creating the first snapshot virtual disk for a source virtual disk called Accounting, the default snapshot virtual disk is Accounting-1, and the associated snapshot repository virtual disk default name is Accounting-R1. The default name of the next snapshot virtual disk you create based on Accounting is Accounting-2, with the corresponding snapshot repository virtual disk named as Accounting-R2 by default.

- Whether you use the software-supplied sequence number that (by default) populates the Snapshot virtual disk name or the Snapshot repository virtual disk name field, the next default name for a snapshot or snapshot repository virtual disk still uses the sequence number determined by the software. For example, if you give the first snapshot of source virtual disk Accounting the name Accounting-8, and do not use the software-supplied sequence number of 1, the default name for the next snapshot of Accounting is still Accounting-2.

- The next available sequence number is based on the number of existing snapshots of a source virtual disk. If you delete a snapshot virtual disk, its sequence number becomes available again.

- You must choose a unique name for the snapshot virtual disk and the snapshot repository virtual disks, or an error message is displayed.

- Names are limited to 30 characters. After you reach this limit in either the snapshot virtual disk name or the Snapshot repository virtual disk name fields, you can no longer type in the field. If the source virtual disk is 30 characters, the default names for the snapshot and its associated snapshot repository virtual disk use the source virtual disk name truncated enough to add the sequence string. For example, for Host Software Engineering Group GR-1, the default snapshot name is Host Software Engineering GR-1, and the default repository name would be Host Software Engineering GR-R1.

Snapshot Repository Capacity

If you receive a warning that the capacity for the snapshot repository virtual disk is approaching its threshold, you can increase the capacity of a snapshot repository virtual disk by using one of the following methods:

- Use the free capacity available on the disk group of the snapshot repository virtual disk.
- Add unconfigured capacity to the disk group of the snapshot repository virtual disk. Use this option when no free capacity exists on the disk group.

You cannot increase the storage capacity of a snapshot repository virtual disk if the snapshot repository virtual disk has any one of the following conditions:

- The virtual disk has one or more hot spare drives in use.
- The virtual disk has a status other than Optimal.
- Any virtual disk in the disk group is in any state of modification.
- The controller that has ownership of this virtual disk is currently adding capacity to another virtual disk. Each controller can add capacity to only one virtual disk at a time.
- No free capacity exists in the disk group.
- No unconfigured capacity is available to add to the disk group.

✍ **NOTE:** You can add a maximum of two physical disks at one time to increase snapshot repository virtual disk capacity.

To expand the snapshot repository virtual disk:

1 In the AMW, select the **Logical** tab.

2 Select the snapshot repository virtual disk for which you want to increase the capacity.

3 Select **Virtual Disk→ Increase Capacity**.

   ✍ **NOTE:** If no free capacity or unconfigured capacity is available, the Increase Capacity option is disabled.

   The **Increase Snapshot Repository Capacity** window displays the Virtual disk attributes. The snapshot repository virtual disk name, the associated snapshot virtual disk name, the associated source virtual disk capacity and name, the current capacity, and the amount of free capacity that is available for the selected snapshot repository virtual disk is displayed. If

free capacity is available, the maximum free space is displayed in the **Increase capacity by** field.

If free capacity does not exist on the disk group, the free space that is displayed in the Increase capacity by spinner box is 0. You must add physical disks to create free capacity on the disk group.

4 To increase capacity of the snapshot repository virtual disk, use one of these methods:

   • Use the free capacity on the disk group of the snapshot repository virtual disk – Go to step 5.

   • Add unconfigured capacity, or physical disks to the disk group of the snapshot repository virtual disk – Go to step 7.

5 In **Increase capacity by**, enter or select the appropriate capacity.

6 Click **OK**.

   The **Logical** tab is updated. The snapshot repository virtual disk having its capacity increased shows a status of **Operation in Progress**. In addition, the snapshot repository virtual disk shows its original capacity and the total capacity being added. The virtual disk involved shows a reduction in capacity. If all of the free capacity is used to increase the size of the virtual disk, the **Free Capacity** node involved is removed from the **Logical** tab.

7 If unassigned physical disks are not available, do you have empty slots in the expansion enclosures?

   • Yes, there are empty slots—Insert new physical disks by using the information in the initial setup guide for your expansion enclosure. Go to step 9.

   • No, there are no empty slots—Install another expansion enclosure and additional physical disks. Use the information in the initial setup guides for your RAID controller module and your expansion enclosure. Go to step 9.

   ✍ **NOTE:** The physical disks that you add must be of the same media type and interface type as the physical disks that already make up the disk group of the snapshot repository virtual disk.

8 Click **Add Physical Disks**.

💋 **NOTE:** The physical disks that is displayed has a capacity that is either the same size or larger than the capacity of the physical disks already being used by the disk group.

9   Select either a single physical disk to add or two physical disks to add.

10 Click **Add**.

The **Add Physical Disks** window closes.

11 Check the **Physical Disks to add** [enclosure, slot] area to ensure that the correct physical disks are added.

12 Either accept the final capacity, or enter or select the appropriate capacity in **Increase capacity by** field.

13 Click **OK**.

The **Logical** tab is updated. The snapshot repository virtual disk that is having its capacity increased shows a status of **Operation in Progress**. In addition, the snapshot repository virtual disk shows its original capacity and the total capacity being added. The **Free Capacity** node involved in the increase shows a reduction in capacity. If all of the free capacity is used to increase the size of the virtual disk, the **Free Capacity** node involved is removed from the **Logical** tab.

A new **Free Capacity** node is created and shown in the **Logical** tab if these conditions exist:

•   A **Free Capacity** node did not exist prior to the addition of capacity.

•   Not all of the capacity that is added is used to increase the capacity of the snapshot repository virtual disk.

On the **Physical** tab, the unassigned physical disks or unconfigured capacity that you added to increase the capacity of the snapshot repository virtual disk change to assigned physical disks. The new assigned physical disks are associated with the disk group of the snapshot repository virtual disk.

Disabling a Snapshot Virtual Disk

Disable a snapshot virtual disk if:

- You do not need the snapshot now.
- You intend to re-create the snapshot at a later time and want to retain the associated snapshot repository virtual disk so that you do not need to create it again.
- You want to maximize storage array performance by stopping copy-on-write activity to the snapshot repository virtual disk.

✏ **NOTE:** If you do not intend to re-create the snapshot virtual disk at a later time, in the **Logical** pane, select the snapshot virtual disk, and select **Virtual Disk**→ **Delete** to remove it. The associated snapshot repository virtual disk is also removed. For more information on removing a snapshot virtual disk, see the *PowerVault Modular Disk Storage Manager online help* topics.

✏ **NOTE:** The SMdevices utility displays the snapshot virtual disk in its output, even after the snapshot virtual disk is disabled.

To disable a snapshot virtual disk:

1 In the AMW, select the **Logical** tab, select the snapshot virtual disk, and select **Virtual Disk**→ **Snapshot**→ **Disable**.

2 In the text box, type **yes** and click **OK**.

   The snapshot virtual disk is disabled. The associated snapshot repository virtual disk does not change status. The copy-on-write activity to the snapshot repository virtual disk stops until the snapshot virtual disk is re-created.

Preparing Host Servers to Re-create a Snapshot Virtual Disk

NOTE: Before you create a new snapshot of a source virtual disk, stop any data access (I/O) activity or suspend data transfer to the source virtual disk and snapshot virtual disk to ensure that you capture an accurate snapshot of the source virtual disk. Close all applications, including Windows Internet Explorer, to ensure all I/O activity has stopped.

NOTE: Removing the drive letter of the associated virtual disk in Windows or unmounting the virtual drive in Linux helps to guarantee a stable copy of the drive for the Snapshot.

Before re-creating a snapshot virtual disk, both the host server and the associated virtual disk you are re-creating have to be in the proper state.

To prepare your host server and virtual disk:

1   Stop all I/O activity to the source and snapshot virtual disk (if mounted).

2   Using your Windows system, flush the cache to both the source and the snapshot virtual disk (if mounted). At the host prompt, type

```
SMrepassist -f <filename-identifier>
```

and press <Enter>. For more information, see "SMrepassist Utility" on page 289.

3   Click the **Summary** tab, then click **Disk Groups & Virtual Disks** to ensure that the snapshot virtual disk is in **Optimal** or **Disabled** status.

4   Remove the drive letter(s) of the source and (if mounted) snapshot virtual disk in Windows or unmount the virtual drive(s) in Linux to help guarantee a stable copy of the drive for the Snapshot. If this is not done, the snapshot operation reports that it has completed successfully, but the snapshot data is not updated properly.

5   Follow any additional instructions for your operating system. Failure to follow these additional instructions can create unusable snapshot virtual disks.

NOTE: If your operating system requires additional instructions, you can find those instructions in your operating system documentation.

Re-creating Snapshot Virtual Disks

You can re-create a snapshot virtual disk that you have previously disabled.

⚠ **CAUTION: Possible loss of data redundancy – If the snapshot virtual disk is in Optimal status, it is first disabled prior to being re-created. This action invalidates the current snapshot.**

Keep these important guidelines in mind when you re-create a snapshot virtual disk:

- To re-create the snapshot virtual disks correctly, follow the instructions for your operating system.

  📝 **NOTE:** Failing to follow these additional instructions could create unusable snapshot virtual disks. For more information, see the *PowerVault Modular Disk Storage Manager online help* topics.

- To use this option, the snapshot virtual disk must be either in an **Optimal** status or **Disabled** status.

- When using this option, the previously configured snapshot name parameters and snapshot repository virtual disk are used.

To recreate the snapshot virtual disk:

1 In the AMW, select the **Logical** tab, select a snapshot virtual disk.

2 Select **Virtual Disk→ Snapshot→ Re-create**.

3 Type **yes**, and click **OK**.

Snapshot Rollback

The snapshot rollback feature allows you to revert the contents of a virtual disk to match a point-in-time image existing on a snapshot virtual disk.

During a rollback, the host server can still write data to the base virtual disk. However, the snapshot virtual disk is set as read-only during the rollback operation and becomes available for write operations immediately after the rollback is complete. The snapshot virtual disk cannot be restarted, deleted, or disabled during the rollback operation.

The associated snapshot repository virtual disk must have sufficient capacity to process the rollback operation and the write operations from the host. At most, the snapshot repository virtual disk may need twice the size of the base disk, plus additional metadata space equaling approximately 1/1000th (that is, 0.1%) of the Base volume capacity.

NOTE: Due to host server write operations, the content in the snapshot virtual disk may have changed since creation of the snapshot. The rollback operation includes any changes made to the snapshot after it was created.

You can set priority for a rollback operation. Higher priority allocates more system resources for the rollback operation and affects overall system performance.

Rules and Guidelines for Performing a Snapshot Rollback

The following rules and guidelines apply to performing a snapshot rollback:

- Rolling back a base virtual disk to a snapshot virtual disk does not affect the contents of the snapshot virtual disks.
- Only one snapshot rollback operation can be performed at a time.
- While a base virtual disk that is undergoing a rollback, you cannot create a new snapshot virtual disks from that base virtual disk.
- A snapshot rollback cannot be started while any of these operations are underway:
  - Virtual Disk Capacity Expansion
  - Virtual Disk Expansion (VDE)
  - RAID Level Migration
  - Segment Size Migration
  - Virtual Disk Copy
  - Role Reversal (in remote replication)
- If the base virtual disk is a secondary virtual disk in a remote replication, you cannot perform a snapshot rollback.
- If any capacity used in the associated snapshot repository virtual disk contains unreadable sectors, the snapshot rollback fails.

If an error occurs during the rollback, the operation is paused and the base virtual disk and snapshot virtual disk displays Needs Attention icons. The RAID controller module also logs the event to the Major Event Log (MEL). Follow the Recovery Guru procedure to correct the error and repeat the rollback operation.

⚠ **WARNING: Risk of data loss: If you cancel a snapshot rollback in progress, the base virtual disk may remain in an unusable state and the snapshot virtual disk is displayed as failed in the MD storage management software. Therefore, do not cancel a snapshot rollback unless reliable recovery options exist for restoring the content of the base virtual disk.**

Command Line Options

Optionally, you also can use the command line interface (CLI) to start, cancel, resume or modify priority of a snapshot rollback. For more information, see the *CLI Guide*.

Protecting Against a Failed Snapshot Rollback

To protect your base virtual disk data, it is recommended that you create a new snapshot virtual disk from the base virtual disk before beginning a rollback operation. If the snapshot rollback fails, use this new snapshot virtual disk to restore your base virtual disk.

Previous Versions of the MD Storage Manager

Snapshot virtual disks created using previous versions of MD Storage Manager that did not support snapshot rollback do not have to be recreated or changed to be used in a subsequent rollback operation. Once the latest versions of the MD storage management software and RAID controller module firmware are installed, snapshot virtual disks created under previous versions support the snapshot rollback feature. However, if you revert to an older version of the MD storage management software after you have performed a snapshot rollback, the older MD storage management software does not support the snapshot virtual disk.

Starting a Snapshot Rollback

To start a snapshot rollback:

1. In the array management window, select the **Logical** tab.

2 Choose one:

- Select the snapshot virtual disk, and select **Virtual Disk**→ **Snapshot**→ **Rollback**.

- Right-click the snapshot virtual disk and select **Rollback**.

The **Confirm Rollback Snapshot Virtual Disk** dialog is displayed.

3 In the Select rollback priority area, use the slider bar to set rollback priority.

> ✍ **NOTE:** If priority is set at the lowest rate, normal data write activity is highest priority and the rollback operation takes longer to complete. If the priority is at the highest rate, the rollback operation is highest priority and data write activity is reduced.

4 To start the snapshot rollback, type `yes` in the confirmation box and click **OK**.

Rollback status is shown in the **Properties** pane for the base virtual disk and snapshot virtual disk.

Resuming a Snapshot Rollback

If an error occurs during the snapshot rollback and the operation is paused, you can resume the rollback using the following steps:

1 In the array management window, select the **Logical** tab.

2 Choose one:

- Select the snapshot virtual disk, and select **Virtual Disk**→ **Snapshot**→ **Resume Rollback**.

- Right-click the snapshot virtual disk and select **Resume Rollback**.

The **Resume Rollback** dialog is displayed.

3 Click **OK**.

If the snapshot rollback resumed successfully, status is displayed in Properties pane of the base virtual disk or snapshot virtual disk.

If the snapshot rollback did not resume successfully, the rollback operation pauses again and both virtual disks display **Needs Attention** icons. Check the Major Event Log (MEL) for details and follow the Recovery Guru procedure to correct the issue.

Canceling a Snapshot Rollback

⚠ **WARNING: Risk of data loss: If you cancel a snapshot rollback in progress, the base virtual disk may remain in an unusable state and the snapshot virtual disk is displayed as failed in the MD storage management software. Therefore, do not cancel a snapshot rollback unless reliable recovery options exist for restoring the content of the base virtual disk.**

1 In the array management window, select the **Logical** tab.

2 Choose one:

   • Select the snapshot virtual disk, and select **Virtual Disk→ Snapshot→ Cancel Rollback**.

   • Right-click the snapshot virtual disk and select **Cancel Rollback**.

   The **Confirm Cancel Rollback** dialog is displayed.

3 To cancel the snapshot rollback, type yes in the confirmation box and click **OK**.

4 Click **Yes** to cancel the rollback operation.

12

Configuration: Premium Feature— Upgrading to High–Performance Tier

The High Performance Tier premium feature on an MD3600f Series array increases the performance of the system beyond that of an MD3600f Series array operating at the standard performance level.

If this feature is ordered, a Premium Feature Activation card is placed in the box with the storage array. After reading the information below, follow the instructions on the card to obtain a key file and enable the feature.

⚠ **CAUTION: Loss of data access – The storage array automatically restarts when the High-Performance-Tier feature is enabled or disabled. During the restart, data is unavailable. Data availability is restored when the array restarts.**

To upgrade from a standard-performance-tier storage array to a high-performance-tier storage array, you enable the high-performance-tier premium feature, using the Dell PowerVault Modular Disk Storage Management (MDSM) software.

When the high performance tier feature is enabled or disabled the array restarts. During this time data access and management access to the controller is temporarily lost.

It is recommended that all I/O to the array be stopped before this feature is enabled or disabled.

While the array is restarting, the state of the array in the MDSM application may change from **Optimal** to **Unresponsive**. When the restart completes the status returns to **Optimal**.

When the array status returns to **Optimal**, I/O to the array can be restarted.

11

Configuration: Premium Feature— Virtual Disk Copy

NOTE: A virtual disk copy overwrites data on the target virtual disk. Before starting a virtual disk copy, ensure that you no longer need the data or back up the data on the target virtual disk.

NOTE: If you ordered this feature, you received a Premium Feature Activation card that shipped in the same box as your Dell PowerVault MD storage array. Follow the directions on the card to obtain a key file and to enable the feature.

NOTE: The preferred method for creating a virtual disk copy is to copy from a snapshot virtual disk. This allows the original virtual disk used in the snapshot operation to remain fully available for read/write activity while the snapshot is used as the source for the virtual disk copy operation.

When you create a virtual disk copy, you create a copy pair that has a source virtual disk and a target virtual disk on the same storage array.

The source virtual disk is the virtual disk that contains the data you want to copy. The source virtual disk accepts the host I/O read activity and stores the data until it is copied to the target virtual disk. The source virtual disk can be a standard virtual disk, a snapshot virtual disk, or the source virtual disk of a snapshot virtual disk. When you start a virtual disk copy, all data is copied to the target virtual disk, and the source virtual disk permissions are set to read-only until the virtual disk copy is complete.

The target virtual disk is a virtual disk to which you copy data from the source virtual disk. The target virtual disk can be a standard virtual disk, or the source virtual disk of a failed or disabled snapshot virtual disk.

After the virtual disk copy is complete, the source virtual disk becomes available to host applications for write requests. To prevent error messages, do not attempt to access a source virtual disk that is participating in a virtual disk copy while the virtual disk copy is in progress.

Reasons to use virtual disk copy include:

- Copying data for improved access—As your storage requirements for a virtual disk change, you can use a virtual disk copy to copy data to a virtual disk in a disk group that uses drives with larger capacity within the same storage array. Copying data for larger access capacity enables you to move data to greater capacity physical disks (for example, 61 GB to 146 GB).

- Restoring snapshot virtual disk data to the source virtual disk—The Virtual Disk Copy feature enables you first to restore the data from a snapshot virtual disk and then to copy the data from the snapshot virtual disk to the original source virtual disk.

- Creating a backup copy—The **Virtual Disk Copy** feature enables you to create a backup of a virtual disk by copying data from one virtual disk (the source virtual disk) to another virtual disk (the target virtual disk) in the same storage array, minimizing the time that the source virtual disk is unavailable to host write activity. You can then use the target virtual disk as a backup for the source virtual disk, as a resource for system testing, or to copy data to another device, such as a tape drive or other media.

✎ **NOTE:** Recovering from a backup copy — You can use the **Edit Host-to-Virtual Disk Mappings** feature to recover data from the backup virtual disk you created in the previous procedure. The **Mappings** option enables you to unmap the source virtual disk from its host and then to map the backup virtual disk to the same host.

Types of Virtual Disk Copies

You can perform either offline or online virtual disk copies. To ensure data integrity, all I/O to the target virtual disk is suspended during either type of virtual disk copy operation. After the virtual disk copy is complete, the target virtual disk automatically becomes read-only to the hosts.

Offline Copy

An offline copy reads data from the source virtual disk and copies it to a target virtual disk, while suspending all updates to the source virtual disk when the copy is in progress. In an offline virtual disk copy, the relationship is between a source virtual disk and a target virtual disk. Source virtual disks that are participating in an offline copy are available for read requests, while the virtual disk copy displays the **In Progress** or **Pending** status. Write requests are allowed only after the offline copy is complete. If the source virtual disk is

formatted with a journaling file system, any attempt to issue a read request to the source virtual disk may be rejected by the storage array RAID controller modules and result in an error message. Make sure that the Read-Only attribute for the target virtual disk is disabled after the virtual disk copy is complete to prevent error messages from being displayed.

Online Copy

An online copy creates a point-in-time snapshot copy of any virtual disk within a storage array, while still allowing writes to the virtual disk when the copy is in progress. This is achieved by creating a snapshot of the virtual disk and using that snapshot as the actual source virtual disk for the copy. In an online virtual disk copy, the relationship is between a snapshot virtual disk and a target virtual disk. The virtual disk for which the point-in-time image is created (the source virtual disk) must be a standard virtual disk in the storage array.

A snapshot virtual disk and a snapshot repository virtual disk are created during the online copy operation. The snapshot virtual disk is not an actual virtual disk containing data; instead, it is a reference to the data contained on the virtual disk at a specific time. For each snapshot taken, a snapshot repository virtual disk is created to hold the copy-on-write data for the snapshot. The snapshot repository virtual disk is used only to manage the snapshot image.

Before a data block on the source virtual disk is modified, the contents of the block to be modified are copied to the snapshot repository virtual disk. Because the snapshot repository virtual disk stores copies of the original data in those data blocks, further changes to those data blocks write only to the source virtual disk.

NOTE: If the snapshot virtual disk that is used as the copy source is active, the source virtual disk performance degrades due to copy-on-write operations. When the copy is complete, the snapshot is disabled and the source virtual disk performance is restored. Although the snapshot is disabled, the repository infrastructure and copy relationship remain intact.

Creating a Virtual Disk Copy for an MSCS Shared Disk

To create a virtual disk copy for a Microsoft Cluster Server (MSCS) shared disk, create a snapshot of the virtual disk, and then use the snapshot virtual disk as the source for the virtual disk copy.

NOTE: An attempt to directly create a virtual disk copy for an MSCS shared disk, rather than using a snapshot virtual disk, fails with the following error:

```
The operation cannot complete because the selected
virtual disk is not a source virtual disk
candidate.
```

NOTE: When creating a snapshot virtual disk, map the snapshot virtual disk to only one node in the cluster. Mapping the snapshot virtual disk to the host group or both nodes in the cluster may cause data corruption by allowing both nodes to concurrently access data.

Virtual Disk Read/Write Permissions

After the virtual disk copy is complete, the target virtual disk automatically becomes read-only to the hosts. The target virtual disk rejects read and write requests while the virtual disk copy operation has a status of **Pending** or **In Progress** or if the operation fails before completing the copy. Keep the target virtual disk Read-Only enabled if you want to preserve the data on the target virtual disk for reasons such as the following:

- If you are using the target virtual disk for backup purposes.
- If you are using the data on the target virtual disk to copy back to the source virtual disk of a disabled or failed snapshot virtual disk.

If you decide not to preserve the data on the target virtual disk after the virtual disk copy is complete, change the write protection setting for the target virtual disk to Read/Write.

To set the target virtual disk read/write permissions:

1 In the AMW, select **Virtual Disk**→ **Copy Manager**.

   The **Copy Manager** window is displayed.

2 Select one or more copy pairs in the table.

3 Perform one of these actions:

- To enable Read-Only permission, select **Change**→ **Target Virtual Disk Permissions**→ **Enable Read-Only**.

  📝 **NOTE:** Write requests to the target virtual disk are rejected when the Read-Only permission is enabled on the target virtual disk.

- To disable Read-Only permission, select **Change**→ **Target Virtual Disk Permissions**→ **Disable Read-Only**.

Virtual Disk Copy Restrictions

Before you perform any virtual disk copy tasks, understand and adhere to the restrictions listed in this section. The restrictions apply to the source virtual disk, the target virtual disk, and the storage array.

- While a virtual disk copy has a status of **In Progress**, **Pending**, or **Failed**, the source virtual disk is available for read I/O activity only. After the virtual disk copy is complete, read and write I/O activity to the source virtual disk are permitted.

- A virtual disk can be selected as a target virtual disk for only one virtual disk copy at a time.

- A virtual disk copy for any virtual disk cannot be mounted on the same host as the source virtual disk.

- Windows does not allow a drive letter to be assigned to a virtual disk copy.

- A virtual disk with a **Failed** status cannot be used as a source virtual disk or target virtual disk.

- A virtual disk with a **Degraded** status cannot be used as a target virtual disk.

- A virtual disk participating in a modification operation cannot be selected as a source virtual disk or target virtual disk. Modification operations include the following:
  - Capacity expansion
  - RAID-level migration
  - Segment sizing
  - Virtual disk expansion
  - Defragmenting a virtual disk

📓 **NOTE:** The following host preparation sections also apply when using the virtual disk copy feature through the CLI interface.

Creating a Virtual Disk Copy

Source virtual disks that are participating in a virtual disk copy are available for read I/O activity only while a virtual disk copy has a status of In Progress or Pending. Write requests are allowed after the virtual disk copy has completed. If the source virtual disk is formatted with a journaling file system, any attempt to issue a read request to the source virtual disk may be rejected by the storage array, and an error message may be displayed. The journaling file system driver issues a write request before it attempts to issue the read request. The storage array rejects the write request, and the read request may not be issued due to the rejected write request. This condition may result in an error message being displayed, which indicates that the source virtual disk is write protected.

To prevent this issue from occurring, do not attempt to access a source virtual disk that is participating in a virtual disk copy while the virtual disk copy has a status of In Progress. Also, ensure that the Read-Only attribute for the target virtual disk is disabled after the virtual disk copy has completed to prevent error messages from being displayed.

The Virtual Disk Copy premium feature includes:

- The **Create Copy Wizard**, which assists in creating a virtual disk copy
- The **Copy Manager**, which monitors virtual disk copies after they are created

Before You Begin

A virtual disk copy fails all snapshot virtual disks that are associated with the target virtual disk, if any exist. If you select a source virtual disk of a snapshot virtual disk, you must disable all of the snapshot virtual disks that are associated with the source virtual disk before you can select it as a target virtual disk. Otherwise, the source virtual disk cannot be used as a target virtual disk.

A virtual disk copy overwrites data on the target virtual disk and automatically makes the target virtual disk read-only to hosts.

If 16 virtual disk copies with the status of **In Progress** exist, any subsequent virtual disk copy has the status **Pending**, which stays until one of the 16 virtual disk copies complete.

Virtual Disk Copy and Modification Operations

If a modification operation is running on a source virtual disk or a target virtual disk, and the virtual disk copy has a status of **In Progress**, **Pending**, or **Failed**, the virtual disk copy does not take place. If a modification operation is running on a source virtual disk or a target virtual disk after a virtual disk copy is created, the modification operation must complete before the virtual disk copy can start. If a virtual disk copy has a status of **In Progress**, any modification operation does not take place.

Create Copy Wizard

The **Create Copy Wizard** guides you through:

- Selecting a source virtual disk from a list of available virtual disks.
- Selecting a target virtual disk from a list of available virtual disks.
- Setting the copy priority for the virtual disk copy.

When you have completed the wizard dialogs, the virtual disk copy starts, and data is read from the source virtual disk and written to the target virtual disk.

The **Operation in Progress** icon is displayed on the source virtual disk and the target virtual disk while the virtual disk copy has a status of **In Progress** or **Pending**.

Failed Virtual Disk Copy

A virtual disk copy can fail due to:

- A read error from the source virtual disk.
- A write error to the target virtual disk.
- A failure in the storage array that affects the source virtual disk or the target virtual disk.

When the virtual disk copy fails, a critical event is logged in the **Event Log**, and a **Needs Attention** icon is displayed in the AMW. While a virtual disk copy has this status, the host has read-only access to the source virtual disk. Read requests from and write requests to the target virtual disk do not take place until the failure is corrected by using the Recovery Guru.

Preferred RAID Controller Module Ownership

During a virtual disk copy, the same RAID controller module must own both the source virtual disk and the target virtual disk. If both virtual disks do not have the same preferred RAID controller module when the virtual disk copy starts, the ownership of the target virtual disk is automatically transferred to the preferred RAID controller module of the source virtual disk. When the virtual disk copy is completed or is stopped, ownership of the target virtual disk is restored to its preferred RAID controller module. If ownership of the source virtual disk is changed during the virtual disk copy, ownership of the target virtual disk is also changed.

Failed RAID Controller Module

You must manually change RAID controller module ownership to the alternate RAID controller module to allow the virtual disk copy to complete under all of these conditions:

- A virtual disk copy has a status of **In Progress**
- The preferred RAID controller module of the source virtual disk fails
- The ownership transfer does not occur automatically in the failover

Copy Manager

After you create a virtual disk copy by using the **Create Copy Wizard**, you can monitor the virtual disk copy through the **Copy Manager**. From the **Copy Manager**, a virtual disk copy may be re-copied, stopped, or removed. You can also modify the attributes, such as the copy priority and the target virtual disk Read-Only attribute. You can view the status of a virtual disk copy in the **Copy Manager**. Also, if you need to determine which virtual disks are involved in a virtual disk copy, you can use the **Copy Manager** or the storage array profile.

Copying the Virtual Disk

You can create a virtual disk copy by using the **Create Copy Wizard**.

⚠ **CAUTION: Possible loss of data access – A virtual disk copy overwrites data on the target virtual disk.**

A virtual disk copy automatically makes the target virtual disk read-only to hosts. You may want to keep this attribute enabled to preserve the data on the target virtual disk.

⚠ **CAUTION: If you decide not to preserve the data on the target virtual disk after the virtual disk copy has completed, disable the Read-Only attribute for the target virtual disk. For more information on enabling and disabling the Read-Only attribute for the target virtual disk, see "Virtual Disk Read/Write Permissions" on page 172.**

To prevent write-protected error messages from being displayed, do not try to access a source virtual disk that is participating in a virtual disk copy while the virtual disk copy has a status of **In Progress**. Also, ensure that the Read-Only attribute for the target virtual disk is disabled after the virtual disk copy has completed to prevent error messages from being displayed.

To copy the virtual disk:

1 Stop all I/O activity to the source virtual disk and the target virtual disk.

2 Unmount any file systems on the source virtual disk and the target virtual disk.

3 In the AMW, select the **Logical** tab and select the source virtual disk.

4 Select **Virtual Disk→ Create Copy**.

The **Select Source Virtual Disk and Copy Type** window is displayed.

5 In the **Select source virtual disk** area, select the appropriate virtual disk.

6 In the **Select Copy Type** area, select either **Offline** or **Online Copy Type**.

✐ **NOTE:** An online virtual disk copy overwrites data on the target virtual disk and automatically makes the target virtual disk read-only to hosts. After the online virtual disk copy completes, use Copy Manager to disable the Read-Only attribute for the target virtual disk. If you have used the target virtual disk in a virtual disk copy before, make sure that you no longer need that data or have backed it up in an accessible location.

The **Select Target Virtual Disk** window is displayed.

7 In the **Select target virtual disk** area, select the appropriate virtual disk

8 In the **Select copy priority area**, select the relevant copy priority and click **Next**.

The **Confirmation** window displays the summary of your selections.

9 Type **yes** and click **Finish**.

🖉 **NOTE:** The **Operation in Progress** icon is displayed on the source virtual disk and the target virtual disk while the virtual disk copy has a status of **In Progress** or **Pending**.

Storage Array Performance During Virtual Disk Copy

The following factors contribute to the overall performance of the storage array:

- I/O activity
- Virtual disk RAID level
- Virtual disk configuration—Number of drives in the virtual disk groups
- Virtual disk type—Snapshot virtual disks may take more time to copy than standard virtual disks

During a virtual disk copy, resources for the storage array are diverted from processing I/O activity to completing a virtual disk copy. This affects the overall performance of the storage array. When you create a new virtual disk copy, you define the copy priority to determine how much controller processing time is diverted from I/O activity to a virtual disk copy operation.

Setting Copy Priority

You can use the **Copy Manager** to select the rate at which a virtual disk copy completes for a selected copy pair. You can change the copy priority for a copy pair at any of these times:

- Before the virtual disk copy begins
- While the virtual disk copy has a status of **In Progress**
- When you re-create a virtual disk copy

To set copy priority:

1 In the AMW, select **Virtual Disk→ Copy Manager**.

The **Copy Manager** window is displayed.

2 In the table, select one or more copy pairs.

3 Select **Change**→ **Copy Priority**.

The **Change Copy Priority** window is displayed.

4 In the **Copy priority** area, select the appropriate copy priority, depending on your system performance needs.

> ✎ **NOTE:** There are 5 copy priority rates available: lowest, low, medium, high, and highest. If the copy priority is set at the lowest rate, I/O activity is prioritized and the virtual disk copy takes longer.

Stopping a Virtual Disk Copy

You can stop a virtual disk copy operation that has an **In Progress** status, a **Pending** status, or a **Failed** status. Stopping a virtual disk copy that has a **Failed** status clears the **Needs Attention** status displayed for the storage array.

Keep these guidelines in mind when you stop a virtual disk copy:

- To use this option, select only one copy pair in the **Copy Manager**.
- When the virtual disk copy is stopped, all of the mapped hosts have write access to the source virtual disk. If data is written to the source virtual disk, the data on the target virtual disk no longer matches the data on the source virtual disk.

To stop a virtual disk copy, complete the following steps:

1 In the AMW, select **Virtual Disk**→ **Copy Manager**.

The **Copy Manager** window is displayed.

2 Select the copy pair in the table.

3 Select **Copy**→ **Stop**.

4 Click **Yes**.

Recopying a Virtual Disk

You can recopy a virtual disk when you have stopped a virtual disk copy and you want to start it again or when a virtual disk copy has failed.

The **Recopy** option overwrites existing data on the target virtual disk and makes the target virtual disk read-only to hosts. This option fails all snapshot virtual disks associated with the target virtual disk, if any exist.

Preparing Host Servers to Recopy a Virtual Disk

✐ **NOTE:** Before you create a new copy of a source virtual disk, stop any data access (I/O) activity or suspend data transfer to the source virtual disk (and, if applicable, the target disk) to ensure that you capture an accurate point-in-time image of the source virtual disk. Close all applications, including Windows Internet Explorer, to ensure all I/O activity has stopped.

✐ **NOTE:** Removing the drive letter of the associated virtual disk(s) in Windows or unmounting the virtual drive in Linux helps to guarantee a stable copy of the drive for the virtual disk copy.

Before creating a new virtual disk copy for an existing copy pair, both the host server and the associated virtual disk you are recopying must be in the proper state. Perform the following steps to prepare your host server and virtual disk:

1 Stop all I/O activity to the source and target virtual disk.

2 Using your Windows system, flush the cache to both the source and the target virtual disk (if mounted). At the host prompt, type

```
SMrepassist -f <filename-identifier>
```

and press <Enter>. For more information, see "SMrepassist Utility" on page 289.

3 To ensure that the virtual disk is in **Optimal** or **Disabled** status, select the **Summary** tab, then click **Disk Groups & Virtual Disks**.

4 Remove the drive letter(s) of the source and (if mounted) virtual disk in Windows or unmount the virtual drive(s) in Linux to help guarantee a stable copy of the drive for the virtual disk. If this is not done, the copy operation reports that it has completed successfully, but the copied data is not is not updated properly.

5 Follow any additional instructions for your operating system. Failure to follow these additional instructions can create unusable virtual disk copies.

✐ **NOTE:** If your operating system requires additional instructions, you can find those instructions in your operating system documentation.

Re-Copying a Virtual Disk

You can use the **Copy Manager** to create a new virtual disk copy for a selected source virtual disk and a target virtual disk. Use this option when you have stopped a virtual disk copy and want to start it again or when a virtual disk copy has failed or completed. The virtual disk copy starts over from the beginning.

- Possible loss of data—The re-copying operation overwrites existing data on the target virtual disk.

- Possible loss of data access—While a virtual disk copy has a status of **In Progress** or **Pending,** source virtual disks are available for read I/O activity only. Write requests are allowed after the virtual disk copy has completed.

Keep these guidelines in mind when re-copying a virtual disk:

- If hosts are mapped to the source virtual disk, the data that is copied to the target virtual disk when you perform the re-copy operation may have changed since the previous virtual disk copy was created.

- Select only one virtual disk copy in the **Copy Manager** dialog.

To re-copy the virtual disk:

1. Stop all I/O to the source virtual disk and the target virtual disk.

2. Unmount any file systems on the source virtual disk and the target virtual disk.

3. In the AMW, select **Virtual Disk** → **Copy Manager**.

   The **Copy Manager** window is displayed.

4. Select the copy pair in the table.

5. Select **Copy**→ **Re-Copy**.

   The **Re-Copy** window is displayed.

6. Set the copy priority.

🖉 **NOTE:** There are 5 copy priority rates available: lowest, low, medium, high, and highest. If the copy priority is set at the lowest rate, I/O activity is prioritized, and the virtual disk copy takes longer. If the copy priority is set to the highest priority rate, the virtual disk copy is prioritized, but I/O activity for the storage array may be affected.

Removing Copy Pairs

You can remove one or more virtual disk copies by using the **Copy Manager.**
Any virtual disk copy-related information for the source virtual disk and the
target virtual disk is removed from the **Virtual Disk Properties** dialog and the
Storage Array Profile dialogs. When you remove a virtual disk copy from the
storage array, the Read-Only attribute for the target virtual disk is also removed.
After the virtual disk copy is removed from the **Copy Manager**, you can either
select the target virtual disk as a source virtual disk or the target virtual disk for
a new virtual disk copy. If you remove a virtual disk copy, the source virtual disk
and the target virtual disk are no longer displayed in the **Copy Manager**.

Keep these guidelines in mind when you remove copy pairs:

* Removing copy pairs does not delete the data on the source virtual disk or
  target virtual disk.
* If the virtual disk copy has a status of **In Progress**, you must stop the
  virtual disk copy before you can remove the copy pair.

To remove copy pairs:

1 In the AMW, select **Virtual Disk→ Copy Manager**.

  The **Copy Manager** window is displayed.

2 In the table, select one or more copy pairs.

3 Select **Copy→ Remove Copy Pairs**.

  The **Remove Copy Pairs** dialog is displayed.

4 Click **Yes**.

13

Configuration: Premium Feature— Remote Replication

The Remote Replication premium feature enables real-time replication of data between two storage arrays in separate locations. Using remote replication, up to 16 replicated virtual disk pairs can be created on a storage array.

A replicated virtual disk pair is created from two standard virtual disks, both of which are logical data storage structures created on a storage array. A standard virtual disk can only be a member of one replicated pair. A pair consists of a primary virtual disk at a local storage array and a secondary virtual disk at a remote storage array. Data written by the host to the primary virtual disk is replicated directly from the primary storage array to the secondary virtual disk. No host server or application resources are used during replication. The data can be written either synchronously or asynchronously.

In the event of a disaster, loss of communications or catastrophic failure in the local storage array, the secondary virtual disk in the remote storage array can be promoted to the role of primary virtual disk and assume responsibility for maintaining business operations.

For detailed steps on creating virtual disks and implementing remote replication, see "Setting Up and Managing Remote Replication" on page 191.

Switchable Host Access Configuration Required With Remote Replication

Access to the host server from the local and/or remote storage arrays using remote replication must be established through a supported Fibre Channel switch configuration. Direct-attached configurations between the host server and Dell PowerVault MD storage arrays are not supported.

For more information on configuring the switch environment, see "Required Switch Zoning Configurations" on page 190.

Activating Remote Replication

📝 **NOTE:** If you ordered the Remote Replication feature, you received a Premium Feature Activation card in the Dell PowerVault MD storage array shipping box. Follow the directions on the card to install and enable this feature.

After the Remote Replication feature is installed following the steps on the Premium Feature Activation card, it must be activated using a wizard-based process. As part of activation, the following actions occur:

- Two replication repository virtual disks are created using existing capacity on either a disk group or storage array specified by the user.

- All hosts currently using the highest-numbered Fibre Channel host port (Port 3) on each RAID controller are logged out.

- The highest-numbered Fibre Channel host port (Port 3) on each RAID controller is reserved for replication data transmission.

- The highest-numbered Fibre Channel host port (Port 3) on each RAID controller is not available to direct host access when the Remote Replication premium feature is activated. Any attempted host communication to this port is rejected.

- Communication between the replication-reserved Fibre Channel ports requires established zoning of Fibre Channel switches in the fabric, see "Required Switch Zoning Configurations" on page 190.

- Before establishing virtual disk replication pairs, the Remote Replication premium feature must be activated on each of the storage arrays participating in replication.

Replication Repository Virtual Disks

Using the Remote Replication premium feature on the storage array requires that two replication repository virtual disks be created on one of the disk groups on the storage array. The RAID controller module stores replication information on these virtual disks, including information about incomplete remote writes that can be used to recover from RAID controller module resets and accidental storage array shutdown.

Other information about replication repository virtual disks include:

- Specifying that replication repository virtual disks be created from unconfigured free capacity on the disk group, or creating a new disk group and its member replication repository virtual disks from unconfigured free capacity on the storage array.

  📝 **NOTE:** The disk group containing replication repository virtual disks cannot be deleted without disabling the Remote Replication premium feature.

- Default names of the replication repository virtual disks are **Replication Repository 1** and **Replication Repository 2**. These names cannot be changed.

- The activation process creates replication repository virtual disk with identical capacity. In a dual RAID controller module storage array, the default capacity for each replication repository virtual disks is 128 MB per controller (requiring a total of 256 MB in the RAID group). This default capacity cannot be modified.

RAID Levels for Replication Repository Virtual Disks

If replication repository virtual disks are created from unconfigured free capacity on the storage array, you can select a RAID level for the disk group of either RAID 1/10, RAID 5, or RAID 6. However, if the replication repository virtual disks are created from an existing disk group, RAID levels are not user configurable.

Primary and Secondary Virtual Disk Pairs

After the Remote Replication feature is activated and configuration of the FC fabric, individual virtual disk replication pairs are be created. Up to 16 replication virtual disk pairs can be established. Each pair consists of a primary (online) virtual disk and a secondary (offline) virtual disk, both of identical capacity. During operation, primary and secondary roles can be interchanged to test the replicated data sets without breaking the relationship. Independent copies may be created by deleting the relationship, in which case both virtual disks are available online to their configured hosts.

Before creating a virtual disk pair, verify the following:

1 The Remote Replication premium feature is installed, enabled and activated on the primary and secondary storage arrays.

2 The storage arrays containing the two virtual disks you want to replicate are connected through a Fibre Channel fabric interface.

3 The secondary storage array contains a virtual disk that is equal (or larger) in capacity than the primary virtual disk in the replicated pair.

4 Both virtual disks are visible to the storage management software.

5 All I/O into the secondary virtual disk is stopped and all file systems are unmounted.

⚠ **CAUTION: Creating a replicated virtual disk pair starts a synchronous process between the primary and secondary virtual disks. This process overwrites all existing data on the secondary virtual disk and sets the disk to read-only access. If you have existing data on your secondary virtual disk, back up the data before creating the replicated virtual disk pair.**

Using Remote Replication With Other Features

Remote replication can be used with other disk utility features such as:

- Storage Partitioning
- Snapshot Virtual Disk Premium Feature
- Virtual Disk Copy Premium Feature
- Virtual Disk Expansion (VDE)

Snapshot Virtual Disk and Virtual Disk Copy are premium features that must be enabled and activated on the primary storage array.

Storage Partitioning With Remote Replication

Storage partitioning allows hosts to share access to virtual disks in a storage array. A storage partition is created when a collection of hosts (host group) or a single host is defined, then a virtual disk-to-logical unit number (LUN) mapping is defined. This mapping allows you to define which host or host group has access to a particular virtual disk in the storage array.

Storage partition definitions for the local storage array and remote storage array are independent. Establishing similar definitions on the disk serving in a secondary role reduces administrative effort required during site recovery.

Snapshot Virtual Disk With Remote Replication

A snapshot virtual disk is a point-in-time image of a virtual disk. In a remote replication, do not mount a snapshot virtual disk on the same server as the primary virtual disk.

Virtual Disk Copy With Remote Replication

The virtual disk copy premium feature copies data from a source virtual disk to a target virtual disk within the same storage array.

A primary virtual disk in a remote replication can be either a source virtual disk or a target virtual disk in a virtual disk copy.

You can create a virtual disk copy on the primary virtual disk in a replicated pair, but you cannot create a virtual disk copy of a secondary virtual disk in a replicated pair. Instead, you can copy a secondary virtual disk in one of two ways:

- Promote the secondary virtual disk to the primary role.

⚠ **CAUTION: If a primary/secondary role reversal is begun while a virtual disk copy is in progress, the virtual disk copy fails and be unable to restart.**

- Create a snapshot virtual disk of the secondary virtual disk, then perform a virtual disk copy of the snapshot virtual disk.

Virtual Disk Expansion With Remote Replication

Virtual Disk Expansion (VDE) increases the capacity of an existing virtual disk. This increased capacity is achieved by using the free capacity available on either:

- The disk group of the standard virtual disk, or
- The snapshot repository virtual disk.

Performing a VDE operation does not interrupt access to data on disk groups, virtual disks or physical disks.

A VDE operation can be performed on a primary virtual disk or a secondary virtual disk of a replicated pair. It cannot be performed on a replication repository virtual disk. No other action can be performed on a virtual disk while it is undergoing expansion (except normal disk I/O).

📝 **NOTE:** To perform a VDE operation, the remote replication must be in Optimal status. The properties pane in the logical view displays the status of a virtual disk.

Required Switch Zoning Configurations

🖉 **NOTE:** A limited set of Fibre Channel switches are supported by the storage array. For a complete list of supported switches and their required driver, BIOS and firmware levels, see the *Configuring Fibre Channel with the Dell MD3600f-Series Storage Array.*

Due to potential restrictions at the host level, remote replication configurations must contain Fibre Channel switches. Fibre Channel switches must be zoned so that a single host adapter accesses only one RAID controller module in a storage array. Additionally, all zone configurations must specify separate zones for ports that are reserved for remote replication functions.

🖉 **NOTE:** Do not zone the uplink port (E_port) that connects (cascades) switches within a fabric.

Switch zoning configurations are typically set using the switch management software provided by the Fibre Channel switch vendor. This software is either included with the switch itself, or downloaded from the vendor's website.

Zoning Guidelines for Remote Replication

Enabling remote replication on your storage array automatically dedicates one port per RAID controller exclusively to replication. These dedicated ports, Port 3 on RAID controller 0 and Port 3 on RAID controller 1, must be in different zones on the Fibre Channel switch. Generally, switch zoning must be established so that:

- All RAID Controller 0 Fibre Channel Port 3s are in a common zone.
- All RAID Controller 1 Fibre Channel Port 3s are in a common zone, but separately zoned from RAID Controller 0 ports.
- Host Fibre Channel access ports are in separate zones from those established for remote replication.

Switch Cascading

When two or more Fibre Channel switches are cascaded, the switch management software combines the ports for all of the switches that are linked.

Journaling File Systems and Remote Replication

When using a journaling file system, you cannot gain read-only access to a remote virtual disk. A journaling file system does not let you mount the remote virtual disk in Windows NTFS. However, you can mount a snapshot of the remote virtual disk, if available.

Setting Up and Managing Remote Replication

Prerequisites

Before setting up remote replication between two storage arrays, ensure that:

- The Remote Replication premium feature is activated.

- The local storage array contains two replication repository virtual disks (created during activation).

- The local storage array contains the primary virtual disk to be replicated.

- The remote storage array contains the secondary virtual disk to serve as the replication target, and:

  - Both virtual disks are visible to a single management station; replicated disk pairs cannot be managed on separate instances of storage management software.

  - The RAID level of the secondary virtual disk can be different than that of the primary virtual disk.

  - The capacity of the secondary virtual disk (after creation and formatting) must be equal to or greater than the capacity of the primary virtual disk.

- The storage arrays containing the virtual disks you want to replicate are connected to each other through a Fibre Channel switch fabric and are accessible from a single management station.

If either the primary or secondary virtual disk does not exist, you must create it before proceeding with remote replication setup. See "Activating the Remote Replication Premium Feature and Creating Replication Virtual Disks" on page 192 and complete the necessary steps.

If all virtual disks are created, skip to "Creating a Remote Replication" on page 194.

Activating the Remote Replication Premium Feature and Creating Replication Virtual Disks

To activate the Remote Replication feature and create the two replication repository virtual disks required, you can either:

- Activate the feature, then create the two replication virtual disks from the total unconfigured capacity on the storage array, or
- Activate the feature, then create the two replication virtual disks from the unconfigured capacity in an already-existing disk group

Activating and Creating Replication Repository Virtual Disks From the Storage Array

To activate the Remote Replication feature and create the two replication repository virtual disks required, you can either:

1. In the Activate Remote Replication wizard, select Unconfigured capacity (create a new disk group) and click **Next.**

   The **Activate Remote Replication** - **Create Disk Group** wizard is displayed.

2. In **Disk Group Name**, type a unique name for the disk group.

3. Select one to select a physical disk:

   - **Automatic**—The storage management software generates a list of available capacity and physical disk options for each available RAID level.

   - **Manual**—The storage management software generates a list of unselected physical disks.

4. Click **Next**.

   - If you selected **Automatic**, an empty **Select Capacity** table and a drop-down list of available RAID levels is displayed. Go to Step 5.

   - If you selected Manual, a populated **Unselected Physical Disks** table, empty **Selected Physical Disks** table, and a drop-down list of available RAID levels is displayed.

5. In the **Select RAID level** dropdown list, select the RAID level for the disk group. The **Select capacity** table displays the available virtual disks for the RAID level.

6 In the **Select capacity** table, select the physical disks and capacities for the new disk group, then click **Next.**

The **Preview (Activate Remote Replication)** wizard is displayed.

7 Click **Finish.**

The Completed (Activate Remote Replication) message is displayed.

8 Click **OK**.

The Remote Replication premium feature is now active. The **Logical** pane displays the new disk group and the two replication repository virtual disks.

Activating and Creating Replication Repository Virtual Disks From an Existing Disk Group

Using this method for creating the replication repository virtual disks, free capacity in an existing disk group is used. By default, the replication repository virtual disks each have either 128 MB or 256 MB capacity. You cannot create the replication repository virtual disks on a disk group with insufficient capacity, or change the default capacities of the replication repository virtual disks.

1 In the Introduction (Activate Remote Replication) wizard, select **Free capacity on existing disk groups.**

2 From the list of available disk groups, select a disk group to contain the replication repository virtual disks, then click **Next**.

The **Preview (Activate Remote Replication)** wizard is displayed.

3 Click **Finish**.

The **Completed (Activate Remote Replication)** message is displayed.

4 Click **OK**.

The Remote Replication premium feature is now active. The **Logical** pane displays the two replication repository virtual disks as part of the disk group.

Creating a Remote Replication

Before creating a remote replication, ensure that all prerequisites are met. See "Prerequisites" on page 191 for more information.

1 Open the AMW of both the local and remote storage array.

2 Verify that the Remote Replication premium feature is activated on both storage arrays.

3 In the **AMW** of the local storage array, select the **Logical** tab.

4 In the **Logical** pane of the local storage array, select the virtual disk you created to serve as your primary virtual disk. This is the disk that is replicated to the secondary disk.

5 In the menu bar in the **AMW**, select **Virtual Disk**→ **Remote– Replication**→ **Create.**

   The **Introduction (Activate Remote Replication)** wizard is displayed.

6 Click **Next.**

   The **Select Storage Array (Create Remote Replication)** dialog is displayed. The Storage Arrays list shows the remote storage arrays.

7 Select the storage array where you created the secondary virtual disk, then click **Next.**

   The **Select Secondary Virtual Disk (Create Remote Replication)** wizard is displayed. Go to "Selecting the Secondary Virtual Disk" on page 195.

Selecting the Secondary Virtual Disk

⚠ **CAUTION: Creating a replicated virtual disk pair starts a process between the primary and secondary virtual disks that overwrites all existing data on the secondary virtual disk and set the disk to read-only access. If you have existing data on your secondary virtual disk stop all I/O to the disk, back up the data, and unmount any file systems mounted to the secondary virtual disk before creating the replicated virtual disk pair.**

1   Select the remote storage array and the primary virtual disk.

2   In the **Select Secondary Virtual Disk (Create Remote Replication)** wizard, select the secondary virtual disk.

  📝 **NOTE:** The secondary virtual disk must have a capacity equal to or greater than the capacity of the primary virtual disk.

3   Click **Next**.

  The **Set Write Mode (Create Remote Replication)** wizard is displayed. Go to "Setting the Write Mode" on page 195.

Setting the Write Mode

The secondary host ports on the storage arrays are reserved for data synchronization between the primary virtual disk and the secondary virtual disk in a replicated virtual disk pair. You can set remote replication to write either synchronously or asynchronously.

•   In the synchronous mode, the RAID controller module on the primary virtual disk sends an I/O completion message back to the host storage array after the data is successfully copied to the secondary storage array. The synchronous mode is the preferred mode of operation, since it offers the best chance of full data recovery from the secondary storage array in the event of a disaster. However, a synchronous replication mode can degrade the I/O performance of the host when long-distance data transfer is involved.

•   In the asynchronous mode, the RAID controller module on the primary storage array sends an I/O completion message to the host storage array before the data is successfully copied to the secondary storage array. The asynchronous mode offers faster host I/O performance; however, it does not guarantee that data was successfully written to the secondary virtual disk or that the write requests were completed on the secondary virtual disk in the same order they were initiated. In asynchronous mode, you can

also specify whether to use a write consistency group option. A write consistency group ensures that the secondary virtual disk receives write requests in the sequence initiated by the RAID controller module of the primary virtual disk. The secondary virtual disk can also be added to a write consistency group.

To set the write mode for the remote replication:

1 In the **Set Write Mode (Create Remote Replication)** wizard, select either the **Synchronous** or **Asynchronous** mode.

2 If you select asynchronous mode, select whether to use the **Add to write consistency group** option.

3 Click **Next**.

   The **Select Synchronization Settings (Create Remote Replication)** wizard is displayed. Go to "Setting Synchronization Priority and Synchronization Method" on page 196.

Setting Synchronization Priority and Synchronization Method

You can set the priority for allocating system resources to synchronizing the remote replication.

• Higher synchronization priorities allocate more resources to the process and may degrade I/O performance.

• Lower synchronization priorities allocate fewer resources to the process and have less impact on normal I/O performance.

The initial synchronization priority and synchronization method you select can be changed later. For more information about resynchronizing virtual disks in a remote replication, go to "Resynchronizing Virtual Disks" on page 204.

To set Synchronization Priority and Synchronization Method:

1 In the **Select Synchronization Settings (Create Remote Replication)** wizard, select the synchronization priority on the slide bar.

2 Select either **Manual** or **Automatic** resynchronization.

   • **Automatic resynchronization**—Resynchronization starts immediately after communication is restored between unsynchronized replicated virtual disks.

- **Manual resynchronization** (recommended)—The replicated pair must be manually resynchronized each time communication is restored between unsynchronized replicated virtual disks.

3 Click **Next**.

The **Preview (Create Remote Replication)** wizard is displayed. Go to "Completing the Remote Replication" on page 197.

Completing the Remote Replication

⚠ **CAUTION: Creating a replicated virtual disk pair starts a process between the primary and secondary virtual disks that overwrites ALL existing data on the secondary virtual disk and set the disk to read-only access. If you have existing data on your secondary virtual disk: (1) stop all I/O to the disk, (2) back up the data and (3) unmount any file systems mounted to the secondary virtual disk before creating the replicated virtual disk pair.**

After selecting the synchronization settings, perform these steps to complete the remote replication:

1 In the **Preview (Create Remote Replication)** wizard, type **Yes** to confirm the synchronization. Then, click **Finish**.

The **Creation Successful (Create Remote Replication)** confirmation message is displayed.

2 If you want to create another replication pair, click **Yes** and repeat the steps. Otherwise, click **No** to exit.

- In the Array Management Windows of both the local storage array and the remote storage array, the **Logical** panes show the replicated virtual disk pairs as members of their disk groups.

- In the local storage array, the **Properties** pane displays the Replication status as Synchronizing and the Synchronization – Progress bar shows estimated time to completion.

To view detailed information about the virtual disks in a remote replication, go to either "Viewing Information About a Remote Replication or Replication Repository Virtual Disk" on page 198 in the Properties Pane or Viewing Information about a Remote Replication or Replication Repository Virtual Disk in the Storage Array Profile. Each virtual disk is also represented by a status icon in the Devices tab on the Enterprise Management Window. For a description of each icon, see "Virtual Disk Status Icons" on page 202.

RAID Controller Module Ownership/Preferred Path

During a remote replication operation, the RAID controller module that owns the primary virtual disk must correspond to the RAID controller owning the secondary volume on the remote array. If both virtual disks do not have the same preferred RAID controller module when a remote replication begins, the ownership of the secondary virtual disk is automatically transferred to the preferred RAID controller module of the primary virtual disk.

- When the remote replication is completed or is stopped, ownership of the secondary virtual disk is restored to its preferred RAID controller module.
- If ownership of the primary virtual disk is changed during the remote replication, ownership of the secondary virtual disk is also changed.

If any of the following conditions apply to your RAID controller module, you must manually change module ownership to the alternate RAID controller module to allow the remote replication to finish:

- The status of the remote replication is **In Progress**.
- The preferred RAID controller module of the primary virtual disk fails.
- Transfer of ownership does not occur automatically during failover.

⚠ **CAUTION: To avoid possible loss of data, verify that either the virtual disks are not in use, or a multi-path driver is installed on the host. If you change the RAID controller module ownership/preferred path while an application is using one of the virtual disks, I/O activity is disrupted and I/O errors may occur.**

Viewing Information About a Remote Replication or Replication Repository Virtual Disk

In MD Storage Manager, the **Storage Array Profile** and the **Properties** pane display information about the physical and logical characteristics of the remote replication.

Viewing the Storage Array Profile

The Storage Array Profile displays the most detailed information about the components of a remote replication and the replication repository virtual disks. You can:

- View detailed information about individual and paired virtual disks in a remote replication
- View detailed information about the replication repository virtual disks in the remote replication and storage array
- Save the Storage Array Profile information as a text file

To view the Storage Array Profile:

1 In the **Array Management Window** of the local or remote storage array, select the **Summary** tab.

2 In the **Status** area, click **Storage Array Profile**.

    The **Storage Array Profile** is displayed.

3 Select the **Virtual Disks** tab.

4 Select either the **Replications** tab or the **Repositories** tab. The Profile for Storage array page is displayed.

5 Perform one of the following:

    - To return to the **Array Management Window** without saving the information, click **Close**.
    - To save the information, click **Save As**, then go to Step 6.

6 In the **Section Selection** area of the **Save Profile** window, select each section you want to save.

7 To save the file, either:

    - Save the file in the default directory, or
    - Save the file in another directory by choosing it in the Look in drop-down list.

8 In **File name**, type a name for the file and click **Save**.

    The file is saved as a .txt file.

Viewing the Properties Pane

The Properties pane is a view-only display of the physical and logical characteristics of a single virtual disk in a replicated pair, or a single replication repository virtual disk.

To view the Properties pane:

1   In the **AMW**, select the **Logical** tab.

2   Select either the primary or secondary virtual disk of the replicated pair.

The Properties pane shows properties for the selected virtual disk. Under Replication properties, the Replication status displays the synchronization status of the replicated pair. When the primary and secondary virtual disks are synchronizing, the Replication status shows a synchronizing icon see "Virtual Disk Status Icons" on page 202).

Viewing Logical Elements of the Secondary Virtual Disk

To view the logical elements of the secondary virtual disk of a remote replication:

1   In the AMW of the local storage array, select the **Logical** tab.

2   Right-click the secondary virtual disk of the remote replication.

3   Select **View Associated Logical Elements**.

The **View Associated Logical Elements** pop-up is displayed, revealing these logical elements:

•   Primary and secondary virtual disks and their locations

•   Replication repository virtual disks and their locations

Viewing the Physical Components or Logical Elements of the Primary Virtual Disk

To view the physical components or logical elements of the primary virtual disk in a remote replication:

1 In the AMW of the storage array that contains the primary virtual disk, select the **Logical** tab.

2 Right-click the primary virtual disk and either:

- View the logical elements of the primary virtual disk by selecting **Select View→ Associated Logical Elements**. The **View Associated Logical Elements** pop-up is displayed, showing visual representations of these elements: the primary and secondary virtual disk in the remote replication (and their locations) and the replication repository virtual disks in the storage array (and their locations). Or,

- View the physical components of the primary virtual disk by clicking **View Associated Physical Components** in the Properties pane. The **View Associated Physical Components** pop-up is displayed, showing a visual representation of the primary virtual disk in the remote replication.

Virtual Disk Status Icons

In the **Devices** tab on the EMW, the following icons depict the status of each virtual disk:

Changing Write Mode and Consistency Group Membership

Table 13-1. Write Mode and Consistency Group Membership

Icon	Status
	Replication repository virtual disk
	Primary virtual disk
	Secondary virtual disk
	Primary virtual disk, synchronization in progress
	Primary virtual disk, replication suspended
	Secondary virtual disk, replication suspended
	Primary virtual disk, unsynchronized
	Secondary virtual disk, unsynchronized with communications loss

The write mode of a remote replication is selected when it is created. However, you can change the write mode in a remote replication at a later time, as well as change the secondary virtual disk's membership in a write consistency group. For more information about write modes and write consistency groups, go to Setting the Write Mode for Creating a Remote Replication online help topic.

To change Write Mode and Consistency Group Membership:

1 In the **AMW** the storage array containing the primary virtual disk, select the **Logical** tab.

2 In the **Logical** pane, right-click the primary virtual disk of the replicated pair.

3 Select **Change→ Write Mode**.

   The **Change Write Mode** dialog is displayed. The Replicated pairs table reflects all replicated pairs in both the local and remote storage arrays.

4 Select one or more replicated pairs. To select all replicated pairs, click **Select All**.

5 Select either the **Synchronous** or **Asynchronous** write mode.

6 If you are you adding the secondary virtual disk of the replicated pair to a write consistency group, select **Yes** and then select the **Add to consistency group** check box.

   If you are not adding the secondary virtual disk of the replicated pair to a write consistency group, go to Step 7.

7 Click **OK**.

   The **Change Write Mode confirmation** message is displayed.

8 In the **Change Write Mode confirmation** message, click **Yes**.

The Replication properties section (the Properties panes in the Array Management Window for the local storage array) displays the following information:

• Replication status is Synchronized.

• Write mode is either Synchronous or Asynchronous, depending on your selection above.

• The secondary virtual disk is either write consistent/not write consistent.

• Resynchronization method is either Manual or Automatic.

Resynchronizing Virtual Disks

There are two resynchronization methods for remote replications:

- Manual resynchronization—See "Manually Resynchronizing Virtual Disks" on page 208.
- Automatic resynchronization—See "Automatically Resynchronizing Virtual Disks" on page 208.

For more information about synchronization and resynchronization in remote replications, see:

- "Normally Synchronized Virtual Disks" on page 204.
- "Unsynchronized Virtual Disks" on page 207.
- "Setting Synchronization Priority and Synchronization Method" on page 196.
- "Changing Synchronization Priority and Method" on page 205.
- "Resynchronizing Virtual Disks" on page 204.

You may need to periodically test the communication between the primary virtual disk and the secondary virtual disk in a remote replication, especially after resynchronizing virtual disks. For information about testing the communication, go to Testing Communication between the Primary and Secondary Virtual Disk.

Normally Synchronized Virtual Disks

In a normally synchronized remote replication, RAID controller module owners manage the transfer of data from the primary to the secondary virtual disk. In a normal remote replication, these events occur:

1 The primary virtual disk receives a write request from a host.

2 The RAID controller module owner on the storage array logs information about the write operation to a replication repository virtual disk in the storage array.

3 The RAID controller module owner writes the data to the primary virtual disk.

4 The RAID controller module owner starts a data transfer operation to the secondary virtual disk on the secondary storage array.

The communication between a primary and secondary virtual disks can be suspended or become unsynchronized. If communication between the primary virtual disk and the secondary virtual disk is disrupted:

1 The status of the replicated pair changes to **Unsynchronized**.

2 A **Needs Attention** status is displayed for the storage array.

3 Data is written to the primary virtual disk.

4 Write requests to the primary virtual disk are logged.

5 The RAID controller module owner sends an I/O completion message to the host sending the write request. Although the host can continue to send write requests to the primary virtual disk, no data transfer takes place to the secondary virtual disk. Writes to the secondary virtual disk are suspended pending restoration of communications between the primary virtual disk and the secondary virtual disk.

When connectivity is restored between the primary virtual disk and the secondary virtual disk, the replicated pair is ready to be resynchronized.

> **NOTE:** When the primary virtual disk and the secondary virtual disk are resynchronized, only data that has changed on the primary virtual disk after the break in communication is transferred to the secondary virtual disk.

> **CAUTION: You must be aware of a possible loss of data if communication is broken after resynchronization starts between the primary storage array and the secondary storage array. New data may mix with old data on the secondary virtual disk and render the data unusable in a disaster recovery situation.**

Changing Synchronization Priority and Method

Synchronization priority defines how much processing time and resources are allocated to synchronizing the primary and secondary virtual disk of a remote replication relative to system performance. Increasing the synchronization priority of a remote replication may degrade system performance. Synchronization priorities can be reset at any time.

Synchronization priorities can affect these operations:

- Performing a copyback
- Performing a Virtual Disk Expansion (VDE)
- Reconstructing a virtual disk
- Initializing a virtual disk

- Changing the segment size of a virtual disk
- Defragmenting a disk group
- Adding free capacity to a disk group
- Changing the RAID level of a disk group

To change the synchronization priority and method for an existing remote replication:

1 In the **AMW** of the storage array containing the primary virtual disk of the replicated pair, right-click the **Logical** tab.

2 Select **Change**→ **Synchronization** Settings.

   The **Change Synchronization Settings** dialog is displayed.

3 In the Replicated pairs table, select the primary virtual disk and the remote virtual disk you want to change. To select all virtual disks, click **Select All**.

4 On the **Select Synchronization Priority** slide bar, select the synchronization priority for the replicated pair.

5 Select either **Manual** or **Automatic** resynchronization.

   - **Automatic** —Resynchronization starts immediately after communication is restored between unsynchronized replicated virtual disks.

   - **Manual**—The replicated pair must be manually resynchronized each time communication is restored between unsynchronized replicated virtual disks.

6 Click **OK**.

   The **Change Synchronization Settings confirmation** message is displayed.

7 In the **Change Synchronization Settings confirmation** message, click **Yes**.

   The Change Synchronization Priority - Progress bar displays the progress of the resynchronization.

8 Click **OK**.

Unsynchronized Virtual Disks

Communication between a primary and secondary virtual disks can be either suspended, or become unsynchronized. If this communication between both virtual disks is disrupted:

1   The status of the replicated pair changes to **Unsynchronized**.

2   A **Needs Attention** status is displayed for the storage array.

3   Data is written to the primary virtual disk.

4   Write requests to the primary virtual disk are logged.

5   The RAID controller module owner sends an I/O completion message to the host sending the write request. Although the host can continue to send write requests to the primary virtual disk, no data transfer takes place to the secondary virtual disk. Writes to the secondary virtual disk are suspended pending restoration of communications between the primary and secondary virtual disks.

When connectivity is restored between the primary and secondary virtual disks, the replicated pair is ready to be resynchronized.

NOTE: When the primary virtual disk and the secondary virtual disk are resynchronized, only data that has changed on the primary virtual disk after the break in communication is transferred to the secondary virtual disk.

CAUTION: Loss of data can occur if communication is broken after resynchronization starts between the primary storage array and the secondary storage array. In addition to resulting in an incomplete resynchronization that may render the secondary storage array unusable, an interrupted resynchronization could cause a loss of data integrity on both arrays.

Automatically Resynchronizing Virtual Disks

When automatic resynchronization is selected, the RAID controller module owning the primary virtual disk automatically starts resynchronizing the data on the remote replication pair immediately after communication is restored.

💿 **NOTE:** When choosing automatic resynchronization, a possible loss of data can occur if a resynchronization is interrupted in progress. Since another resynchronization begins automatically after communication is restored between the primary virtual disk and the remote virtual disk, data integrity could be affected.

With automatic resynchronization set, you cannot add a secondary virtual disk to a write consistency group. Therefore, write consistency during the resynchronization process is not preserved. The write order is not consistent until the entire group achieves **Optimal** status.

Manually Resynchronizing Virtual Disks

When manual resynchronization is selected, you must manually resynchronize and resume the data transfer on a remote replication after communication is restored between the primary and remote virtual disks. Manual resynchronization is the recommended setting for all remote replications for these reasons:

- You determine when resynchronization starts, which allows you to better manage and/or mitigate impact on I/O performance.
- In a disaster recovery situation, manual resynchronization offers the best chance of retrieving valid data.
- When the secondary virtual disk is in a write consistency group, manual resynchronization preserves the write order.

Reversing Roles Between the Primary and Secondary Virtual Disks

Reversing the roles between the virtual disks promotes the secondary virtual disk to the role of primary virtual disk and demotes the primary virtual disk to the role of secondary virtual. If the primary virtual disk in a remote virtual disk replication fails, you can reverse the roles of the primary virtual disk and the secondary virtual disk to transfer the data back to the restored virtual disk.

⚠ **CAUTION: A potential loss of data access can occur if you reverse roles between the secondary virtual disk and the primary virtual disk while a virtual disk copy is in progress. The role reversal may succeed, but the virtual disk copy fails and cannot be restarted.**

- You cannot perform a virtual disk copy on a secondary virtual disk in a virtual disk remote replication. To create a virtual disk copy of a secondary virtual disk, you must reverse the roles of the secondary virtual disk and the primary virtual disk, then perform the virtual disk copy on the new primary virtual disk.

- While a virtual disk remote replication is synchronizing, you cannot perform a virtual disk copy on either the primary or secondary virtual disk.

- If you reverse roles between a secondary virtual disk with less capacity than the primary virtual disk, the role reversal succeeds but the usable capacity of the new primary virtual disk equals the total capacity of the secondary virtual disk.

Promoting the Secondary Virtual Disk or Demoting the Primary Virtual Disk

To promote the secondary virtual disk to the role of primary virtual disk, or demote the primary virtual disk to the role of secondary virtual disk:

1 In the **AMW** of the storage array that contains the virtual disk you are changing, click the **Logical** tab.

2 Right-click the virtual disk you are changing.

   - If you are promoting the secondary virtual disk to the role of primary virtual disk, select **Change**→ **Role** to Primary. The Change to Primary message is displayed. Clicking **Yes** reverses the roles of the primary and secondary virtual disks.

   - If you are demoting the primary virtual disk to the role of secondary virtual disk, select **Change**→ **Role** to Secondary. The Change to Secondary message is displayed. Clicking **Yes** reverses the roles of the primary and secondary virtual disks.

Suspending a Remote Replication

1 In the AMW of the storage array with the primary virtual disk, select the **Logical** tab.

2 In the **Logical** tab, right-click the primary virtual disk of a replicated pair, then select **Suspend Replication**.

   The **Suspend Replicated Pair** dialog is displayed. The Replicated pairs table shows all replicated pairs in the local and remote storage arrays.

3 Select one or more of the replicated pairs to suspend. To select all replicated pairs, click **Select All**.

4 Click **Suspend.**

   The **Suspend Replication Relationship** - **Confirmation** message is displayed.

5 In the **Confirmation** text box, type **Yes**, then click **OK**.

The Suspend Replicated Pair - Progress bar indicates the progress of the suspension. The **Properties** pane in the AMW that contains the suspended primary virtual disk shows the Replication status as **Suspended**. The suspended icon is displayed next to both the primary virtual disk icon and the secondary virtual disk icon in the **Logical** pane in the AMW.

Resuming a Remote Replication

1 In the **AMW** of the storage array with the primary virtual disk, select the **Logical** tab.

2 In the **Logical** pane, right-click the primary virtual disk of the replicated pair, then select **Resume Replication**.

The **Resume Replicated Pair** dialog is displayed. The **Replicated pairs** table shows all suspended replicated pairs in the local and remote storage arrays.

3 Select one or more replicated pairs. To select all replicated pairs, click **Select All**.

4 Click **Resume**.

The `Resume Replication Pair - Confirmation` message is displayed.

5 In the **Confirmation** text box, type **Yes**.

The remote virtual disk replication resumes.

The Properties panes in the AMW for the local storage array and the remote storage array show the Replication status as Synchronized for both the primary and secondary virtual disks.

About Resumed Remote Replications

When a remote virtual disk replication is suspended, data continues to write to the primary virtual disk, but is not replicated to the secondary virtual disk. Writes to the primary virtual disk are persistently logged to the replication repository virtual disks.

After communications are restored, data transfer between the primary virtual disk and the secondary virtual disk must be resynchronized.

Depending on which resynchronization method you specify, the following occurs:

- **Automatic resynchronization**—Data transfer automatically starts immediately after the virtual disks are resynchronized.

- **Manual resynchronization**—You must manually resume the remote virtual disk replication to restart the data transfer. A suspended remote virtual disk replication stays in a Suspended status until it is manually resumed.

After the remote virtual disk replication resumes, data is automatically written to the secondary virtual disk. Only the regions of the primary virtual disk that changed since the replicated pair was suspended are written to the secondary virtual disk.

⚠ **CAUTION: Possible loss of data access: When you resume a remote virtual disk replication involving a primary or secondary virtual disk that is a member of a write consistency group, any other suspended remote replications for replicated pairs in the write consistency group also resumes.**

✎ **NOTE:** If the write mode is synchronous, you do not need to resynchronize the primary and secondary virtual disk after resuming a remote virtual disk replication.

Testing Communication Between the Primary and Secondary Virtual Disks

When a resynchronization is manual or during a disaster recovery scenario, you may want to test the communication between the primary virtual disk and secondary virtual disk in a remote replication. For more information about synchronization and resynchronization in remote virtual disk replications, go to these topics:

- "Normally Synchronized Virtual Disks" on page 204.
- "Unsynchronized Virtual Disks" on page 207.
- "Setting Synchronization Priority and Synchronization Method" on page 196.
- "Changing Synchronization Priority and Method" on page 205.
- "Resynchronizing Virtual Disks" on page 204.
- "Automatically Resynchronizing Virtual Disks" on page 208.

To test the communication between virtual disks in a remote replication:

1 In the **AMW** of either the primary or secondary virtual disk, select the **Logical** tab.

2 In the **Logical** pane, right-click the virtual disk.

3 Select **Test Replication Communication**.

The Replication Communication Test Progress message is displayed.

Deleting a Virtual Disk From a Replicated Pair in a Storage Array

Follow these steps to delete either a primary virtual disk, a secondary virtual disk, or both virtual disks from a replicated pair in a storage array.

NOTE: Do not remove a replication relationship to back up a replicated virtual disk. To perform backups of either the primary virtual disk or the secondary virtual disk, suspend the remote virtual disk replication so that the replication relationship is not broken.

Deleting a Primary Virtual Disk

CAUTION: Depending on which premium features are enabled on the storage array, deleting a primary virtual disk may delete all associated virtual disks, resulting in a possible loss of data on those virtual disks.

NOTE: You cannot delete a primary virtual disk while it is synchronizing.

When a primary virtual disk is deleted from a remote virtual disk replication, the following occurs:

- The primary virtual disk is deleted from the storage array.

    NOTE: The virtual disk is permanently deleted from the storage array and all data on the primary virtual disk is permanently lost.

- The replication relationship breaks.

- The capacity of the deleted virtual disk becomes unconfigured free capacity in the storage array and is available for creation of new virtual disks.

- The secondary virtual disk becomes a regular, standard virtual disk and is can accept both reads and writes.

- For instructions on how to delete a primary virtual disk, go to Deleting a Primary Virtual Disk in a Replicated Pair from a Storage Array online help topic.

Deleting a Secondary Virtual Disk

✎ **NOTE:** Depending on which premium features are enabled on the storage array, deleting a secondary virtual disk may delete all associated virtual disks, resulting in a possible loss of data on those virtual disks.

✎ **NOTE:** You cannot delete a secondary virtual disk while it is synchronizing.

When a secondary virtual disk is deleted, the replication relationship is removed and the remote virtual disk replication is destroyed.

✎ **NOTE:** Deleting a secondary virtual disk results in the permanent loss of the data on the secondary virtual disk.

For steps describing how to delete a secondary virtual disk, see "Deleting a Virtual Disk From a Replicated Pair in a Storage Array" on page 213.

Deleting a Primary Virtual Disk in a Replicated Pair From a Storage Array

1   Stop all I/O activity to the primary virtual disk and unmount any file systems on the primary virtual disk.

2   In the **AMW** of the storage array containing the primary virtual disk, select the **Logical** tab.

3   In the **Logical** pane, right-click the primary virtual disk, then select **Delete**.

   The **Delete Virtual Disks** dialog is displayed.

4   Select one or more virtual disks to delete, then click **Delete**.

   The **Confirm Delete Virtual Disk(s)** message is displayed.

5   In the confirmation text box, type **Yes**. Then, click **OK**.

   The **Delete Virtual Disks** - **Progress** bar is displayed.

6   When the deletion is complete, click **OK**.

The primary virtual disk is deleted from the storage array. The secondary virtual disk in the replicated pair now becomes a regular standard virtual disk in the storage array.

✎ **NOTE:** The primary virtual disk is deleted from the storage array and all data on the virtual disk is permanently lost.

Deleting a Secondary Virtual Disk in a Replicated Pair From a Storage Array

1 Stop all I/O activity on the secondary virtual disk and unmount any file systems on the secondary virtual disk.

2 In the **AMW** of the storage array that contains the secondary virtual disk, select the **Logical** tab.

3 In the **Logical** pane, right-click the secondary virtual disk, then select **Delete**.

   The **Delete Virtual Disks** dialog is displayed.

4 Select one or more virtual disks to delete, then click **Delete**.

   The **Confirm Delete Virtual Disk(s)** message is displayed.

5 In the confirmation text box, type **Yes**. Then, click **OK**.

   The **Delete Virtual Disks** - **Progress** bar is displayed.

6 When the deletion is complete, click **OK**.

   The replication relationship is removed and the remote virtual disk replication is destroyed.

NOTE: The secondary virtual disk is deleted from the storage array and all data on the virtual disk is permanently lost.

Removing a Remote Replication From a Storage Array

Removing a remote virtual disk replication from a storage array returns both the primary virtual disk and the secondary virtual disk to regular standard virtual disks. After the removal is complete, normal I/O operations continue on the former primary virtual disk; the former secondary virtual disk also becomes available for normal I/O operations. Both virtual disks become read-write enabled. Unless one of disks is deleted, a replication relationship between the two virtual disks can be re-created.

NOTE: No data on either virtual disk is deleted.

CAUTION: Do not remove a replication relationship to back up a replicated virtual disk. To back up either the primary virtual disk or the secondary virtual disk, suspend the remote virtual disk replication so that the replication relationship is not broken.

To remove remote replication from a storage array:

1 In the **AMW** of the storage array containing the primary virtual disk, select the **Logical** tab.

2 In the **Logical** pane, right-click the primary virtual disk of a replicated pair, then select **Remove Replication Relationship**.

The **Remove Replication Relationship** dialog is displayed. The Replicated pairs table shows all replicated pairs in both the local and remote storage array.

3 To remove a replication relationship, select one or more replicated pairs. To select all replicated pairs, click **Select All**.

4 Click **Remove**.

The **Remove Replication Relationship** - **Confirmation** message is displayed.

5 Click **Yes**.

The **Remove Replicated Pair** - **Progress** bar shows the progress of the removal process.

Deactivating the Remote Replication on the Storage Array

Before you can permanently disable the Remote Replication premium feature you must deactivate it on the storage array, see "Disabling the Remote Replication Premium Feature" on page 217. Deactivating the Remote Replication premium feature on this storage array does not affect remote virtual disk replications or the Remote Replication premium features on other storage arrays. However, another storage array cannot use this storage array as a remote storage array for creating a separate remote virtual disk replication.

To deactivate remote replication in a storage array:

1 In the **Array Management Window,** select **Storage Array**→ **Remote Replication**→ **Deactivate**.

The **Deactivate Remote Replication** confirmation message is displayed.

2 Click **Yes**.

The Remote Replication premium feature is deactivated and the two replication repository virtual disks are deleted from the storage array.

Disabling the Remote Replication Premium Feature

Before disabling the Remote Replication premium feature, all remote virtual disk replications must are removed on the storage array. See "Disabling the Remote Replication Premium Feature" on page 217.

To disable the remote replication premium feature:

1   In the **Array Management Window**, select **Storage Array** > **Premium Features**.

    The **Premium Features and Feature Pack Information** screen is displayed.

2   Select **Remote Replication** in the Premium Features window, then click **Disable**.

    A warning message is displayed.

3   If you are sure you want to disable the Remote Replication feature, click **Yes** to confirm.

    The Remote Replication premium feature is disabled.

14

Configuration: Device Mapper Multipath for Linux

Overview

The MD3600f Series storage array uses a Linux operating system software framework, known as Device Mapper (DM), to enable multipath capabilities on Linux Host Servers. The DM multipath functionality is provided by a combination of drivers and utilities. This chapter describes how to use those utilities to complete the process of enabling an MD3600f Series storage array on a Linux system.

NOTE: The Device Mapper technology replaces an earlier, proprietary technology, known as MPP. MPP was used to enable multipathing for a previous MD generation, the MD3000 Series of storage arrays.

NOTE: The required Device Mapper software components are installed on a Linux host server by running the MD3600f Series Resource media installation program on the server, and selecting either the Full or Host install option. For detailed installation procedures, see the *Dell PowerVault MD3600f and MD3620f storage arrays Deployment Guide* at **dell.com/support/manuals**.

Benefits of using DM Multipath include:

- Detects path failure and re-routes I/O to other available paths
- Revalidates failed paths after path restoration
- Utilizes multiple available paths to maximize performance
- Reconfigures path usage based on path states and error conditions
- Unifies multiple device nodes into a single logical multipath device node
- Indentifies a new multipathed LU and automatically configures a new multipath node
- Provides device name persistency for DM devices under /dev/mapper/

Using DM Multipathing Devices

NOTE: Using or modifying any nodes other than the multipathing device nodes can result in array or file system problems, including loss of communication with the array and corruption of the file system. Avoid accessing any device other than the multipathing device.

NOTE: After creating a partition on a multipathing device, all IO operations, including filesystem creation, raw I/O and file system I/O, must be done through the partition node, not the multipathing device nodes.

Prerequisites

The following tasks must be completed before proceeding. For more information about step 1and step 2, see the *MD3600f and MD3620f Storage Arrays Deployment Guide* at **dell.com/support/manuals**. For more information about step 3, see "Creating Virtual Disks" on page 104.

1 **Install the host software from MD3600f Series resource media**—Insert the resource media to start the installation of Modular Disk Storage Manager (MDSM).

   NOTE: Red Hat install of 5.x needs a remount of the DVD media to make contents executable.

2 **Reboot when prompted by the install program**—The installation program prompts for and needs a reboot at completion of install.

3 **Create and map virtual disks using MDSM**—After configuring the arrays run the MDSM to create and map virtual disks.

Using the MDSM software:

1 Map the host server to the MD3600f Series storage array.

2 Create the Virtual Disks.

3 Map newly created arrays to your host server.

   NOTE: Any array configured with MDCU automatically adds to the list of Devices in the PowerVault Modular Disk Storage Manager Enterprise Management Window (EMW).

Device Mapper Configuration Steps

To complete the DM multipathing configuration and make the storage available to the Linux host server:

1 Scan for virtual disks. See "Scan for Newly Added Virtual Disks" on page 221.

2 Display the multipath device topology. See "Display the Multipath Device Topology Using the Multipath Command" on page 222.

3 Create a partition on a multipath device node. See "Create a New fdisk Partition on a Multipath Device Node" on page 223.

4 Add a partition to DM. See "Add a New Partition to Device Mapper" on page 223.

5 Create a file system on a DM partition. See "Create a File System on a Device Mapper Partition" on page 224.

6 Mount a DM partition. See "Mount a Device Mapper Partition" on page 224.

✎ **NOTE:** In the following command descriptions <x> is used to indicate where a substitution must be made. On RHEL systems <x> is the number assigned to the device. On SLES systems <x> is the letter(s) assigned to the device.

Scan for Newly Added Virtual Disks

The rescan_dm_devs command scans the host server system looking for existing and newly added virtual disks mapped to the host server.

```
# rescan_dm_devs
```

If an array virtual disk (VD) is mapped to the host server at a later time, the rescan_dm_devices command must be run again to make the VD a visible LUN to the operating system.

Display the Multipath Device Topology Using the Multipath Command

The multipath command adds newly scanned and mapped virtual disks to the Device Mapper tables and creates entries for them in the **/dev/mapper** directory on the host server. These devices are the same as any other block devices in the host.

To list all the multipath devices, run the following command:

```
# multipath –ll
```

The output must be similar to this example, which shows the output for one mapped virtual disk.

```
mpath1 (3600a0b80005ab177000017544a8d6b92) dm-0 DELL,
MD36xxf
```

```
[size=5.0G][features=3 queue_if_no_path
pg_init_retries 50][hwhandler=1 rdac][rw]
```

```
\_ round-robin 0 [prio=6][active]
 \_ 5:0:0:0   sdc  8:32   [active][ready]
\_ round-robin 0 [prio=1][enabled]
 \_ 4:0:0:0   sdb  8:16   [active][ghost]
```

where:

`mpath1` is the name of the virtual device created by device mapper. It is located in the **/dev/mapper** directory.

`DELL` is the vendor of the device

`MD36xxf` is the model of the device

`Sdc` is the physical path to the owning controller for the device

`Sdb` is the physical path to the non-owning controller for the device

The following is an example of SLES output:

```
mpathb(360080e500017b2f80000c6ca4a1d4ab8) dm-21
DELL,MD36xxf
```

```
[size=1.0G][features=3 queue_if_no_path
pg_init_retries 50][hwhandler=1 rdac][rw]
```

```
\_ round-robin 0 [prio=6][active]
 \_ 4:0:0:22 sdx  65:112 [active][ready]
```

```
\_ round-robin 0 [prio=1][enabled]
 \_ 6:0:0:22 sdcl 69:144 [active][ghost
```

where:

mpathb is the name of the virtual device created by device mapper. It is located in the /dev/mapper directory.

DELL is the vendor of the device

MD36xxf is the model of the device

Sdx is the physical path to the owning controller for the device

Sdcl is the physical path to the non-owning controller for the device

Create a New fdisk Partition on a Multipath Device Node

The fdisk command allows you to create a partition space for a file system on the newly scanned and mapped virtual disks that are presented to Device Mapper.

To create a partition with the multipathing device nodes /dev/mapper/mpath< x>, for example, use the following command, where mpath< x> is the multipathing device node on which you want to create the partition:

```
# fdisk /dev/mapper/mpath<x>
```

📝 **NOTE:** The <x> value is an alphanumeric operating system dependent format. The corresponding value for mapped virtual disks can be seen using the previously run multipath command. Please see your operating system documentation for additional information on fdisk usage.

Add a New Partition to Device Mapper

The kpartx command adds the newly fdisk created partition to the Device Mapper list of usable partitions. See examples below, where mpath< x> is the device node on which the partition was created.

```
# kpartx -a /dev/mapper/mpath<x>
```

If successful, the command does not show an output. To verify success and view exact partition naming, the user can use these commands to see the full partition names assigned.

```
# cd /dev/mapper
```

```
# ls
```

The following examples provide information about the general mapping formats:

On RHEL hosts, a partition node has the format /dev/mapper/mpath<x>p<y>, where <x> is the alphabetic number for the multipathing device, <y> is the partition number for this device.

On SLES 11.x hosts, a partition node has the format /dev/mapper/mpath<x>-part<y>, where <x> is letter(s) assigned to the multipathing device and <y> is the partition number.

On SLES 10.3 hosts, a partition node has the format /dev/mapper/mpath<x>_part<y>, where <x> is the letter(s) assigned to the multipathing device and <y> is the partition number.

> **NOTE:** After creating a partition on a multipathing device, all I/O operations, including file system creation, raw I/O and file system I/O, must be done through the partition node, and not through the multipathing device nodes.

Create a File System on a Device Mapper Partition

Use the standard `mkfs` command to create the file system on the newly created Device Mapper partition

For example:

```
# mkfs –t <filesystem type> /dev/mapper/<partition
node>
```

where <partition node> is the partition on which the file system is created.

Mount a Device Mapper Partition

Use the standard `mount` command to mount the Device Mapper partition:

```
# mount /dev/mapper/<partition_node> <mounting_point>
```

Ready for Use

The newly created virtual disks created on the MD3600f Series array are now setup and ready to be used. Future reboots automatically finds multipathing devices along with their partitions.

Blacklist Local Drive in Multi-path Driver

If your multipath drivers are connecting to storage area networks (SANs), it may be useful to be able to exclude or "blacklist" certain devices in your `/etc/multipath.conf` file. Blacklisting prevents the multipath driver from attempting to use those local devices.

To blacklist a local drive or device:

1 Run the `multipath -l` command to determine the local drive or device WWID (World-Wide Identifier) or vendor/model string.

2 Edit the **/etc/multipath.conf** file as follows:

```
blacklist {

wwid    drive_wwid

…

}
```

or

```
blacklist {

 device {

      vendor  vendor_string

     model model_string

 };
```

📖 **NOTE:** RedHat version 6.0 and 6.1 users must rebuild the **initramfs** root file image to include the updated configuration file by running the `#dracut -force` command.

3 Reboot the host.

Important Information About Special Partitions

When using Device Mapper with the MD3600f Series array, all physical disks are assigned a disk device node. This includes a special device type used for in-band management of the MD3600f Series array, known as the Access Disk or Universal Xport device.

⚠ **CAUTION: Certain commands, such as lsscsi, displays one or more instances of Universal Xport devices. These device nodes must never be accessed, mounted, or used in any way. Doing so could cause loss of communication to the storage array and possibly cause serious damage to the storage array, potential making data stored on the array inaccessible.**

Only multipathing device nodes and partition nodes created using the directions provided above must be mounted or in any way accessed by the host system or its users.

Table 14-1.   Device Mapper Commands

Command	Descriptions
multipath -h	Prints usage information.
multipath -ll	Displays the current multipath topology using all available information (sysfs, the device mapper, path checkers, and so on).
multipath	Re-aggregates multipathing device with simplified output.
multipath -f < multipath_dev_node>	Flushes out Device Mapper for the specified multipathing device. Used if the underlying physical devices are deleted/unmapped.
multipath -F	Flushes out all unused multipathing device maps.
rescan_dm_devs	Dell provided script. Forces a rescan of the host's SCSI bus and aggregates multipathing devices as needed. For use when: • LUNs are dynamically mapped to the hosts. • New targets are added to the host. • Failback of the storage array is required.

Limitations and Known Issues

- In certain error conditions with the no_path_retry or the queue_if_no_path feature is set, applications may hang. To overcome these conditions the you must enter the following command for each affected multipath device:

  ```
  dmsetup message [device] 0 "fail_if_no_path"
  ```

  where `[device]` is the multipath device name (for example. mpath2; do not specify the path)

- I/O may hang when a Device Mapper device is deleted before the volume is unmounted.

- If the scsi_dh_rdac module is not included in initrd, slower device discovery may be seen and the syslog may become populated with buffer I/O error messages.

- I/O may hang if the host server or storage array is rebooted while I/O is active. All I/O to the storage array must be stopped before shutting down or rebooting the host server or storage array.

- After a failed path is restored on an MD3600i Series array, failback does not automatically occur because the driver cannot auto-detect devices without a forced rescan. Run the command `rescan_dm_devs` to force a rescan of the host server. This restores the failed paths enabling failback to occur.

- Failback can be slow when the host system is experiencing heavy I/O. The problem is exacerbated if the host server is also experiencing very high processor utilization.

- The Device Mapper Multipath service can be slow when the host system is experiencing heavy I/O. The problem is exacerbated if the host server also experiences very high processor utilization.

- If the root disk is not blacklisted in the multipath.conf file, a multipathing node may be created for the root disk. The command `multipath -ll` lists vendor/product ID which can help identify this issue.

Troubleshooting

Question	Answer
How can I check if **multipathd** is running?	Run the following command `/etc/init.d/multipathd status`
Why does the `multipath -ll` command output not show any devices?	First verify if the devices are discovered or not. The command `#cat /proc/scsi/scsi` displays all the devices that are already discovered. Then verify the multipath.conf to ensure that it is updated with proper settings. After this, run `multipath`. Then run `multipath -ll`, the new devices must be displayed.
Why is a newly-mapped LUN not assigned a multipathing device node?	Run "rescan_dm_devs" in any directory. This must bring up the devices.
I have no LUNs mapped before. Then I map a few LUNs. After running `rescan-scsi-bus.sh`, LUN 0 does not show up.	Run `rescan_dm_devs` instead of `rescan-scsi-bus` for LUN 0 reconfiguration.
I removed a LUN, but the multipathing mapping is still present.	The multipathing device is still there after you remove the LUNs. Run `multipath -f` <device node for the deleted LUN> to remove the multipathing mapping. For example, if device devices related with /dev/dm-1 are deleted, you must run `multipath -f /dev/dm-1` to remove /dev/dm-1 from DM mapping table. If multipathing daemon is stopped/restarted, run `multipath -F` to flush out all stale mappings.
Failback does not happen as expected with the array.	Sometimes the low level driver cannot auto-detect devices coming back with the array. Run `rescan_dm_devs` to do a rescan of the host server SCSI bus and re-aggregate devices at multipathing layer.

Management: Firmware Downloads

Downloading RAID Controller and NVSRAM Packages

A version number exists for each firmware file. The version number indicates whether the firmware is a major version or a minor version. You can use the Enterprise Management Window (EMW) to download and activate both the major firmware versions and the minor firmware versions. You can use the Array Management Window (AMW) to download and activate only the minor firmware versions.

> **NOTE:** Firmware versions are of the format aa.bb.cc.dd.
> Where, aa is the major firmware version and bb.cc.dd is the minor firmware version. Depending on which one changes, firmware can be updated from EMW and AMW or only EMW.

You can activate the files immediately or wait until a more convenient time. You may want to activate the firmware or NVSRAM files at a later time because of these reasons:

- **Time of day**—Activating the firmware and the NVSRAM can take a long time, so you can wait until I/O loads are lighter. The RAID controller modules goes offline briefly to load the new firmware.

- **Type of package**—You may want to test the new firmware on one storage array before loading the files onto other storage arrays.

The ability to download both files and activate them later depends on the type of RAID controller module in the storage array.

> **NOTE:** You can use the command line interface to download and activate the firmware to several storage arrays by using a script. For more information on the command line interface, see the *PowerVault Modular Disk Storage Manager online help* topics.

Downloading Both RAID Controller and NVSRAM Firmware

📝 **NOTE:** I/O to the array can continue while you are upgrading RAID controller and NVSRAM firmware.

📝 **NOTE:** Dell recommends that the firmware and NVSRAM be upgraded during a maintenance period when the array is not being used for I/O.

📝 **NOTE:** The RAID enclosure must contain at least two disk drives in order to update the firmware on the controller.

To download RAID controller and NVSRAM firmware in a single operation:

1 If you are using the EMW, go to step 9. If you are using the AMW, go to step 2.

2 Perform one of these actions:

   • Select **Advanced**→ **Maintenance**→ **Download**→ **RAID Controller Module Firmware**.

   • Select the **Support** tab, and click **Download Firmware**. In **Select download task**, select the **Download RAID controller module firmware** and click **OK**.

   📝 **NOTE:** The RAID Controller Module Firmware area and the NVSRAM area list the current firmware and the current NVSRAM versions respectively.

3 To locate the directory in which the file to download resides, click **Select File**.

4 In the **File Selection** area, select the file to download.

   By default, only the downloadable files that are compatible with the current storage array configuration are displayed.

   When you select a file in the **File Selection** area of the dialog, applicable attributes (if any) of the file are displayed in the **File Information** area. The attributes indicate the version of the file.

5 If you want to download an NVSRAM file with the firmware, select **Transfer NVSRAM file with RAID controller module firmware**, and click **Select File** next to **Selected NVSRAM file**.

6 To transfer the files to the RAID controller module without activating them, click **Transfer files but don't activate them (activate later)**.

7 Click **Transfer**.

Keep these guidelines in mind:

- If the **Transfer** button is inactive, ensure that you either select an NVSRAM file or deselect the **Transfer NVSRAM file with RAID controller module firmware** option.

- If the file selected is not valid or is not compatible with the current storage array configuration, the **File Selection Error** dialog is displayed. Click **OK** to close it, and choose a compatible firmware or NVSRAM file.

8 In the **Confirm Download** dialog, click **Yes**.

The download starts.

9 Perform one of these actions:

- Select **Tools→ Upgrade RAID Controller Module Firmware**.

- Select the **Setup** tab, and click **Upgrade RAID Controller Module Firmware**.

10 In the **Storage array** pane, select the storage array for which you want to upgrade the RAID controller module firmware or the NVSRAM.

You can select more than one storage array.

NOTE: The Details pane shows the details of only one storage array at a time. If you select more than one storage array in the Storage Array pane, the details of the storage arrays are not shown in the Details pane.

11 Click **Firmware** in the **Download** area.

If you select a storage array that cannot be upgraded, the **Firmware** button is disabled.

The **Download Firmware** dialog is displayed. The current firmware version and the NVSRAM version of the selected storage arrays is displayed.

NOTE: If you select the storage arrays with different RAID controller module types that cannot be updated with the same firmware or NVSRAM file and click Firmware, the Incompatible RAID Controller Modules dialog is displayed. Click OK to close the dialog and select the storage arrays with similar RAID controller module types.

12 To locate the directory in which the file to download resides, click **Browse** in the **Select files** area.

The **Select File** dialog is displayed.

13 Select the file to download.

14 Click **OK**.

15 If you want to download the NVSRAM file with the RAID controller module firmware, select **Download NVSRAM file with firmware** in the **Select files** area.

Attributes of the firmware file are displayed in the Firmware file information area. The attributes indicate the version of the firmware file.

Attributes of the NVSRAM file are displayed in the NVSRAM file information area. The attributes indicate the version of the NVSRAM file.

16 If you want to download the file and activate the firmware and NVSRAM later, select the **Transfer files but don't activate them (activate later)** check box.

    ✎ **NOTE:** If any of the selected storage arrays do not support downloading the files and activating the firmware or NVSRAM later, the **Transfer files but don't activate them (activate later)** check box is disabled.

17 Click **OK**.

The **Confirm Download** dialog is displayed.

18 Click **Yes**.

The download starts and a progress indicator is displayed in the Status column of the **Upgrade RAID Controller Module Firmware** window.

Downloading Only NVSRAM Firmware

You also can use the command line interface (CLI) to download and activate NVSRAM to several storage arrays. For more information, see the *PowerVault Modular Disk Storage Manager online help* topics.

To download only NVSRAM firmware:

1 To download the NVSRAM firmware from:

- EMW—Go to step 8.
- AMW—Go to step 10

2 Perform one of these actions:

- Select **Advanced**→ **Maintenance**→ **Download**→ **RAID Controller Module NVSRAM**

- Select the **Support** tab, and click **Download Firmware**. In **Select download task**, select **Download RAID controller module NVSRAM** and click **OK**.

An error message is displayed. Click **OK** to close it and select a compatible file.

3 To locate the directory in which the file to download resides, click **Select File**.

4 Select the file to download in the **File selection** area and click **OK**.

By default, only downloadable files that are compatible with the current storage array configuration are displayed.

When you select a file in the File selection area, applicable attributes (if any) of the file are displayed in the NVSRAM File information area. The attributes indicate the version of the NVSRAM file.

5 Click **Transfer**.

> ✍ **NOTE:** If the file selected is not valid or is not compatible with the current storage array configuration, the File Selection Error dialog is displayed. Click OK to close it, and choose a compatible NVSRAM file.

6 Click **Yes** in the **Confirm Download** dialog.

The download starts.

7 Perform one of these actions:

- Select **Tools**→ **Upgrade RAID Controller Module Firmware**.

- Select the **Setup** tab, and click **Upgrade RAID Controller Module Firmware.**

The **Upgrade RAID Controller Module Firmware** window is displayed.

The Storage array pane lists the storage arrays. The **Details** pane shows the details of the storage array that is selected in the Storage array pane.

8 In the **Storage array** pane, select the storage array for which you want to download the NVSRAM firmware.

You can select more than one storage array.

NOTE: The Details pane shows the details of only one storage array at a time. If you select more than one storage array in the Storage array pane, the details of the storage arrays are not shown in the Details pane.

9 Click **NVSRAM** in the **Download** area.

NOTE: If you select a storage array that cannot be upgraded, the NVSRAM button is disabled.

The **Download NVSRAM** dialog is displayed. The current firmware version and the NVSRAM version of the selected storage arrays are displayed.

NOTE: If you select the storage arrays with different RAID controller module types that cannot be updated with the same NVSRAM file and click NVSRAM, the **Incompatible RAID Controller Modules** dialog is displayed. Click **OK** to close the dialog and select the storage arrays with similar RAID controller module types.

10 To locate the directory in which the NVSRAM file to download resides, click **Browse** in the **Select file** area.

The **Select File** dialog is displayed.

11 Select the file to download and click **OK**.

Any attributes of the NVSRAM file are displayed in the NVSRAM file information area. The attributes indicate the version of the NVSRAM file.

12 Click **OK**.

The **Confirm Download** dialog is displayed.

13 Click **Yes**.

The download starts and a progress indicator is displayed in the Status column of the **Upgrade RAID Controller Module Firmware** window.

Downloading Physical Disk Firmware

⚠ **CAUTION: When updating physical disk firmware, you must stop all I/O activity to the array to prevent data loss.**

The physical disk firmware controls various features of the physical disk. The disk array controller (DAC) uses this type of firmware. Physical disk firmware stores information about the system configuration on an area of the physical disk called DACstore. DACstore and the physical disk firmware enable easier reconfiguration and migration of the physical disks. The physical disk firmware performs these functions:

- The physical disk firmware records the location of the physical disk in an expansion enclosure. If you take a physical disk out of an expansion enclosure, you must insert it back into the same physical disk slot, or the physical disk firmware cannot communicate with the RAID controller module or other storage array components.

- RAID configuration information is stored in the physical disk firmware and is used to communicate with other RAID components.

⚠ **CAUTION: Risk of application errors – Downloading the firmware could cause application errors.**

Keep these important guidelines in mind when you download firmware to avoid the risk of application errors:

- Downloading firmware incorrectly could result in damage to the physical disks or loss of data. Perform downloads only under the guidance of your Technical Support representative.

- Stop all I/O to the storage array before the download.

- Ensure that the firmware that you download to the physical disks are compatible with the physical disks that you select.

- Do not make any configuration changes to the storage array while downloading the firmware.

✐ **NOTE:** Downloads can take several minutes to complete. During a download, the **Download Physical Disk - Progress** dialog is displayed. Do not attempt another operation when the **Download Physical Disk - Progress** dialog is shown.

To download Physical Disk Firmware:

1   From the AMW, select **Advanced**→ **Maintenance**→
    Download→ **Physical Disk.**

    The **Download Physical Disk** - **Introduction** window is displayed.

2   Click **Next**.

    The **Download Physical Disk Firmware** - **Add Packages** window is displayed.

3   In the **Selected Packages** area, click **Add**. Navigate to the location of the
    packages and click **OK**.

    The selected package is added to the **Packages to be transferred** area.

4   Click **Next**.

    The **Download Physical Disk Firmware** - **Select Physical Disks** window
    is displayed.

5   In the **Compatible Physical Disks** tab, select the appropriate physical
    disks or **Select all** the physical disks.

    The **Confirm Download** dialog is displayed.

6   Type **yes** and click **OK**.

    The **Download Physical Disk Firmware** - **Progress** window is displays the
    progress of physical disk firmware download.

7   After the firmware download is complete, click **Close**.

For more information, see the *PowerVault Modular Disk Storage Manager
online help* topics.

Downloading MD1200 Series Expansion Module EMM Firmware

> **NOTE:** Due to a limitation with Linux, expansion enclosure EMM firmware updates must be performed using out-of-band management only. Failure to do so may result in the host server becoming unresponsive, and it may require a reboot.

You can transfer a downloadable firmware file to the expansion enclosure EMM in the expansion enclosures attached to the storage array.

> **CAUTION: Risk of possible loss of data or risk of damage to the storage array –** Downloading the expansion enclosure EMM firmware incorrectly could result in loss of data or damage to the storage array. Perform downloads only under the guidance of your Technical Support representative.

> **CAUTION: Risk of making expansion enclosure EMM unusable –** Do not make any configuration changes to the storage array while downloading expansion enclosure EMM firmware. Doing so could cause the firmware download to fail and make the selected expansion enclosure unusable.

To download the EMM firmware:

1. Perform one of these actions:
   - In the AMW, select **Advanced**→ **Maintenance**→ **Download**→ **EMM Firmware**.
   - Select the **Support** tab, and click **Download Firmware**. In the dialog that is displayed, select the **EMM firmware**, and click **OK**.

   The **Download Environmental (EMM) Firmware** dialog is displayed.

2. In the **Select enclosures** area, either select each expansion enclosure to which you want to download firmware, or select **Select All** to select all of the expansion enclosures in the storage array.

   Each selected expansion enclosure must have the same product ID.

3. Click **Select File**.

   The **Select Environmental (EMM) Card Firmware File** dialog is displayed.

4. Select the file to download and click **OK**.

5. Click **Start**.

6. Click **Yes** to continue with the firmware download.

Monitor the progress and completion status of the download to the expansion enclosures. The progress and status of each expansion enclosure that is participating in the download is displayed in the Status column of the Select enclosures table.

7 Perform one of these actions depending on whether the download succeeded:

- The download succeeded – The statuses of all the expansion enclosures show Complete. You can close the Download environmental (EMM) Card Firmware dialog by clicking **Close**. The expansion enclosure EMM cards are now operating with the new firmware.

- The download failed – The status of one expansion enclosure shows Failed and the remainder of the expansion enclosures show Canceled. Ensure that the new firmware file is compatible before attempting another firmware download.

Self-Monitoring Analysis and Reporting Technology (SMART)

Self-Monitoring Analysis and Reporting Technology (SMART) monitors the internal performance of all physical disk components to detect faults indicating the potential for physical disk failure. SMART uses this information to report whether failure is imminent so that a physical disk can be replaced before failure occurs. The RAID controller monitors all attached drives and notifies users when a predicted failure is reported by a physical disk.

Media Errors and Unreadable Sectors

If the RAID controller detects a media error while accessing data from a physical disk that is a member of a disk group with a redundant RAID level (RAID 1, RAID 5 or RAID 10), the controller tries to recover the data from peer disks in the disk group and uses recovered data to correct the error. If the controller encounters an error while accessing a peer disk, it is unable to recover the data and affected sectors are added to the unreadable sector log maintained by the controller. Other conditions under which sectors are added to the unreadable sector log include:

- A media error is encountered when trying to access a physical disk that is a member of a non-redundant disk group (RAID 0 or degraded RAID 1, RAID 5 or RAID 10).

- An error is encountered on source disks during rebuild.

    NOTE: Data on an unreadable sector is no longer accessible.

16

Management: Installing Array Components

Recommended Tools

You may need the following items to perform the procedures in this section:

- Key to the system keylock
- #2 Phillips screwdriver
- Wrist grounding strap

Front Bezel (Optional)

Removing the Front Bezel

1 Using the system key, unlock the front bezel (if locked).

2 Lift up the release latch next to the keylock.

3 Rotate the left end of the bezel away from the front panel.

4 Unhook the right end of the bezel and pull the bezel away from the system.

Figure 16-1. Removing and Installing the Front Bezel

1	bezel	2	keylock
3	release latch	4	hinge tab

Installing the Front Bezel

1 Hook the right end of the bezel onto the chassis.

2 Fit the free end of the bezel onto the system.

3 Secure the bezel with the keylock. See Figure 16-1.

Hard Drives

SAFETY: Models AMT, E03J, and E04J

Models AMT, E03J, and E04J are intended for installation only in restricted access locations as defined in cl 1.2.7.3 of IEC 60950-1:2005.

Depending on your configuration, your array either supports up to twenty four 2.5" SAS hard drives or up to twelve 3.5" SAS hard drives in internal drive bays. Hard drives are connected to a backplane through hard-drive carriers and can be configured as hot-swappable.

Removing a Hard-Drive Blank

⚠ CAUTION: To maintain proper system cooling, all empty hard-drive bays must have drive blanks installed.

1 If installed, remove the front bezel. See "Removing the Front Bezel" on page 242.

2 Press the release tab and slide the hard-drive blank out until it is free of the drive bay. See Figure 16-2 for PowerVault MD3600f and Figure 16-3 for PowerVault MD3620f.

Figure 16-2. Removing and Installing a 3.5" Hard-Drive Blank (MD3600f Only)

| 1 | hard-drive blank | 2 | release tab |

| 1 | hard-drive blank | 2 | release tab |

Installing a Hard-Drive Blank

1   If installed, remove the front bezel. See "Removing the Front Bezel" on page 242.

2   Insert the hard-drive blank into the drive bay until the blank is fully seated.

3   Close the handle to lock the blank in place.

4   If applicable, replace the front bezel. See "Installing the Front Bezel" on page 242.

Removing a Hard Drive

⚠ CAUTION: Many repairs may only be done by a certified service technician. You must only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that came with the product.

1   If installed, remove the front bezel. See "Removing the Front Bezel" on page 242.

2   From the Modular Disk Storage Manager (MDSM) software, prepare the drive for removal. Wait until the hard-drive indicators on the drive carrier signal that the drive can be removed safely. See your controller documentation for information about hot-swap drive removal.

If the drive is online, the green activity/fault indicator flashes as the drive is powered down. When the drive indicators are off, the drive is ready for removal.

3   Press the release button to open the drive carrier release handle. See Figure 16-4.

4   Slide the hard drive out until it is free of the drive bay.

⚠ **CAUTION: To maintain proper system cooling, all empty hard-drive bays must have drive blanks installed.**

5   Insert a drive blank in the empty drive bay. See "Installing a Hard-Drive Blank" on page 244.

6   If applicable, replace the front bezel. See "Installing the Front Bezel" on page 242.

Figure 16-4.   Removing and Installing a Hard Drive

1 release button	2 hard-drive carrier handle

Installing a Hard Drive

⚠ **CAUTION: Many repairs may only be done by a certified service technician. You must only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that came with the product.**

⚠ **CAUTION: Use only hard drives that are tested and approved for use with the MD3600f Series.**

⚠ **CAUTION: When installing a hard drive, ensure that the adjacent drives are fully installed. Inserting a hard-drive carrier and attempting to lock its handle next to a partially installed carrier can damage the partially installed carrier's shield spring and make it unusable.**

1 If applicable, remove the front bezel. See "Removing the Front Bezel" on page 242.

2 If applicable, remove the drive blank from the bay. See "Removing a Hard-Drive Blank" on page 243.

3 Press the release button to open the drive carrier release handle.

4 Insert the hard-drive carrier into the drive bay until the carrier contacts the backplane.

5 Close the handle to lock the drive in place.

Removing a Hard Drive From a Hard-Drive Carrier

Remove the screws from the slide rails on the hard-drive carrier and separate the hard drive from the carrier. See Figure 16-5 for PowerVault MD3600f and Figure 16-6 for PowerVault MD3620f.

Figure 16-5.   Removing and Installing a Hard Drive Into a 3.5" Hard-Drive Carrier

1	screws (4)	2	hard drive
3	SAS screw hole	4	hard-drive carrier

Figure 16-6.   Removing and Installing a Hard Drive Into a 2.5" Hard-Drive Carrier

1	screws (4)	2	hard drive
3	SAS screw hole	4	hard-drive carrier

Installing a Hard Drive Into a Hard-Drive Carrier

1   Insert the hard drive into the hard-drive carrier with the connector end of the drive at the back. See Figure 16-5.

2   Align the screw holes on the hard drive with the back set of holes on the hard-drive carrier.

    When aligned correctly, the back of the hard drive is flush with the back of the hard-drive carrier.

3   Attach the four screws to secure the hard drive to the hard-drive carrier.

RAID Controller Module

An MD3600f series storage array supports single as well as dual RAID controller configurations. If only one RAID controller module is installed in your array, it must be installed in slot 0. You must install the RAID controller module blank in slot 1.

⚠ **CAUTION: RAID controller modules can be removed and installed without turning off the array. It is recommended that you do not remove the RAID controller module while data is being transferred. Replacing or installing a RAID controller module that is connected to a host server causes it to loose communication with the array and may require a reboot of the host server.**

Removing a RAID Controller Module Blank

⚠ **CAUTION: To maintain proper system cooling, you must install a RAID controller module blank in the empty slot.**

1   Turn off the array and host server.

2   Disconnect all the power cables connected to the array.

3   To remove the RAID controller module blank, press down on the release latch and pull the blank away from the array. See Figure 16-7.

4   Install RAID controller modules in slot 0 and 1. See "Installing a RAID Controller Module" on page 252.

5 Connect all the power cables to the array.

6 Turn on the array and the host server.

Figure 16-7.   Removing and Installing a RAID Controller Module Blank

| 1 | release latch | 2 | RAID controller module blank |

Installing a RAID Controller Module Blank

To install a RAID controller module blank:

1 Align the blank with the RAID controller module bay

2 Insert the blank into the chassis until it clicks into place.

Removing a RAID Controller Module

⚠ **CAUTION: Many repairs may only be done by a certified service technician. You must only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that came with the product.**

⚠ **CAUTION: If your configuration uses fiber optic cables, remember that they are fragile. Bending, twisting, folding, or pinching fiber optic cables can cause damage to the cables, degraded performance, or loss of data. To prevent damage, do not twist, fold, pinch, or step on the cables. Do not bend the cables in less than a 5 cm (2") radius.**

1   If small form-factor pluggable (SFP) transceivers are present, record the ports in which they are installed, and remove them.

2   Take the desired RAID controller module offline.

3   Disconnect the cables connected to the RAID controller module.

    If fiber optic cables are present, it may be necessary to partially remove the RAID controller module using the two release levers. Opening the release levers makes it easier to press down the fiber optic cable release tab. If the system is still running, do not disturb the second RAID controller module.

4   Push down on the release tab and pull the release lever away from the chassis. See Figure 16-8.

5   Grasp the release lever and pull the module away from the chassis.

    📝 **NOTE:** To avoid damage to the sensitive EMI contacts on the RAID controller module, do not stack RAID controller modules.

Figure 16-8. Removing and Installing a RAID Controller Module

1	RAID controller module	2	release tab
3	release lever		

Installing a RAID Controller Module

⚠ **WARNING: Fiber optic cables are fragile. Bending, twisting, folding, or pinching fiber optic cables can cause damage to the cables, degraded performance, or loss of data. To prevent damage, do not twist, fold, pinch, or step on the cables. Do not bend the cables in less than a 5 cm (2") radius.**

⚠ **CAUTION: Many repairs may only be done by a certified service technician. You must only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that came with the product.**

   1 Insert the RAID controller module into the RAID controller module bay until it seats into place.

   2 Push the release lever toward the chassis until it clicks into place.

3 Connect all the cables to the RAID controller module. If SFP transceivers and fiber optic cables were present, reinstall them in the correct ports.

4 If applicable, update the firmware for the RAID controller module. For information about the latest firmware, see **dell.com/support**.

Opening the RAID Controller Module

⚠ **CAUTION: Many repairs may only be done by a certified service technician. You must only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that came with the product.**

1 Turn off the array and host server.

2 Disconnect all the power cables connected to the array.

3 Remove the RAID controller module. See "Removing a RAID Controller Module Blank" on page 249.

4 Remove the screws from the sides of the RAID controller module. See Figure 16-9.

5 While pressing the indent, slide the cover in the direction of the arrow and lift it away from the RAID controller module. See Figure 16-9.

Figure 16-9.   Opening and Closing the RAID Controller Module

1	screws (2)	2	RAID controller module
3	indent		

Closing the RAID Controller Module

⚠ **CAUTION: Many repairs may only be done by a certified service technician. You must only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that came with the product.**

1 Place the cover onto the RAID controller module and offset it slightly towards the back, so that the hooks on the cover fit over the corresponding slots on the RAID controller module.

2 Slide the cover toward the front till it snaps into place. See Figure 16-9.

3 Replace the screws on the RAID controller module. See Figure 16-9.

4 Connect all the cables to the array. If SFP transceivers and fiber optic cables were previously in place, reinstall them in the correct ports.

5 Turn on the array and the host server.

RAID Controller Module Backup Battery Unit

Removing the RAID Controller Module Backup Battery Unit

⚠ **CAUTION: Many repairs may only be done by a certified service technician. You must only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that came with the product.**

1 Turn off the array and host server.

2 Disconnect all the cables connected to the array.

3 Remove the RAID controller module. See "Removing a RAID Controller Module" on page 251.

4 Open the RAID controller module. See "Opening the RAID Controller Module" on page 253.

5 Loosen the screw that secures the backup battery unit to the RAID controller module. See Figure 16-10.

6 Slide the backup battery unit in the direction of the arrow and lift it out of the RAID controller module. See Figure 16-10.

Figure 16-10.   Removing and Installing the RAID Controller Module Backup Battery Unit

| 1 | backup battery unit | 2 | screw |

Installing the RAID Controller Module Backup Battery Unit

⚠ CAUTION: Many repairs may only be done by a certified service technician. You must only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that came with the product.

1 Align the backup battery unit with the slots on the RAID controller module.

2 Slide the backup battery unit toward the connector on the RAID controller module.

3 Tighten the screw that secures the backup battery unit to the RAID controller module.

4 Close the RAID controller module. See "Closing the RAID Controller Module" on page 254.

5 Replace the RAID controller module. See "Installing a RAID Controller Module" on page 252.

6 Connect all the cables to the array.

7 Turn on the array and the host server.

Power Supply/Cooling Fan Module

NOTE: Your storage array includes two integrated, hot-swappable power supply/cooling fan modules.

The array supports two hot-swappable power supply/cooling fan modules. While the array can operate temporarily with one module, both the modules must be present for proper system cooling.

CAUTION: A single power supply/cooling fan module can be removed from a powered-on array for a maximum period of 5 minutes. Beyond that time, the array may automatically shut down to prevent damage.

Removing a Power Supply/Cooling Fan Module

CAUTION: Many repairs may only be done by a certified service technician. You must only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that came with the product.

NOTE: If you remove a fully functioning power supply/cooling fan module, the fan speed in the remaining module increases significantly to provide adequate cooling. The fan speed decreases gradually when a new power supply/cooling fan module is installed.

1   Turn off the power supply/cooling fan module.

2   Disconnect the power cable from the power source.

3   Remove the straps that secure the power cable and then disconnect the power cable from the power supply/cooling fan module.

WARNING: The power supply/cooling fan modules are heavy. Use both hands while removing the module.

4   Press the release tab and pull the power supply out of the chassis.

Figure 16-11.   Removing and Installing a Power Supply/Cooling Fan Module

1	release tab	2	power supply
3	power supply handle		

Installing a Power Supply/Cooling Fan Module

⚠ **CAUTION: Many repairs may only be done by a certified service technician. You must only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that came with the product.**

1 Slide the power supply/cooling fan module into the chassis until it is fully seated and the release tab clicks into place. See Figure 16-11.

2 Connect the power cable to the power supply/cooling fan module and plug the cable into a power outlet.

3 Secure the power cable using the strap. See Figure 16-12.

Figure 16-12.   Securing the Power Cable

1   restraining strap

⚠ **CAUTION: When connecting the power cable, secure the cable with the velcro strap.**

🖉 **NOTE:** If the array is powered on, all the power supply LEDs remain off until the AC power cable is connected to the power supply/cooling fan module and the power switch is turned on.

4 Turn on the power supply/cooling fan module.

Control Panel

Removing the Control Panel

1 Turn off the array and host server.

2 Disconnect all the power cables connected to the array.

3 Remove the hard drives from:

– slots 0 to 2 in *PowerVault MD3600f*

– slots 0 to 5 in *PowerVault MD3620f*

See "Removing a Hard Drive" on page 244.

🖉 **NOTE:** Mark each hard drive with it's slot position as you remove it.

4 Slide the control panel out of the chassis after:

 – Pushing the release tab toward the front of the array in *PowerVault MD3600f*. See Figure 16-13.

 – Pulling the release pin toward the front of the array in *PowerVault MD3620f*. See Figure 16-14.

Figure 16-13.   Removing and Installing the Control Panel-PowerVault MD3600f

1	control panel	2	release tab

Figure 16-14.   Removing and Installing the Control Panel-PowerVault MD3620f

| 1 | release pin | 2 | control panel |

Installing the Control Panel

1   Align the control panel with the slot on the array.

2   Slide the control panel into the array until:

  –   The release tab clicks into place in *PowerVault MD3600f*. See Figure 16-13.

  –   The release pin clicks into place in *PowerVault MD3620f*. See Figure 16-14.

3   Replace the hard drives in their respective slots. See "Installing a Hard Drive" on page 246.

4   Connect all the power cables to the array.

5   Turn on the array and the host server.

Backplane

⚠ **CAUTION: Many repairs may only be done by a certified service technician. You must only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that came with the product.**

Removing the Backplane

1 Turn off the array and disconnect it from the electrical outlet.

2 Disconnect all the cables connected to the array.

3 Remove the hard drives. See "Removing a Hard Drive" on page 244.

4 Remove the RAID controller modules. See "Removing a RAID Controller Module" on page 251.

5 Remove the power supply/cooling fan modules. See "Removing a Power Supply/Cooling Fan Module" on page 257.

6 Remove the control panel. See "Removing the Control Panel" on page 259.

7 Remove the screws that secure the RAID controller module/power supply cage to the chassis.

8 Grasp the cage removal ring at the bottom center of the array and pull the RAID controller module/power supply cage toward the back of the chassis. See Figure 16-15.

9 Lift the RAID controller module/power supply cage away from the chassis. See Figure 16-15.

10 Loosen the captive screw that secures the backplane to the chassis. See Figure 16-16 for PowerVault MD3600f or Figure 16-17 for PowerVault MD3620f.

11 Remove the screws that secure the backplane and pull the backplane out of the array. See Figure 16-16 for PowerVault MD3600f or Figure 16-17 for PowerVault MD3620f.

Figure 16-15. Removing and Installing the RAID Controller Module/Power Supply Cage

| 1 | screws (6) | 2 | RAID controller module/power supply cage |

Figure 16-16.   Removing and Installing the Backplane-PowerVault MD3600f

1	screws (5)	2	backplane
3	captive screw		

Figure 16-17.   Removing and Installing the Backplane-PowerVault MD3620f

1	screws (4)	2	backplane
3	captive screw		

Installing the Backplane

1 Align the holes on the backplane with the holes on the array.

2 Tighten the captive screw to secure the backplane to the chassis. See Figure 16-16 for PowerVault MD3600f or Figure 16-17 for PowerVault MD3620f.

3 Replace the screws that secure the backplane to the chassis. See Figure 16-16 for PowerVault MD3600f or Figure 16-17 for PowerVault MD3620f.

4 Align the slots on the RAID controller module/power supply cage with the tabs on the chassis. See Figure 16-15.

5 Push the RAID controller module/power supply cage toward the front of the array.

6 Replace the screws that secure the RAID controller module/power supply cage to the chassis.

7 Replace the control panel. See "Installing the Control Panel" on page 261.

8 Replace the power supply/cooling fan modules. See "Installing a Power Supply/Cooling Fan Module" on page 258.

9 Replace the hard drives. See "Installing a Hard Drive" on page 246.

10 Connect all the cables to the array.

11 Turn on the array and the host server.

17

Management: Firmware Inventory

A storage array is made up of many components, which may include RAID controller modules, physical disks, and enclosure management modules (EMMs). Each of these components contains firmware. Some versions of firmware are dependent on other versions of firmware. To capture information about all of the firmware versions in the storage array, view the firmware inventory.

If the firmware inventory does not contain information for a particular storage array, the firmware inventory service is not available on that storage array.

You can also save the firmware inventory to a text file. You can then send the file to your Technical Support representative to detect any firmware mismatches.

Viewing the Firmware Inventory

1   Perform one of these actions based on whether you want to view the firmware information for one storage array or all storage arrays:

   •   One storage array—From the **Array Management Window**, select **Advanced**→ **Maintenance**→ **Firmware Inventory**.

   •   All storage arrays—From the **Enterprise Management Window**, select **Tools**→ **Firmware Inventory**.

2   To save the firmware inventory to a text file, click **Save As**.

3   In **Specify filename**, enter a name for the file to be saved. You may also specify another physical disk and directory if you want to save the file in a location other than the default.

   📝 **NOTE:** The suffix *.txt is added to the file name automatically if you do not specify a suffix for the file name.

4   Click **Save**.

   An ASCII text file that contains the firmware inventory is saved to the designated directory.

18

Management: System Interfaces

Microsoft Services

Virtual Disk Service

The Microsoft Virtual Disk Service (VDS) is a component of the Windows operating system. The VDS component utilizes third-party vendor specific software modules, known as providers, to access and configure third-party storage resources, such as MD3600f storage arrays. The VDS component exposes a set of application programming interfaces (APIs) that provides a single interface for managing disks and other storage hardware. The MD3600f Series VDS Provider enables Windows tools, including the Disk Manager, to access and configure storage array virtual disks.

The VDS Provider for the MD3600f Series arrays is available on the MD3600f Series resource media. For more information on VDS, see **microsoft.com**.

Volume Shadow-Copy Service

The Microsoft Volume Shadow-copy Service (VSS) is a component of the Microsoft Windows operating system. The VSS component utilizes third-part vendor specific software modules, known as providers, to access and utilize snapshot and disk copy functionality provided by third-party storage resources, such as MD3600f Series storage arrays. The combination of the VSS component and the VSS Provider, included on the MD3600f Resource media, enable the MD3600f Series arrays to be utilized by third-party and Windows backup and snapshot applications.

> **NOTE:** Virtual disks that is used as source virtual disks for VSS snapshots must have names no longer than 16 characters.

The VSS hardware provider uses the source virtual disk name as a prefix for the snapshot and repository virtual disk names. The resulting snapshot and repository names are too long if the source virtual disk name exceeds 16 characters.

VSS attaches to the service and uses it to coordinate the creation of snapshot virtual disks on the storage array. VSS-initiated snapshot virtual disks can be triggered through backup tools, known as requestors. The VSS Provider Configuration Tool offers the following configuration options:

- Snapshot Repository Virtual Disk Properties—This section contains a drop-down list for the RAID level and a field for entering source virtual disk capacity percentage for snapshot repositories.

- Snapshot Repository Virtual Disk Location—This section contains a list of preferences for the location of the snapshot repository virtual disk. These preferences are honored whenever conditions permit.

The Microsoft VSS installer service for storage provisioning is available on the MD3600f resource media in the **\windows\VDS_VSS** directory.

NOTE: When registering VSS during your Windows setup, the registration graphical user interface (GUI) prompts you to provide the name of your array because settings in the GUI are array-specific, not host-specific.

Storage management VSS hardware provider tips:

- The number of snapshot virtual disks that can be created using a single snapshot set varies with the I/O load on the RAID controller modules. Under little or no I/O load, the number of virtual disks in a snapshot set must be limited to 16. Under high I/O loads, the limit is 3.

- The snapshot virtual disks created in the storage management software are differential snapshots. Plex snapshots are not supported.

- Virtual disks to be used as source virtual disks for VSS snapshots must have names no longer than 16 characters. The VSS hardware provider uses the source virtual disk name as a prefix for the snapshot and repository virtual disk names. The resulting snapshot and repository names are too long if the source virtual disk name exceeds 16 characters.

NOTE: A volume is another term for virtual disk.

For more information on VDS and VSS, see **microsoft.com**.

Troubleshooting: Your Storage Array Software

Start–Up Routine

Look and listen during the array's start-up routine for the indications described in Table 19-1. For a description of the front- and back-panel indicators, see "Planning: About Your Storage Array" on page 25.

Table 19-1.   Start-Up Routine Indications

Look/Listen for	Action
Alert messages.	See your storage management documentation.
An unfamiliar constant scraping or grinding sound when you access a physical disk.	See "Getting Help" on page 309.

> **NOTE:** At least two physical disks must be installed in the array.

Device Health Conditions

When you open the Enterprise Management Window (EMW), the Dell PowerVault Modular Disk Storage Management software (MDSM) establishes communication with each managed storage array and determines the current storage array status. The current status is represented by icons next to the managed storage array.

The status icons shown in the Tree view in the EMW represent a summary status for each storage array. If a storage array has a status of Needs Attention or a status of Fixing, you must determine the condition that is causing this status before attempting any management actions. You can determine the condition causing the Needs Attention status or the Fixing status by selecting the storage array and launching its Array Management Window (AMW).

To launch the AMW, perform one of these actions:

- On the **Devices** tab, in either the Tree view or the Table view, double-click a storage array. Alternatively, you can right-click a storage array and select **Manage Storage Array** from the pop-up menu.

- On the **Setup** tab, select **Manage a Storage Array**.

After the AMW is displayed, select the **Physical** tab to see the components in the storage array. A component that has a problem is indicated by a status icon.

The status icons indicate the status of the components that comprise the storage array. Also, the Recovery Guru option provides a detailed explanation of the conditions and the applicable steps to remedy any Needs Attention status. For more information, see "Recovery Guru" on page 278.

For the status of a storage array, the icons shown in the following table are used in the Tree view, the Table view, and both the EMW Status Bar and the AMW Status Bar.

Table 19-2.   Status Icon

Status	Icon	Description
Optimal		Each component in the managed storage array is in the desired working condition.
Needs Attention		There is a problem with the managed storage array that requires your intervention to correct it.
Unresponsive		The storage management station cannot communicate with the storage array or one RAID controller module or both RAID controller modules in the storage array.
Fixing Status		A Needs Attention status is corrected and the managed storage array is currently transitioning to an Optimal state.
Unsupported		The node is currently not supported by this version of MDSM.
Software Unsupported		The storage array is running a level of software that is no longer supported by MDSM.

In the Table view, every managed storage array is listed once, regardless of the number of attachments it has in the Tree view. After the storage array is contacted by MDSM, an icon representing its hardware status is displayed. Hardware status can be Optimal, Needs Attention, or Fixing. If, however, all of the network management connections from the storage management station to the storage array shown in the Tree view are Unresponsive, the storage array status is represented as Unresponsive.

In the EMW Status Bar and the AMW Status Bar, the icons also have these behaviors:

- Hold the mouse over the icon in the EMW Status Bar and the AMW Status Bar to display a tooltip with a brief description of the status.

- The icons for the Needs Attention status and Unresponsive status are displayed in EMW Status Bar and the AMW Status Bar if there are discovered storage arrays with either condition.

The EMW Tree view has additional status icons that are shown in the following table.

Table 19-3. Additional Status Icons

Status	Icon	Description
Unsupported Alerts with a Needs Upgrade Status		Setting an alert on a s storage array with a Needs Upgrade status is not supported. In this case, the storage array shows both a Needs Upgrade status and an Unsupported Alerts icon in the Tree view. The Unsupported Alerts icon indicates that the storage array cannot be monitored.
Alert Set		If you installed the Event Monitor with MDSM, and if you have set alerts, the Alert Set icon is displayed next to the storage array status in the Tree view for which the alerts are set.
Setting an Alert at the Parent Node Level		You can set alerts at any of the nodes in the Tree view. Setting an alert at a parent node level, such as at a host level, sets alert fir any child nodes. If you set an alert at a parent node level and any of the in-band storage array child nodes have a Needs Upgrade status, the Alert Disables status icon is displayed next to the parent node in the tree view.

Table 19-3.   Additional Status Icons

Status	Icon	Description
Adding a Storage Array		The Contacting Storage Array icon is shown in the Tree view and Table view until the current status of each managed storage array is known.
		The Contacting Storage Array icon is shown in the EMW Status Bar and the AMW Status Bar and the tooltip shows Contacting Storage arrays.
		As each storage array is contacted, its current status is obtained and shown in the Tree view and Table view. The applicable statuses are the Optimal, Needs Attention, Fixing, or Unresponsive.
Adding a Storage Array OK		No problems were encountered while adding the storage array.
		MDSM software continues to check for any status change events.
Adding a Storage Array Error		Displayed only when an error occurs.

✎ **NOTE:** MDSM may take a few minutes to update a status change to Unresponsive or from Unresponsive. A status change from or to Unresponsive depends on the network link to the storage array. All other status change updates faster.

Storage Array Support Data

You can gather various types of inventory, status, and performance data that can help troubleshoot any problem with the storage array. All the files are compressed into a single archive in a zipped-file format. You can forward the archive file to your Technical Support representative for troubleshooting and further analysis.

To generate the support data report:

1   In the AMW, perform one of these actions:

   •   Select **Advanced**→ **Troubleshooting** → **Support Data** → **Collect**.

   •   Select the **Support** tab, and click **Gather Support Information**.

   The **Collect All Support Data** window is displayed.

2 Enter a name for the support data file in **Specify filename** or click **Browse** to navigate to a previously saved file to overwrite an existing file.

The suffix .zip is added automatically to the file if you do not specify a suffix for the file.

3 Enter the **Execution summary**.

4 Click **Start**.

After all of the support files are gathered, they are archived using the file name that you specified.

5 Click **OK**.

NOTE: If a support data operation is running, it must complete before another support data operation can begin. Concurrent collections are not supported and results in an error message.

Automatically Collect the Support Bundle Data

You can use the Collect Support Bundle option to automatically save a copy of the support bundle when the client monitor process detects a critical event.

During a critical event, the support bundle is saved to the local physical disk of the client system in the same area that is used for other recovery information. This information is not overwritten for at least 72 hours.

WARNING: Use this option only under the guidance of your Technical Support representative.

To automatically collect the support bundle data:

1 In the AMW, select **Advanced**→ **Troubleshooting**→ **Support Data**→ **Automatic Settings**.

2 Select **Automatically collect support data for critical events**.

3 To change the location of the saved support bundle, click **Change**.

The **Change Folder Location** window is displayed, navigate to the relevant folder and click **OK**.

4 To reset the default location, click **Reset**.

5 Click **OK**.

Collecting the Physical Disk Data

You can use the Collect Physical Disk Data option to collect log sense data from all the physical disks on your storage array.

Log sense data consists of statistical information that is maintained by each of the physical disks in your storage array. Your Technical Support representative can use this information to analyze the performance of your physical disks and for troubleshooting problems that may exist.

⚠ **WARNING: Use this option only under the guidance of your Technical Support representative.**

To collect physical disk data:

1 In the AMW, select **Advanced**→ **Troubleshooting**→ **Collect Physical Disk Data**.

The **Collect Physical Disk Data** window is displayed.

2 Enter a name for the physical disk data filename in **Specify filename** or click **Browse** to navigate to a previously saved file to overwrite an existing file.

The suffix *.**bin** is added to the file automatically if you do not specify a suffix for the file.

3 Click **Start.**

The physical disk data collection is completed and saved at the location that you entered.

4 Click **OK**.

Event Log

You can use the **Event Log Viewer** to view a detailed list of events that occur in a storage array. The event log is stored on reserved areas on the storage array disks. It records configuration events and storage array component failures. The event log can be used as a supplementary diagnostic tool to the Recovery Guru for tracing storage array events. Always see the Recovery Guru first when attempting to recover from component failures in the storage array.

⚠ **WARNING: Use this option only under the guidance of your Technical Support representative.**

The event log stores approximately 8000 events before it replaces an event with a new event. If you want to keep the events, you may save them, and clear them from the event log.

The event log window shows two types of event views:

- Summary view—Shows an event summary in a table form.
- Detail view—Shows details about a selected event.

To view the event log:

1 In the AMW, select **Advanced**→ **Troubleshooting**→ **View Event Log**.

  The **Event Log** is displayed. By default, the summary view is displayed.

2 To view the details of each selected log entry, select **View details**.

  A detail pane is added to the event log that contains detailed information about the log item. You can view the details about a single log entry at a time.

3 To save the event log, click **Save As**.

  The **Save Events** dialog is displayed.

4 Navigate to the relevant folder, enter the relevant **file name**, and click **Save**.

5 To erase all log entries from the event log, click **Clear All**.

6 To exit the event log, click **Close**.

For more information, see the *PowerVault Modular Disk Storage Manager online help* topics.

Recovery Guru

The Recovery Guru is a component of MDSM that diagnoses critical events on the storage array and recommends step-by-step recovery procedures to resolve the problems.

In the AMW, to display the Recovery Guru, perform one of these actions:

- Click **Recovery Guru** .
- In the **Support** tab, click the **Recover from Failure** link.
- From the **Status** pane on the **Summary** tab, click the **Storage Array Needs Attention** link.

You can detect a problem using the following indicators:

- Non-Optimal status icons
- Alert notification messages that are sent to the appropriate destinations
- Hardware indicator lights

The status icons return to Optimal status as problems are resolved.

Storage Array Profile

The storage array profile provides a description of all of the components and properties of the storage array. The storage array profile also provides the option to save the storage array profile information to a text file. You may want to use the storage array profile as an aid during recovery or as an overview of the current configuration of the storage array. Create a new copy of the storage array profile if your configuration changes.

To open the storage array profile:

1 In the AMW, perform one of the following actions:

- Select **Storage Array**→ **View**→ **Profile**.
- Select the **Summary** tab, and click **Storage Array Profile** in the **Status** area.
- Select the **Support** tab, and click **View Storage Array Profile**.

    The **Storage Array Profile** dialog is displayed. The Storage Array Profile dialog contains several tabs, and the title of each tab corresponds to the subject of the information contained.

2 Perform one of these actions in the **Storage Array Profile** dialog:

- View detailed information—Go to step 3.
- Search the storage array profile—Go to step 4.
- Save the storage array profile—Go to step 5.
- Close the storage array profile—Go to step 6.

3 Select one of the tabs, and use the horizontal scroll bar and the vertical scroll bar to view the storage array profile information.

You can use the other steps in this procedure to search the storage array profile, to save the storage array profile, or to close the storage array profile.

4 To search the storage array profile:

a Click ▤.

b Type the term that you want to search for in the **Find** text box.

If the term is located on the current tab, the term is highlighted in the storage array profile information.

    🖉 **NOTE:** The search is limited to the current tab. If you want to search for the term in other tabs, select the tab and click the **Find** button again.

c Click the **Find** button again to search for additional occurrences of the term.

5 To save the storage array profile:

a Click **Save As**.

b To save all sections of the storage array profile, select the **All sections**.

c To save information from particular sections of the storage array profile, select the **Select sections**, and select the check boxes corresponding to the sections that you want to save.

d Select an appropriate directory.

e   In **File Name**, type the file name of your choice. To associate the file
     with a particular software application that is displayed it, specify a file
     extension, such as .txt.

     📖 **NOTE:** The file is saved as ASCII text.

f   Click **Save**.

6   To exit the storage array profile, click **Close**.

For more information, see the *PowerVault Modular Disk Storage Manager
online help* topics.

Viewing the Logical Associations

You can use the **Associated Logical Elements** option to view the logical
associations among different virtual disks in a storage array. To view the
associations for source virtual disks, snapshot virtual disks, and snapshot
repository virtual disks:

1   In the AMW, select the **Logical** tab or in the **Topology** pane of the
    Mappings tab.

2   Select **View→ Associated Logical Elements**.

    or

    Right-click the virtual disk to open a pop-up menu and select **View→
    Associated Logical Elements**.

    If you select a virtual disk that does not have logical associations with other
    virtual disks, the **Associated Logical Elements** option is disabled.

    📖 **NOTE:** The View Associated Logical Elements dialog is displayed, which
         indicates the logical associations for the selected virtual disk.

3   To close the **View Associated Logical Elements** dialog, click **Close**.

Viewing the Physical Associations

You can use the Associated Physical Components option to view the physical components that are associated with source virtual disks, snapshot virtual disks, snapshot repository virtual disks, disk groups, unconfigured capacity, and free capacity in a storage array.

To view the physical associations:

1 In the **AMW**, select a node in the **Logical** pane of the **Logical** tab or in the **Topology** pane of the **Mappings** tab.

2 Select **View**→ **Associated Physical Components**. Alternatively, if the selected node is a virtual disk, you can right-click the node to open a pop-up menu and select **View**→ **Associated Physical Components**. If the selected node is a disk group, unconfigured capacity, or free capacity, you can right-click the node to open a pop-up menu and select **View**→ **Associated Physical Components**.

   The **View Associated Physical Components** dialog is displayed with a green triangle next to the physical components that are associated with the selected node.

3 To close the **View Associated Physical Components** dialog, click **Close**.

Finding Nodes

You can use the Find option to search for a particular node on the **Logical** tab, the **Physical** tab, or the **Mappings** tab of the AMW. The search may be based on a particular node name, the RAID level, virtual disk capacity, or specific free capacity nodes. The search may be based also on one of these combinations:

• The node name and the RAID level

• The node name and the virtual disk capacity

To find nodes:

1 In the AMW, select **View**→ **Find**.

2 Based on the type of search, select one of these options, and go to the indicated step:

- Search by name—see step 3.
- Search by special criteria—see step 4.

3 Type the name of the node to be found in **Find Node**. See step 8.

4 Based on the search criteria, select one of these options, and go to the indicated step:

- Find all virtual disks with RAID level—Go to step 5.
- Find all virtual disks with capacity—Go to step 6.
- Find all free capacity nodes—Go to step 7.

5 To search for all nodes based on their RAID level, perform these steps:

a Select **Find all virtual disks with RAID level**.

b Select the RAID level from the list.

c Go to step 8.

6 To search for all nodes based on their virtual disk capacity, perform these steps:

a Select **Find all virtual disks with capacity**.

b Type the capacity in the GB box.

c Specify that the capacity to be matched is less than, equal to, or greater than the capacity entered in the GB box.

d Go to step 8.

7 To search for all **Free Capacity** nodes with a particular capacity, perform these steps:

✎ **NOTE:** This option is not available when the Search by name option is selected or from the **Mappings** tab. You must cancel the selection of the Search by name option to use this option.

a Select **Find all free capacity nodes**.

b Type the capacity in the GB box.

c Specify that the free capacity to be matched is less than, equal to, or greater than the capacity entered in the GB box.

d Go to step 8.

8 Click **Find Next**.

To see every node that matches the criteria, click **Find Next** repeatedly. If no matches are found, the **Search Failed** dialog is displayed. Click **OK**, and re-enter the search criteria.

9 To close the dialog, click **Cancel**.

To continue searching for nodes with the same criteria after the **Find** dialog is closed, press < F3 > .

Using Go To

Use the **Go To** option to quickly jump to an associated snapshot repository virtual disk, snapshot virtual disk, source virtual disk, or target virtual disk. These virtual disks are displayed in the **Logical** pane of the **Logical** tab.

The Go To option is available only if the Snapshot premium feature or the **Virtual Disk Copy** premium feature is enabled or if snapshot virtual disks or virtual disk copies currently exist on the storage array. The Go To option is not accessible from the **Mappings** tab of the **Array Management Window**.

1 On the **Logical** tab of the AMW, select one of these virtual disks, and go to the indicated step:

- Snapshot virtual disk—Go to step 2.
- Snapshot repository virtual disk—Go to step 3.
- Source virtual disk—Go to step 4.
- Target virtual disk—Go to step 5.

2 Select **View→ Go To→ Snapshot Virtual Disk**.

The selection jumps to the associated snapshot virtual disk in the **Logical** pane.

3 Select **View→ Go To→ Snapshot Repository Virtual Disk**.

The selection jumps to the associated snapshot repository virtual disk in the **Logical** pane.

4   Select **View→ Go To→ Source Virtual Disk**.

The selection jumps to the associated source virtual disk in the **Logical** pane.

5   Select **View→ Go To→ Target Virtual Disk**.

> ✎ **NOTE:** If the source virtual disk has more than one associated target virtual disk, select the target virtual disk that you want from the list, and click OK.

The selection jumps to the associated target virtual disk in the **Logical** pane.

Recovering From an Unresponsive Storage Array Condition

A storage array can have an Unresponsive status for several reasons. Use the procedure in this topic to determine a possible cause and solution.

MDSM can take up to 5 minutes to detect that a storage array has become unresponsive or becomes responsive again. Before completing this procedure, ensure that you wait some time before you decide that the storage array is still unresponsive.

To recover from an unresponsive storage array:

1   Check the Tree View in the EMW to see if all storage arrays are unresponsive.

2   If any storage arrays are unresponsive, check the storage management station network connection to ensure that it can reach the network.

3   Ensure that the RAID controller modules are installed and that there is power to the storage array.

4   If there a problem with the storage array, correct the problem.

5   Perform one of these actions, depending on how your storage array is managed:

   • Out-of-band managed storage array—Go to step 6.

   • In-band managed storage array—Go to step 12.

6   For an out-of-band managed storage array, ensure that the RAID controller modules are network accessible by using the ping command to ensure that the RAID controller module can be reached. Type one of these commands, and press < Enter>.

   • ping < host-name>

   • ping < RAID controller module-IP-address>

7 If the verification is successful, see step 8, if not, see step 9.

8 Remove the storage array with the Unresponsive status from the EMW, and select Add Storage Array to add the storage array again.

9 If the storage array does not return to Optimal status, check the Ethernet cables to ensure that there is no visible damage and that they are securely connected.

10 Ensure the appropriate network configuration tasks are performed. For example, ensure that IP addresses are assigned to each RAID controller module.

11 If there is a cable or network accessibility problem, see step 20, if not step 12.

12 For an in-band managed storage array, ensure that the host is network accessible by using the ping command to verify that the host can be reached. Type one of these commands, and press < Enter >.

```
ping <host-name>

ping <RAID controller module-IP-address>
```

13 If the verification is successful, go to step 14. If the verification is not successful, go to step 15.

14 Remove the host with the Unresponsive status from the EMW, and select **Add Storage Array** to add the host again.

15 If the host does not return to Optimal status, go to step 16.

16 Ensure that the host is turned on and operational and that the host adapters are installed.

17 Check all external cables and switches or hubs to ensure that no visible damage exists and that they are securely connected.

18 Ensure the host context agent software is installed and running.

   If you started the host system before you were connected to the RAID controller module in the storage array, the host context agent software is not able to detect the RAID controller modules. If this is the case, ensure that the connections are secure, and restart the host context agent software.

19 If you have recently replaced or added the RAID controller module, restart the host context agent software so that the new RAID controller module is recognized.

20 If the problem still exists, make the appropriate host modifications, check with other administrators to see if a firmware upgrade was performed on the RAID controller module from another storage management station.

If a firmware upgrade was performed, the EMW on your management station may not be able to locate the new AMW software needed to manage the storage array with the new version of the firmware.

21 If the problem persists contact your Technical Support representative.

22 Determine if there is an excessive amount of network traffic to one or more RAID controller modules.

This problem is self-correcting because the EMW software periodically retries to establish communication with the RAID controller modules in the storage array. If the storage array was unresponsive and a subsequent try to connect to the storage array succeeds, the storage array becomes responsive.

For an out-of-band managed storage array, determine if management operations are taking place on the storage array from other storage management stations. A RAID controller module-determined limit exists to the number of Transmission Control Protocol/Internet Protocol (TCP/IP) connections that can be made to the RAID controller module before it stops responding to subsequent connection attempts. The type of management operations being performed and the number of management sessions taking place together determine the number of TCP/IP connections made to a RAID controller module. This problem is self-correcting because, after some TCP/IP connections terminate, the RAID controller module then becomes responsive to other connection attempts.

23 If the storage array is still unresponsive, a problem may exist with the RAID controller modules. Contact your Technical Support representative.

Locating a Physical Disk

You can use the Locate Physical Disk option to physically locate and identify one or more of the physical disks in an expansion enclosure by activating physical disk LEDs.

To locate the physical disk:

1 Select the **Physical** tab.

2 Select the physical disks that you want to locate.

3 Select **Physical Disk**→ **Blink**→ **Physical Disk**.

   The LEDs on the selected physical disks blink.

4 When you have located the physical disks, click **OK**.

   The LEDs stop blinking. If any other blink operations (Blink Disk Group, Blink Storage Array, Blink Physical Disk Ports, or Blink Expansion Enclosure) are currently being invoked from another storage management station, these LEDs also stop blinking.

5 In the rare case that the LEDs on the physical disks do not stop blinking, in the AMW, select **Storage Array**→ **Blink**→ **Stop All Indications**.

   If the LEDs successfully stop blinking, a confirmation message is displayed.

6 Click **OK**.

Locating an Expansion Enclosure

You can use the Blink option to physically locate and identify an expansion enclosure in the storage array.

The LED activation varies according to the type of expansion enclosure that you have.

• If you have an expansion enclosure with a white LED, the Blink Expansion Enclosure operation causes the white LED on the expansion enclosure to come on. The LED does not blink.

• If you have any other types of expansion enclosures, this operation causes the appropriate LED on all of the physical disks in the expansion enclosure to blink.

To locate the expansion enclosure:

1 Select the **Physical** tab.

2 Select a physical disk in the expansion enclosure that you want to locate.

3 Select **Physical Disk**→ **Blink**→ **Expansion Enclosure**.

   The LED or LEDs on the expansion enclosure or physical disks come on.

4 When you have located the expansion enclosure, click **OK**.

   The LEDs stop blinking. (If you have an expansion enclosure with a blue LED, the LED goes off). If any other blink operations (Blink Storage Array, Blink Disk Group, Blink Physical Disk Ports, Blink Expansion Enclosure, or Blink Physical Disk) are currently being invoked from another storage management station, these LEDs also stop blinking.

5 In the rare case that the LEDs on the expansion enclosure do not stop blinking, from the Array Management Window, select **Storage Array**→ **Blink**→ **Stop All Indications**.

   If the LEDs successfully stop blinking, a confirmation message is displayed.

6 Click **OK**.

Capturing the State Information

Use the **Troubleshooting**→ **Capture State Information** option to capture information about the current state of your storage array and save the captured information to a text file. You can then send the captured information to your Technical Support representative for analysis.

Potential to cause an unresponsive storage array – The State Capture option can cause a storage array to become unresponsive to both the host and the storage management station. Use this option only under the guidance of your Technical Support representative.

1 From the AMW, select **Advanced**→ **Troubleshooting** → **Capture State Information**.

2 Read the information in the **Confirm State Capture** dialog, and type **yes** to continue.

3 In the **Specify filename** text box, enter a name for the file to be saved, or browse to a previously saved file if you want to overwrite an existing file.

Use the convention filename.dmp for the name of the file. The suffix .dmp is added to the file automatically if you do not specify a suffix for the file.

4 Click **Start**.

> ![note icon] **NOTE:** Each test shows a status of Executing while it is in progress. The test then shows Completed when it successfully finishes. If any of the tests cannot be completed, a status of Failed is displayed in the Execution summary window.

5 Monitor the progress and completion status of all of the tests. When they finish, click **OK** to close the **State Capture** dialog.

Clicking **Cancel** stops the state capture process, and any remaining tests do not complete. Any test information that is generated to that point is saved to the state capture file.

> ![note icon] **NOTE:** See the *PowerVault Modular Disk Storage Manager online help* topics for more information on troubleshooting, and recovering from failures.

SMrepassist Utility

SMrepassist (replication assistance) is a host-based utility for Windows platforms. This utility is installed with MDSM. Use this utility before and after you create a virtual disk copy on a Windows operating system to ensure that all the memory-resident data for file systems on the target virtual disk is flushed and that the driver recognizes signatures and file system partitions. You can also use this utility to resolve duplicate signature problems for snapshot virtual disks.

From a command prompt window on a host running Windows, navigate to:

C:\Program Files\Dell\MD Storage Manager\util and run the following command:

```
SMrepassist –f <filesystem-identifier>
```

where -f flushes all the memory-resident data for the file system indicated by < filesystem-identifier>, and < filesystem-identifier> specifies a unique file system in the following syntax:

drive-letter: < mount-point-path>

The file system identifier may consist of only a drive letter, for example:

```
SMrepassist -f E:
```

NOTE: In Windows, the mount point path is a drive letter.

An error message is displayed in the command line when the utility cannot distinguish between the following:

- Source virtual disk and snapshot virtual disk (for example, if the snapshot virtual disk is removed).

- Standard virtual disk and virtual disk copy (for example, if the virtual disk copy is removed).

Unidentified Devices

An unidentified node or device occurs when MDSM cannot access a new storage array. Causes for this error include network connection problems, the storage array is turned off, or the storage array does not exist.

NOTE: Before beginning any recovery procedure, ensure that the host context agent software is installed and running. If you host is started before it is connected to the storage array, the host context agent software is not able to find the storage array. If so, ensure that the connections are tight, and restart the host context agent software.

- If a storage array is managed by using both out-of-band management and in-band management using the same host, a management network connection problem may prevent direct communication with the storage array. However, you may still be able to manage the storage array over the in-band connections. The opposite situation can also occur.

- If a storage array is managed through more than one host, it is possible that the storage array may become unresponsive to communication over the connections given by one host. However, you may still be able to manage the storage array over the connections provided by another host.

Recovering From an Unidentified Storage Array

1 Ensure that the network connection to the storage management station is operating.

2 Ensure that the controllers are installed and that the power is turned on to the storage array. Correct any existing problems before continuing.

3 If you have an in-band storage array, use the following procedure. Click **Refresh** after each step to check the results:

a Ensure that the host context agent software is installed and running. If the host is started before it is connected to the controllers in the storage array, the host context agent software is not able to find the controllers. If so, ensure that the connections are tight, and restart the host context agent software.

b Ensure that the network can access the host by using the ping command in the following syntax:

```
ping <host-name-or-IP-address-of-the-host>.
```

If the network can access the host, continue to step c. If the network cannot access the host, go to step d.

c Remove the host with the unresponsive status from the MDSM, and add that host again.

If the host returns to optimal status, you have completed this procedure.

d Ensure that the power to the host is turned on and that the host is operational.

e If applicable, ensure that the host bus adapters are installed in the host.

f Examine all external cables and switches or hubs to ensure that you cannot see any damage and that they are tightly connected.

g If you have recently replaced or added the controller, restart the host context agent software so that the new controller is found.

If a problem exists, make the appropriate modifications to the host.

4 If you have an out-of-band storage array, use the following procedure. Click **Refresh** after each step to ensure the results:

a Ensure that the network can access the controllers by using the ping command. Use the following syntax:

```
ping <controller-IP-address>
```

If the network can access the controllers, continue to step b. If the network cannot access the controllers, go to step c.

b Remove the storage array with the unresponsive status from MDSM, and add that storage array again.

If the storage array returns to optimal status, you have completed this procedure.

c Examine the ethernet cables to ensure that you cannot see any damage and that they are tightly connected.

d Ensure that the applicable network configuration tasks are done (for example, the IP addresses are assigned to each controller).

5 Ensure that the controller firmware is compatible with MDSM on your management station. If the controller firmware was upgraded, the MDSM may not have access to the storage array. A new version of MDSM may be needed to manage the storage array with the new version of the controller firmware.

If this problem exists, see **dell.com/support**.

6 Look to see if there is too much network traffic to one or more controllers. This problem corrects itself because the MDSM tries to re-establish communication with the controllers in the storage array at regular times. If the storage array was unresponsive and a subsequent try to connect to the storage array succeeds, the storage array becomes responsive.

7 For an out-of-band storage array, look to see if management operations are taking place on the storage array from other storage management stations. The type of management operations being done and the number of management sessions taking place together establish the number of TCP/IP connections made to a controller. When the maximum number of TCP/IP connections are made, the controller stops responding. This problem corrects itself because after some TCP/IP connections complete, the controller then becomes responsive to other connection tries.

8 If the storage array is still unresponsive, problems may exist with the controllers.

If these problems exist, see **dell.com/support**.

Starting or Restarting the Host Context Agent Software

The host context agent software module is the software component that resides on the server or management station that communicates with the MD3600f series storage arrays. The SMagent software automatically starts after you reboot the host.

Windows

To restart the SMagent software in Windows:

1   Click **Start**→ **Settings**→ **Control Panel**→ **Administrative Tools**→ **Services**.

    or

    Click **Start**→ **Administrative Tools**→ **Services**.

2   In the **Services** dialog, select **Modular Disk Storage Manager Agent**.

3   If the modular disk storage manager agent is running, click **Action**→ **Stop**, then wait approximately 5 seconds.

4   Click **Action**→ **Start**.

Linux

To start or restart the host context agent software in Linux, enter the following command at the prompt:

```
SMagent start
```

The SMagent software may take a little time to initialize. The cursor is shown, but the terminal window does not respond. When the program starts, the following message is displayed:

```
SMagent started.
```

After the program completes the startup process, text similar to the following messages is displayed:

```
Modular Disk Storage Manager Agent, Version
90.02.A6.14

Built Wed Feb 03 06:17:50 CST 2010
```

```
Copyright (C) 2009-2010 Dell, Inc. All rights
reserved.
Checking device <n/a> (/dev/sg10): Activating
Checking device /dev/sdb (/dev/sg11): Skipping
Checking device <n/a> (/dev/sg3): Activating
Checking device <n/a> (/dev/sg4): Activating
Checking device <n/a> (/dev/sg5): Activating
Checking device <n/a> (/dev/sg6): Activating
Checking device <n/a> (/dev/sg7): Activating
Checking device <n/a> (/dev/sg8): Activating
Checking device <n/a> (/dev/sg9): Activating
```

20

Troubleshooting: Your Array

Safety First—For You and Your Array

⚠️ **CAUTION: Many repairs may only be done by a certified service technician. You must only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that came with the product.**

Troubleshooting Storage Array Startup Failure

If your system halts during startup, check if:

- The array emits a series of beeps.
- The array fault LEDs are lit. See "RAID Controller Modules" on page 33.
- There is a constant scraping or grinding sound when you access the hard drive. See "Getting Help" on page 309.

Troubleshooting Loss of Communication

For information about troubleshooting loss of communication, see "Troubleshooting Array and Expansion Enclosure Connections" on page 303.

Troubleshooting an SFP Transceiver

⚠️ **CAUTION: When handling static-sensitive devices, take precautions to avoid damaging the product from static electricity.**

📝 **NOTE:** Before proceeding with this section see "Working with SFP Modules" and "Fiber Optic Cables" in the *Dell PowerVault MD3600f Deployment Guide* at **dell.com/support**.

Equip yourself with antistatic protection and a replacement small form-factor pluggable (SFP) transceiver before replacing an SFP transceiver in the RAID controller module. Also, see the initial setup information for the enclosure to verify LED definitions.

1 If possible, use the storage management software to create, save, and print a new storage array profile.

2 If the Recovery Guru directed you to replace a failed RAID controller module, go to step 3, else run the Recovery Guru to identify the failed component.

3 Unpack the new SFP transceiver.

   ✍ **NOTE:** Ensure that the SFP transceiver is replaced by the same type.

   ✍ **NOTE:** Set the new SFP transceiver on a dry, level surface near the RAID controller module, the RAID controller module enclosure, or the expansion enclosure.

   ✍ **NOTE:** Save all the packing materials in case you need to return the SFP transceiver.

4 To locate a failed SFP transceiver, view the FCIN (host channel) speed LEDs on the front of the RAID controller modules. If an SFP transceiver has failed, both FCIN speed LEDs for a particular port are off.

   • If both FCIN speed LEDs are off for a particular port with a known good cable, the SFP transceiver must be replaced. Go to step 6.

   • If at least one FCIN speed LED is on for a particular port, the SFP transceiver is functional. The FCIN speed LEDs indicate a channel speed of 2 Gbps, 4 Gbps, or 8 Gbps (see "Host Channel LED Link/Rate Indications" on page 36). Ensure that the speed indicated by the LEDs is what is expected.

⚠ **WARNING: Data processing environments can contain equipment transmitting on system links with laser modules that operate at greater than Class 1 power levels. Never look into the end of an optical fiber cable or open receptacle.**

⚠ **CAUTION: To prevent degraded performance, do not twist, fold, pinch, or step on fiber optic cables. Do not bend the fiber optic cables tighter than a 5 cm (2") radius.**

5 If present, disconnect the fiber optic cable from the failed SFP transceiver.

6 Remove the failed SFP transceiver from the interface port.

Figure 20-1.   Removing an SFP Module

7   Install the new SFP transceiver into the interface port. Ensure that the transceiver is properly seated.

8   Reconnect the fiber optic cable.

9   View the FCIN speed LEDs and the Physical Disk Channel speed LEDs.

Based on the LED status, perform one of these actions:

- At least one of the FCIN LEDs for each port is on. Go to step 11.

- Both the FCIN LEDs on an MD storage array RAID controller module enclosure for a particular port are off.

- Check that the SFP transceiver is installed correctly. Reinstall the SFP transceiver if necessary.

   ✎ **NOTE:** If the storage array is not in an Optimal status, click the Recovery Guru toolbar button in the Array Management Window to determine if any other actions are required.

10 If the problem is resolved, go to step 11. If the problem persists, see "Getting Help" on page 309.

11 Create, save, and print a new storage array profile.

Troubleshooting External Connections

- Verify that the cables are connected to the correct ports before troubleshooting any external devices. To locate the back-panel connectors on your array, see Figure 3-1.

- Ensure that all the cables are securely attached to the external connectors on your array.

- For information on cabling, see the *Dell PowerVault MD3600f Deployment Guide* at **dell.com/support/manuals**.

Troubleshooting Power Supply/Cooling Fan Module

⚠ **CAUTION: Many repairs may only be done by a certified service technician. You must only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that came with the product.**

⚠ **CAUTION: It is recommended that you turn off the host server before turning off the array to prevent loss of data.**

1 Locate the faulty power supply and determine the status of the LEDs.

- If the AC power LED is not lit, check the power cord and power source into which the power supply is plugged.

  - Connect another device to the power source to verify if it is working.

  - Connect the cable to a different power source.

  - Replace the power cable.

  If the problem is not resolved, see "Getting Help" on page 309.

- If the DC power LED is not lit, verify that the power switch is turned on. If the power switch is turned on, see step 2.

- If the power supply's fault indicator is lit, see "Getting Help" on page 309.

⚠ **CAUTION: Power supply/cooling fan modules are hot-swappable. The array can operate on a single power supply; however both modules must be installed to ensure proper cooling. A single power supply/cooling fan module can be removed from a powered-on array for a maximum period of 5 minutes. Beyond that time, the array may automatically shut down to prevent damage.**

2   Reseat the power supply by removing and reinstalling it. See "Power Supply and Cooling Fan Features" on page 31.

   📝 **NOTE:** After installing a power supply, allow several seconds for the array to recognize the power supply and to determine if it is working properly.

   If the problem is not resolved, see "Getting Help" on page 309.

Troubleshooting Array Cooling Problems

⚠ **CAUTION: Many repairs may only be done by a certified service technician. You must only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that came with the product.**

Ensure that none of the following conditions exist:

- Array cover or drive blank is removed.
- Ambient temperature is too high. See "Technical Specifications" in the *Getting Started Guide*.
- External airflow is obstructed.
- The power supply/cooling fan module is removed or has failed. See "Troubleshooting Power Supply/Cooling Fan Module" on page 298.

If the problem is not resolved, see "Getting Help" on page 309.

Troubleshooting Expansion Enclosure Management Modules

⚠ **CAUTION: Many repairs may only be done by a certified service technician. You must only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that came with the product.**

⚠ **CAUTION: It is recommended that you turn off the host server before turning off the expansion enclosure array to prevent loss of data.**

- If the EMM status LED is solid or blinking amber (2 or 4 times per sequence):

  a   Turn off the server.

  b   Remove the EMM and verify that the pins on backplane and EMM are not bent. See Removing an EMM in the MD1200 and MD1220 Storage Enclosures *Hardware Owner's Manual.*

  c   Reseat the EMM module and wait for 30 seconds. See Removing an EMM in the MD1200 and MD1220 Storage Enclosures *Hardware Owner's Manual.*

  d   Turn on the server.

  e   Check the EMM status LED.

  f   If the LED does not turn Green, replace the EMM.

  If the problem is not resolved, step c.

- If EMM status LED is blinking amber (5 times per sequence), update the firmware to the latest supported firmware on both the EMMs. For more information about downloading the latest firmware, see "Management: Firmware Downloads" on page 229.

- If the link status LEDs are not green:

  a   Turn off the server.

  b   Reseat the cables on the expansion array and the server.

  c   Turn on the expansion arrays and then the storage array and wait until the system is fully booted.

d   Turn on the server.

   e   Check the link status LED. If the link status LED is not green, replace the cables.

   If the problem is not resolved, see "Getting Help" on page 309.

Troubleshooting RAID Controller Modules

⚠ **CAUTION: Many repairs may only be done by a certified service technician. You must only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that came with the product.**

⚠ **CAUTION: In the case of non-redundant configurations, it is recommended that you turn off the host server before turning off the array to prevent loss of data.**

- If the array status LED is solid or blinking amber:

   a   In the AMW, select the **Summary** tab, and click on **Storage Array needs attention**. Follow the listed procedures in the Recovery Guru(s) and wait for up to 5 minutes to check if the LED has turned blue. See "Recovery Guru" on page 278.

   b   If following the recovery guru procedures does not solve the problem, complete the following procedure to further troubleshoot the array.

   c   Turn off the host server as appropriate.

   d   Remove the RAID controller module and verify that the pins on backplane and RAID controller module are not bent. See "Removing a RAID Controller Module Blank" on page 249.

   e   Reinstall the RAID controller module and wait for 30 seconds. See "Installing a RAID Controller Module" on page 252.

   f   Check the RAID controller module status LED.

   g   Replace the RAID controller module.

   h   Turn on the host server.

   If the problem is not resolved, see "Getting Help" on page 309.

- If both LEDs for any given FCIN port are unlit, see "Troubleshooting Array and Expansion Enclosure Connections" on page 303.

    a Turn off the server, storage arrays, and expansion enclosures.

    b Reseat the RAID controller module and reconnect cables on the storage array and the server.

    c Restart the storage array and wait until the array is fully booted.

    d Turn on the server.

    e Recheck the LEDs of the affected port(s). Replace the fiber optic cables of any port(s) where both LEDs are unlit.

    NOTE: The SAS OUT status LED must be green.

    If the problem is not resolved, see "Getting Help" on page 309.

Troubleshooting Hard Drives

CAUTION: Many repairs may only be done by a certified service technician. You must only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that came with the product.

1 Check the storage array profile to ensure that the most current version of the firmware is installed. For more information, see the *Support Matrix* at **dell.com/support/manuals**.

2 Remove the hard drive from the system. See "Removing a Hard Drive" on page 244.

    NOTE: You must ensure that you check the hard drive indicators before removing the faulty hard drive from the system.

3 Check the hard drives and the backplane to ensure that the connectors are not damaged.

4 Reinstall the hard drive.

5 If the problem is not resolved, replace the failed hard drive.

    If the problem persists, see "Troubleshooting Loss of Communication" on page 295, "Troubleshooting an SFP Transceiver" on page 295, or "Getting Help" on page 309.

Troubleshooting Array and Expansion Enclosure Connections

1 Verify that the SAS OUT status LED is green and that for each FCIN port connected to a cable, at least one of the paired LEDs is on. If this is not the case, see "Planning: RAID Controller Modules" on page 33.

2 Ensure that all the cables are attached correctly according to array mode you selected.

3 Turn off the server, storage array, and expansion enclosures.

4 Reseat the RAID controller module and reconnect cables on the storage array and the server.

5 Turn on the expansion arrays and then the storage array and wait until the system is fully booted.

6 Turn on the server.

7 Check the LEDs of all cabled ports. Replace the fiber optic cables of any port(s) where both LEDs are unlit.

   If the problem is not resolved, see "Getting Help" on page 309.

8 Reboot the host server.

   NOTE: You must turn off the host server before resetting the cables on the storage array or expansion enclosure.

   If the problem is not resolved, see "Getting Help" on page 309.

Troubleshooting a Wet Storage Array

⚠ **CAUTION: Many repairs may only be done by a certified service technician. You must only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that came with the product.**

1 Turn off the array and disconnect all the cables.

2 Remove the following components from the array. See "Management: Installing Array Components" on page 241.

   • Hard drives

   • RAID controller modules

   • Power supply/cooling fan modules

   • Control panel

   • Backplane

3 Let the system dry thoroughly for at least 24 hours.

4 Reinstall the components you removed in step 2.

5 Connect all the cables and turn on the array.

   If the array does not start properly, see "Getting Help" on page 309.

Troubleshooting a Damaged Array

⚠ CAUTION: Many repairs may only be done by a certified service technician. You must only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that came with the product.

1 Ensure that the following components are properly installed:

- Hard drives
- RAID controller modules
- Power supply/cooling fan modules
- Control panel
- Backplane

2 Ensure that all the cables are properly connected and that there are no damaged pins in the connectors.

3 Run diagnostics available in Dell PowerVault Modular Disk (MD) Storage Manager software. In the AMW, select a component in the Physical pane of the Physical tab. Select **Advanced**→ **Troubleshooting**→ **Run Diagnostics**.

If the test fails, see "Getting Help" on page 309.

Troubleshooting RAID Controller Modules

Conditions

Certain events can cause a RAID controller module to fail and/or shut down. Unrecoverable ECC memory or PCI errors, or critical physical conditions can cause lockdown. If your RAID storage array is configured for redundant access and cache mirroring, the surviving controller can normally recover without data loss or shutdown.

Typical controller failures are detailed in the following sections.

Invalid Storage Array

The RAID controller module is supported only in a Dell-supported storage array. After installation in the storage array, the controller performs a set of validation checks. The array status LED is lit with a steady amber color while the RAID controller module completes these initial tests and the controllers are booted successfully. If the RAID controller module detects a non-Dell supported storage array, the controller aborts startup. The RAID controller module does not generate any events to alert you in the event of an invalid array, but the array status LED is lit with a flashing amber color to indicate a fault state.

For full details on the LEDs and their interpretation, see "Back Panel Features and Indicators" on page 29.

ECC Errors

RAID controller firmware can detect ECC errors and can recover from a single-bit ECC error whether the RAID controller module is in a redundant or nonredundant configuration. A storage array with redundant controllers can recover from multi-bit ECC errors as well because the peer RAID controller module can take over, if necessary.

The RAID controller module failsover if it experiences up to 10 single-bit errors, or up to 3 multi-bit errors.

PCI Errors

The storage array firmware can detect and only recover from PCI errors when the RAID controller modules are configured for redundancy. If a virtual disk uses cache mirroring, it fails over to its peer RAID controller module, which initiates a flush of the dirty cache.

Critical Conditions

The storage array generates a critical event if the RAID controller module detects a critical condition that could cause immediate failure of the array and/or loss of data. The storage array is in a critical condition if one of the following occurs:

- More than one fan has failed
- Any backplane temperature sensors in the critical range
- Backplane/power supply failure
- Two or more temperature sensors are unreadable
- Failure to detect or unable to communicate with peer port

NOTE: If both RAID controller modules fail simultaneously, the enclosure cannot issue critical or noncritical event alarms for any enclosure component.

When the enclosure is under critical condition, its enclosure status LED blinks amber.

Noncritical Conditions

A noncritical condition is an event or status that does not cause immediate failure, but must be corrected to ensure continued reliability of the storage array. Examples of noncritical events include the following:

- One power supply has failed
- One cooling fan has failed
- One RAID controller module in a redundant configuration has failed
- A battery has failed or is removed
- A physical disk in a redundant virtual disk has failed

When the enclosure is under noncritical condition, its enclosure status LED blinks amber.

Getting Help

Locating Your System Service Tag

Your system is identified by a unique Express Service Code and Service Tag number. The Express Service Code and Service Tag are found on the front of the system by pulling out the information tag. This information is used by Dell to route support calls to the appropriate personnel.

Contacting Dell

NOTE: Dell provides several online and telephone-based support and service options. If you do not have an active Internet connection, you can find contact information on your purchase invoice, packing slip, bill, or Dell product catalog. Availability varies by country and product, and some services may not be available in your area.

To contact Dell for sales, technical support, or customer-service issues:

1 Go to **dell.com/contactdell**.

2 Select your country or region from the interactive world map.

   When you select a region, the countries for the selected regions are displayed.

3 Select the appropriate language under the country of your choice.

4 Select your business segment.

   The main support page for the selected business segment is displayed.

5 Select the appropriate option depending on your requirement.

NOTE: If you have purchased a Dell system, you may be asked for the Service Tag.

Documentation Feedback

If you have feedback for this document, write to
documentation_feedback@dell.com. Alternatively, you can click on the
Feedback link in any of the Dell documentation pages, fill up the form, and
click **Submit** to send your feedback.

Index