Dell OpenManage™ Array Manager 3.6

User's Guide

Information in this document is subject to change without notice.

© 2004 Dell Inc. © 2004 VERITAS Software Corporation. All rights reserved.

Reproduction in any manner whatsoever without the written permission of Dell Inc. and VERITAS Software Corporation is strictly forbidden.

Portions copyright (c) 1998 Chris Maunder.

Some cryptographic code in this product is derived from the SSLeay library, copyright (c) 1995-97 by Eric Young (eay@cryptsoft.com).

All rights reserved.

Trademarks used in this text: Dell, Dell OpenManage, the DELL logo, PowerEdge, and PowerVault are trademarks of Dell Inc.; Microsoft, Windows, Windows NT, Windows 2000, and Windows Server 2003 are registered trademarks of Microsoft Corporation; Novell and NetWare are registered trademarks of Novell, Inc. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell Inc. and VERITAS Software Corporation disclaim any proprietary interest in trademarks and trade names other than their own.

August, 2004

What this Product Does 1

Array Manager is designed for system administrators who implement hardware and software RAID solutions and understand corporate and small business storage environments.

Array Manager enables you to configure the storage components attached to your system. These components include RAID and non-RAID controllers and the channels, enclosures, and disks attached to them. Using Array Manager, you can implement controller functions without requiring access to the BIOS. These functions include configuring virtual disks and applying RAID levels and hot spares for data protection. You can initiate many other controller functions such as rebuilds, troubleshooting, and so on. Most functions can be implemented while the server remains online and continues to process requests.

The Array Manager console reports the status of storage components through error icons and graphical displays. When there is a change in status, Array Manager sends events to the Array Manager event log which is viewable from the console. Most events also generate SNMP traps.

Other than monitoring and reporting status, Array Manager does not automatically initiate actions independent of user input. (Automatic shutdown of enclosures that have exceeded a critical temperature is the only exception. See “Thermal Shutdown.”) Array Manager actions are user-initiated using the wizards, pull-down menus, and right-click functions provided by the console. Array Manager does, however, report the actions taken by the controllers, which include generating events, initiating tasks such as a rebuild, and making state changes.

Note Array Manager reports the state of disks and other storage components as viewed by the controller.

For more information, see the following:

- Understanding Storage Concepts
- Wizards and Common Tasks
- Viewing Storage Status and Configuration
- Event Monitoring

Array Manager Components

Array Manager consists of the following components:

- **Management console or client.** The management console or client provides the graphical user interface. The console is installed on the system from which you wish to manage storage and can connect to remote systems for remote storage management. For more information see:
  - Before You Begin and Installation Overview for installation information
  - Starting Array Manager
  - Using the Console
- **Managed system or server.** The managed system or server is installed as a system service on all systems that have storage you wish to manage. The managed system identifies the local storage and provides this information to the management console. The managed system also implements commands received from the console. For more information see:
  - Before You Begin and Installation Overview for installation information
  - Starting the Array Manager Service for the First Time
  - Starting and Stopping the Array Manager Service
- **Array Manager Utilities.** The Array Manager Utilities enable you to modify the Array Manager configuration. The Array Manager Utilities also display firmware and driver version levels. For more information, see "Array Manager Utilities."

When selecting or deselecting the **Hide unknown Windows disks at startup** check box, you need to restart the console in order for the new setting to take effect.

Note The **Hide unknown Windows disks at startup** feature is not available on NetWare and Windows Server 2003.

Supported Controllers

This release of Array Manager supports the following controllers:

- PERC Subsystem 1 Controllers — PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, PERC 4/IM, 4e/Si, 4e/Di, and CERC ATA100/4ch
- PERC Subsystem 2 Controllers — PERC 2, 2/Si, 3/Si, 3/Di, and CERC SATA1.5/6ch, and CERC SATA1.5/2s
- SCSI Subsystem Controllers — Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030.

- PowerVault 660F storage system with dual fibre channel RAID controllers. For more information, see "The Dell PowerVault 660F Storage System."

The firmware and drivers listed in the Readme file refer to the minimum supported version. Later versions of the firmware and drivers are also supported. Refer to http://support.dell.com for the most recent driver and firmware requirements.

NetWare Support

This release supports NetWare® version 5.1 and 6.5. Array Manager does not support NetWare volume management or clustering.

See the Dell PowerVault Systems Storage Area Network (SAN) documentation for information on supported NetWare configurations.

Note Dell does not offer NetWare in Japan.

Support for Windows Server 2003

Array Manager provides support for the Windows Server 2003 operating system. This support includes the following:

- The PERC 3/QC and PowerVault 660F controllers as well as 64-bit applications are supported on a Windows Server 2003 64-bit system.
- All other controllers supported by Array Manager are supported on a Windows Server 2003 32-bit system.

Array Manager does not support Windows disk or volume management on Windows Server 2003.

For information on Windows Server 2003 and connecting to a remote system, see "Remote Connection and Windows Server 2003."

For information on SNMP support, see "SNMP Support."

Array Manager and the Host Operating System

Array Manager's secured local and remote storage management uses the native security model of the host operating system.

Under Windows 2000 and Windows Server 2003, you can use Array Manager's Microsoft Management Console (MMC) snap-in instead of running the program from the Array Manager console. For more information on the MMC Snap-in, see "MMC Snap-in."

In Windows 2000 and Windows Server 2003, Array Manager is also available as one of the programs or utilities in the Computer Management window. To access the Computer Management window, use **Start | Settings | Control Panel | Administrative Tools | Computer Management**.

For information on support for NetWare and Windows Server 2003, see "NetWare Support" and "Support for Windows Server 2003."

For information on Array Manager and Windows Disk Management, see "Array Manager and Windows Disk Management."

Array Manager and Windows Disk Management

Disk Management is the built-in disk and volume manager that comes with Windows 2000 and Windows Server 2003. Disk Management is also known as Logical Disk Manager, or LDM.

On Windows 2000, installing Array Manager turns off the service for Disk Management. (See "Replacing Windows 2000 Disk Management.") However, if you uninstall Array Manager, the service for Disk Management is turned back on, and the Disk Management program will function normally. The service for Disk Management is not turned off on a system that has a console-only installation of Array Manager.

Because Array Manager and Disk Management are related products, Array Manager has the same disk and volume functions as Disk Management plus the added capability of managing RAID hardware controllers. If you created disks and volumes with Disk Management before installing Array Manager, those same disks and volumes can be managed by Array Manager without any adjustments.

Installation, Upgrades, and Uninstall 2

This chapter explains how to install Dell OpenManage™ Array Manager, including hardware and software requirements, upgrade information, and installation types and components. There will also be a section about uninstall.

Before You Begin

Before installing, you should:

- Read the Installation Requirements to be sure that your system meets or exceeds the minimum requirements.
- Read the Readme file. It contains important information about software, firmware, and driver versions, plus information about known issues.
- Check the Dell support site for more current versions of the Array Manager software or for other updates such as service packs. You can check the Dell support site at: http://support.dell.com.
- Read the installation chapter completely.

Note Array Manager can be installed from more than one source, including the *Dell OpenManage Applications* CD, the Array Manager CD, and from the Dell support site. Throughout this online help, these installation sources may be referred to as the Array Manager CD or the Array Manager installation media.

Installation Overview

This overview section contains the following topics:

- Installation Components
- Replacing Windows 2000 Disk Management
- ASPI Driver and Non-RAID SCSI Disk and Enclosure Management

Installation Components

Array Manager has the following installation components:

- Management console (client) for Windows® 2000, Windows Server 2003®, Windows XP
- Managed system (server) for Windows 2000, and Windows Server 2003
- Managed system (server) for NetWare

You can install these components using the following installation options:

1. Typical installation for Microsoft® Windows
   - Management console for Windows 2000, Windows Server 2003, Windows XP, and NetWare
   - Managed system (server) for Windows 2000, and Windows Server 2003
   - Help files
2. Custom installation
   - You choose which components to install

Replacing Windows 2000 Disk Management

When using Express Install on a Windows 2000 server, Array Manager replaces the native Windows 2000 Disk Management, the built-in disk and volume manager that comes with Windows 2000. Array Manager extends the functionality of Disk Management by providing both hardware and software RAID management, while maintaining the disk and volume functions originally found in the Windows 2000 Disk Management.

If you do not wish to replace the Windows 2000 Disk Management during installation, then use the Custom Install.

After replacing the Windows 2000 Disk Management, if you uninstall Array Manager, the Disk Management program will be restored to its original state.

Note Disk Management is not disabled on a system that has a console-only installation of Array Manager.

For information on Array Manager and Windows Disk Management, see "Array Manager and Windows Disk Management."

ASPI Driver and Non-RAID SCSI Disk and Enclosure Management

In order to provide non-RAID SCSI disk and enclosure management, Array Manager requires the Advanced SCSI Programming Interface (ASPI) driver. The 4.6 version of this driver is installed with the Array Manager management console.

Note The Adaptec aspi32.sys driver version 4.6 that is installed by Array Manager has not passed the WHQL certification tests that are recommended by Microsoft. Running Microsoft's Driver Verifier against this driver will report problems with the driver. It is recommended that Microsoft's Driver verifier not be run once Array Manager has been installed.

Security Considerations and the ASPI Driver

Array Manager installs the ASPI 4.6 driver as part of the management console installation. The ASPI 4.6 driver provides a programming interface to SCSI devices.

The ASPI driver is a legacy driver and does not adhere to Microsoft's standards for security access. For this reason, the presence of the ASPI driver on a system can be considered a security risk. It is not recommended that the ASPI driver be installed on a server that is not secure or that has public access.

If you choose to uninstall the ASPI driver, Array Manager will no longer provide non-RAID SCSI disk and enclosure management. See Uninstalling the ASPI 4.6 Driver for more information.

Installing on a System with the ASPI 4.7.x Driver

Array Manager cannot be installed on a system that has the ASPI 4.7.x driver installed.

If you have the ASPI 4.7.x driver installed on your system, you need to remove this driver and install the ASPI 4.6 driver. You can do this by downloading the ASPI 4.6 package from http://support.dell.com. The ASPI 4.6 package uninstalls the ASPI 4.7.x driver and installs the ASPI 4.6 driver. Refer to the documentation of any other applications installed on the system to see if those applications have known problems with the ASPI 4.6 driver.

Uninstalling the ASPI 4.6 Driver

You can remove the ASPI 4.6 driver that has been installed by Array Manager. To do so, download the ASPI 4.6 package from http://support.dell.com. This package is designed to uninstall existing versions of the ASPI 4.7.x driver and install the ASPI 4.6 driver in its place. This package also contains the remvaspi.bat script which uninstalls the ASPI 4.6 driver.

After downloading the ASPI 4.6 package, do the following to remove the ASPI 4.6 driver from your system:

1. Unzip the download package.
2. Use Windows Explorer to locate the unzipped download package. The default location is:

   C:\temp_aspi.
3. Double-click the **remvaspi.bat** executable.
4. Reboot the server when the remvaspi.bat script has completed.

Installation Requirements

The following sections describe the Array Manager installation requirements:

- General Requirements
- Management Console Requirements (Client)
- Managed System for Windows Requirements (Server)
- Managed System for NetWare Requirements (Server)
- Online Help Requirements

Please see the Readme file for the latest information.

General Requirements

The following requirements apply to any system running the Array Manager console (client) or the Array Manager managed system (server):

- 64 megabytes (MB) of RAM.
- Administrator access to the system.
- For information on supported operating systems, drivers, and firmware for the Array Manager console and managed system, see the following:
  - Management Console Requirements (Client)
  - Managed System for Windows Requirements (Server)
  - Managed System for NetWare Requirements (Server)

- Transmission Control Protocol/Internet Protocol (TCP/IP) connection on the console and server when attempting remote system management.
- Simple Network Management Protocol (SNMP) services are required for SNMP support. If you do not have SNMP, Array Manager will function as normal but will not provide SNMP support.

The Array Manager console and managed system have additional system requirements. For information on these requirements as well as the supported operating systems, see the following:

- Management Console Requirements (Client)
- Managed System for Windows Requirements (Server)
- Managed System for NetWare Requirements (Server)
- Online Help Requirements

Note After Array Manager is installed, the Windows Add/Remove Programs utility may incorrectly report the Array Manager installation to be smaller than it is. For the correct disk space requirements, refer to the Management Console Requirements (Client) and the Managed System for Windows Requirements (Server) sections.

Note If you intend to web serve Array Manager, Internet Information Services (including the World Wide Web Server component) must be installed on the Web server prior to installing Array Manager. Use the Windows server installation media to install Internet Information Services. If you install Internet Information Services after installing Array Manager, you will need to Add/Modify (reinstall) Array Manager before web serving an Array Manager installation.

Note Dell provides configuration utilities that enable you to perform certain RAID management operations using the PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4/IM, 4e/Si, 4e/Di, CERC ATA100/4ch, and PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers on platforms not currently supported by Array Manager. These utilities are available at http://support.dell.com.

Management Console Requirements (Client)

The following are minimum requirements for installing the Array Manager console. The console is the client portion of Array Manager. For the latest information, see the Readme file.

- 64 megabytes (MB) of RAM.

- 40 MB of free hard drive space.
- One of the following versions of Windows:
  - Microsoft Small Business Server
  - Microsoft Windows 2000® Server with service pack 3 or later
  - Microsoft Windows 2000 Advanced Server with service pack 3 or later
  - Microsoft Windows 2000 Professional with service pack 3 or later
  - Microsoft Windows 2000 Small Business Suite
  - Microsoft Windows XP
  - Microsoft Windows Server 2003 Enterprise Server

Note It is highly recommended that you use the most current service pack.

- Administrator rights.
- Transmission Control Protocol / Internet Protocol (TCP / IP) connection on the console and server when attempting remote system management.
- Microsoft Management Console (MMC) Version 1.2 or later for snap-in support on Windows 2000 and Windows Server 2003. For more information, see "MMC Snap-in."
- (Optional). Novell Client for Windows version 4.9 or later if you intend to map a drive to NetWare system and copy the Array Manager install media from the Windows system to the NetWare system.
- (Optional). Microsoft Internet Explorer Version 4.0 or later for viewing the online help and the browser plug-in.

Managed System for Windows Requirements (Server)

The following are minimum requirements for installing a managed system for Windows. For the latest information, see the Readme file.

- A server running one of the following versions of Windows:
  - Microsoft Small Business Server
  - Microsoft Windows 2000 Server with service pack 3 or later
  - Microsoft Windows 2000 Advanced Server with service pack 3 or later
  - Microsoft DataCenter
  - Microsoft Windows Server 2003 Enterprise Server

Note It is highly recommended that you use the most current service pack.

- 64 MB of RAM.
- 15 MB of free hard disk space on the local drive.
- Administrator access to the system.
- Hardware and drivers installed for at least one Dell RAID controller. Drivers and firmware for supported controllers can be downloaded from the Dell support site at http://support.dell.com. See the Readme file for supported drivers and firmware.
- Fibre Channel Host Bus Adapters and associated device drivers are needed to manage the Dell PowerVault 660F Fibre Channel RAID Array in either a direct-attach or SAN environment.
- (Optional) TCP/IP installation if this system will be managed over a network.
- (Optional). Simple Network Management Protocol (SNMP) for SNMP trap forwarding by Array Manager. The SNMP service should be installed prior to installing Array Manager.
- (Optional) On NetWare, Array Manager requires that the Dell OpenManage Server Agent (OMSA) 4.0.0 or later be installed on the managed system to provide thermal shutdown (when an enclosure becomes too hot or cold) for the Dell PowerVault 20xS, 21xS, and 22xS SCSI enclosures. If OMSA is not installed, Array Manager can provide other features of enclosure management, but not thermal shutdown without the risk of data loss.

Note The managed system's RAID controllers must be attached to storage in order to run Array Manager.

Managed System for NetWare Requirements (Server)

The following are minimum requirements for installing a managed system for NetWare. For the latest information, see the Readme file.

- 64 megabytes of RAM.
- 15 megabytes of free hard disk space.
- Administrator access to the system.
- One of the following versions of Novell NetWare:
- Novell NetWare version 5.1 or later with Support Pack 7 or later.

- Novell NetWare version 6.5 or later with Support Pack 2 or later is supported in a SCSI environment. (See the Dell PowerVault Systems Storage Area Network (SAN) documentation for information on SAN support for NetWare.).
- Hardware and drivers installed for at least one Dell RAID controller. Drivers and firmware for supported controllers can be downloaded from the Dell support site at http://support.dell.com. See the Readme file for the supported drivers and firmware.
- Fibre Channel Host Bus Adapters and associated device drivers are needed to manage the Dell PowerVault 660F Fibre Channel RAID Array in either a direct-attach or SAN environment. See the Dell PowerVault Systems Storage Area Network (SAN) documentation for information on supported NetWare configurations.
- Transmission Control Protocol/Internet Protocol (TCP/IP) required for connection to the management console.
- Long file name support.
- (Optional). Novell Client for Windows version 4.9 or later installed on a remote Windows system if you intend copy the Array Manager install media from the Windows system to the NetWare system by mapping a drive.
- (Optional). Simple Network Management Protocol (SNMP) for SNMP trap forwarding by Array Manager.
- (Optional) On NetWare, Array Manager requires that the Dell OpenManage Server Agent (OMSA) 4.0.0 or later be installed on the managed system to provide thermal shutdown (when an enclosure becomes too hot or cold) for the Dell PowerVault 20xS, 21xS, and 22xS SCSI enclosures. If OMSA is not installed, Array Manager can provide other features of enclosure management, but not thermal shutdown without the risk of data loss.

Note The managed system's RAID controllers must be attached to storage in order to run Array Manager.

Note Dell does not offer NetWare in Japan.

Online Help Requirements

Internet Explorer version 4.0 or later is required for launching and viewing the online help from the Array Manager console. In addition to the online help, the Array Manager installation directory includes documentation in HTML and PDF format. The total space requirement for the online help and documentation is 15 MB.

Access Rights

Array Manager uses the host operating system's native security policies. These govern access rights to the Array Manager managed system. You need Administrator level access rights to install and use Array Manager.

Installation Considerations for a SAN

When installing an Array Manager managed system onto a Windows server connected to a Storage Area Network (SAN), you must take the following into consideration:

1. You can have at most one managed system (server) managing a Dell PowerVault 660F Fibre Channel RAID Array.

   During installation on a Windows server connected to a SAN or a PowerVault 660F, you will be asked whether this managed system will manage a PowerVault 660F storage system. Answer "Yes" for only one Windows server on each PowerVault 660F storage system. The only exception to this is if the server is part of an MSCS cluster. You can set up the Array Manager managed system on both servers of the MSCS cluster to manage the PowerVault 660F storage system, but only one server should be used to manage the RAID array at a time.

Note Each PowerVault 660F in a storage area network (SAN) should only be managed by one Array Manager managed system. If you configure more than one Array Manager managed system to manage a PowerVault 660F storage system, then each managed system may not be able to view all events generated by the fibre channel RAID controllers in the PowerVault 660F storage system.

   If you answer "Yes" to manage a PowerVault 660F storage system, you may be asked whether you want to manage it in a Direct Attach configuration. Answer "Yes" only if the following is true:

   - This server is directly attached to the PowerVault 660F and is not part of a SAN.
   - You intend to use the PowerVault 660F in a SAN and you don't need the LUN masking functionality of the Dell OpenManage Storage Consolidation software.

2. If the Array Manager managed system will not manage a Dell PowerVault 660F Array, it can still be installed onto a Windows server connected to the SAN to take advantage of the Online Volume Expansion features. See Working with Dynamic Volumes in the Volumes chapter of this document for more information.

3. If the Windows server contains PERC controllers and is connected to a SAN, Array Manager can be used to simultaneously manage both the SCSI and Fibre Channel RAID adapters and arrays.

Installing Array Manager on Windows

The Setup program takes you through the process of installing the Array Manager software. Before proceeding, you should exit all Microsoft Windows programs and log on with Administrator rights.

This section includes the following subtopics:

- Installing from the Dell OpenManage Systems Management CD
- Installing from the Dell OpenManage Array Manager CD

For other information on installation considerations, see the following:

- Installation Requirements for information on requirements for the Array Manager console and managed system.
- Understanding Array Manager on NetWare Installation for information on installing the Array Manager managed system on NetWare.
- Unattended (Silent) Installation for information on installing Array Manager on a remote machine.
- Installation Considerations for a SAN for information on installing Array Manager on a server connected to a Storage Area Network (SAN).

Installing from the *Dell OpenManage Systems Management* CD

Use the following procedure to install Array Manager using the Dell OpenManage Systems Management CD:

1. Insert the CD. Setup should begin automatically. If not, start setup by double-clicking on the **Setup** icon in the root directory of the CD.

2. Follow the instructions shown in the Setup wizard. You may need to use Custom Setup to install Array Manager.

3. Reboot the system when finished.

Installing from the *Dell OpenManage* Array Manager CD

Use the following procedure to install Array Manager using the Dell OpenManage Array Manager CD:

1. Insert the CD. Setup should begin automatically. If not, start setup by double-clicking on the **Setup** icon in the Setup directory of the CD.

2. Follow the instructions shown in the Setup wizard.

3. Reboot the system when finished.

Note When installing the managed system (server) on a Windows 2000 or Windows Server 2003 system, the Plug and Play Manager will activate after the reboot, indicating that it has discovered new hardware. You may be required to perform a second reboot if dynamic disks are present on your system. Otherwise, allow Plug and Play to register the new driver. (This occurs automatically.)

Understanding Array Manager on NetWare Installation

You can install the Array Manager managed system (server) on a NetWare server and then connect to the NetWare server from the Array Manager console.

Installing Array Manager on NetWare involves the following activities:

- **Installing the NetWare NWASPI.CDM driver.** In order for Array Manager to provide non-RAID SCSI support, the NetWare NWASPI.CDM driver needs to be installed. It is preferable to install the NWASPI.CDM driver before installing Array Manager. The NWASPI.CDM installation requires that the NetWare server be restarted. If you do not meet the minimum driver requirements when installing Array Manager, you will need to load the appropriate drivers after completing installation and then restart the NetWare server. You can avoid restarting the NetWare server twice by installing the NWASPI.CDM driver first. For more information on installing the NWASPI.CDM driver, see "Installing the NetWare NWASPI.CDM Driver."
- One of the following:
  - **When the install media is on a Windows system.** If you have the Array Manager installation media on a Windows system that has the Novell client, you can map a drive from the Windows system to a NetWare server. You can then copy the AMINST.NLM and AMINST.IPS files to the NetWare server and install Array Manager using the NetWare configuration utility. If you do not meet the minimum driver requirements when installing Array Manager, you will need to load the appropriate drivers after completing installation and then restart the NetWare server. For more information, see "Installing on NetWare Using a Mapped Drive from Windows."
  - **When the install media is on a CD.** If you have the Array Manager installation media on a CD, you can mount the CD as a NetWare volume. You can then install Array Manager using the NetWare configuration utility. If you do not meet the

minimum driver requirements when installing Array Manager, you will need to load the appropriate drivers after completing installation and then restart the NetWare server. For more information, see "Installing on NetWare Using an Installation CD."

Note Dell does not offer NetWare in Japan.

Installing the NetWare NWASPI.CDM Driver

Array Manager requires that the NWASPI.CDM driver be installed on NetWare in order to provide non-RAID SCSI support. The NWASPI.CDM driver can be optionally installed when installing the NetWare operating system. If you did not install the NWASPI.CDM driver when you installed NetWare, then you need to install it now. If you do not install the NWASPI.CDM driver, then Array Manager will still function on NetWare, but support for non-RAID SCSI controllers will not be available.

It is preferable to install the NWASPI.CDM driver before installing Array Manager on NetWare.

Depending on which version of NetWare you have, use one of the following procedures:

- Installing the NWASPI.CDM driver on NetWare 5.x
- Installing the NWASPI.CDM driver on NetWare 6.5

Note Dell does not offer NetWare in Japan.

Installing the NWASPI.CDM driver on NetWare 5.x

Use the NetWare configuration utility to install the NWASPI.CDM driver on NetWare 5.x.

1. Start the NetWare configuration utility by entering the following from the NetWare system console:

   nwconfig

2. Select **Product Options** and then **Install a product not listed**.

3. Specify the full path to the NWASPI.CDM file. This file may be located in the NWSERVER/drivers or NWSERVER directory. NetWare completes the installation for you.

4. Do one of the following:

- Continue with the Array Manager installation by either Installing on NetWare Using a Mapped Drive from Windows or "Installing on NetWare Using an Installation CD."

OR

- Restart the NetWare server by entering **restart server** at the NetWare system console command prompt. In order for the installation to complete successfully, the NetWare server must be restarted.

Installing the NWASPI.CDM driver on NetWare 6.5

Use the NetWare hdetect utility to install the NWASPI.CDM driver on NetWare 6.5.

1. Start the NetWare hardware configuration utility be entering the following from the NetWare system console:

   hdetect

   The **hdetect** utility displays a list of device types and drivers.

2. Select the drivers displayed next to the **Storage devices** option and press Enter.

   The **Storage devices** option is displayed under **Device types**. If **Storage devices** is not already displayed, select **Continue** until the **Storage devices** option is displayed.

3. Press the F3 function key to initiate Auto detect. The system locates and displays the available SCSI devices.

4. Select the SCSI devices and press the **Insert** key to display a list of available drivers.

5. Select the NWASPI.CDM from the list and press Enter.

6. Select **Return to Driver Summary**. The NWASPI.CDM driver should now be added to the list of devices.

7. Select **Continue** until the system saves the configuration and exits the **hdetect** utility.

8. Continue with the Array Manager installation by either Installing on NetWare Using a Mapped Drive from Windows or "Installing on NetWare Using an Installation CD."

Installing on NetWare Using a Mapped Drive from Windows

Use this procedure if the Array Manager install media is located on a Windows system. This would be the case, for example, if you have downloaded the Array Manager installation media from the Dell support site to a Windows system.

Installing the Array Manager managed system (server) on a NetWare server requires the AMINST.NLM and AMINST.IPS files. If these files are located on a Windows system, you can map a drive to the NetWare server and then copy the AMINST.NLM and AMINST.IPS files from the Windows system to the NetWare server. You would then go to the NetWare server and complete the installation.

Installing Array Manager using this method requires that the Novell client be installed on the Windows system from which you are mapping a drive to the NetWare server.

For more information on NetWare installation considerations, see the following:

- Managed System for NetWare Requirements (Server)
- Understanding Array Manager on NetWare Installation
- Uninstalling Array Manager on NetWare

Use the following steps to install the Array Manager managed system on the NetWare server from a mapped drive.

1. Install the NetWare NWASPI.CDM driver. You can install the NWASPI.CDM driver before or after installing Array Manager. The NWASPI.CDM installation requires that the NetWare server be restarted. If you do not meet the minimum driver requirements when installing Array Manager, you will need to load the appropriate drivers after completing installation and then restart the NetWare server. You can avoid restarting the NetWare server twice by installing the NWASPI.CDM driver before Array Manager. For more information, see "Installing the NetWare NWASPI.CDM Driver."

2. Insert the *Dell OpenManage Applications* CD or the Array Manager CD into the CDROM drive of the Windows system. Setup should begin automatically.

   OR

   Unzip the Array Manager installation package.

3. Map a drive from the Windows system to the NetWare server.

4. Locate the AMINST.NLM and the AMINST.IPS files on the Array Manager installation media.

   - If you have the *Dell OpenManage Applications* CD, then the AMINST.NLM and AMINST.IPS files are located in the following directory:

     \AM<n>\NetWare

     where <n> represents an Array Manager version number. For example, <n> might be specified as 35, 36, 37, 38 and so on.

   - If you have the Array Manager installation CD or if you unzipped the installation package, then the AMINST.NLM and AMINST.IPS files are located in the NETWARE subdirectory.

5. Copy the AMINST.NLM and AMINST.IPS files from the Windows system to the NetWare server. For example, you can copy the AMINST.NLM and AMINST.IPS files to the NetWare SYS:SYSTEM\ folder.

6. Start the NetWare configuration utility by entering the following from the NetWare system console:

   nwconfig

7. Select **Product Options** and then **Install a product not listed**.

8. Specify the full path to the folder where you copied the AMINST.NLM and AMINST.IPS files. NetWare completes the installation for you.

   The Array Manager installation displays messages indicating if you do not have the required drivers installed on the NetWare system. You may wish to observe the installation so as to take note of any required drives that need to be loaded. If your system has the required drivers, the installation completes with the following message:

   Install successful

   If your system does not have the required drivers, the installation completes with the following message:

   Install successful (Restart required)

   In this case, you must load the required drivers and restart the NetWare server. To restart the NetWare server, enter **restart server** at the NetWare system console command prompt

9. If you have not done so already, install the NWASPI.CDM driver. The NWASPI.CDM driver is required for Array Manager to provide non-RAID SCSI support. See Installing the NetWare NWASPI.CDM Driver for more information.

After installing the Array Manager managed system on a NetWare server, you can connect to the NetWare server from the Array Manager console. For more information on connecting to a managed system from the console, see "Connecting to a Remote System from the Console."

For information on starting the Array Manager service on NetWare, see "Starting the Array Manager Service for the First Time."

For information on enabling NetWare support for the PowerVault 660F, see "NetWare Support for the PowerVault 660F."

Note Dell does not offer NetWare in Japan.

Installing on NetWare Using an Installation CD

If you have an installation CD, you install the Array Manager managed system directly on the NetWare server.

Note If a version of Array Manager is already installed on the NetWare system, then uninstall that version before continuing with the install of a new version of Array Manager. See Uninstalling Array Manager on NetWare for information on uninstallation.

For more information on NetWare installation considerations, see the following:

- Managed System for NetWare Requirements (Server)
- Understanding Array Manager on NetWare Installation
- Uninstalling Array Manager on NetWare

Note Dell does not offer NetWare in Japan.

Use the following steps to install the Array Manager managed system (server) on the NetWare server using an installation CD.

1. Install the NetWare NWASPI.CDM driver. You can install the NWASPI.CDM driver before or after installing Array Manager. The NWASPI.CDM installation requires that the NetWare server be restarted. If you do not meet the minimum driver requirements when installing Array Manager, you will need to load the appropriate drivers after completing installation and then restart the NetWare server. You can avoid restarting the NetWare server twice by installing the NWASPI.CDM driver before installing Array Manager. For more information, see "Installing the NetWare NWASPI.CDM Driver."

2. Insert the Array Manager installation CD into the NetWare server's CDROM drive.

3. Mount the installation CD as a NetWare volume.

4. Locate the AMINST.NLM and AMINST.IPS files on the Array Manager installation CD.

   - If you have the *Dell OpenManage Applications* CD, then the AMINST.NLM and AMINST.IPS files are located in the following directory:

     <*VolumeName*>:\AM<n>\NetWare

     where <n> represents an Array Manager version number. For example, <n> might be specified as 35, 36, 37, 38 and so on.

- If you have the Array Manager installation CD, then the AMINST.NLM and AMINST.IPS files are located in the NETWARE subdirectory.

5. Start the NetWare configuration utility by entering the following from the NetWare system console:

   nwconfig

6. Select **Product Options** and then **Install a product not listed**.

7. Specify the full path to the directory containing the AMINST.NLM and AMINST.IPS files on the Array Manager installation CD. NetWare completes the installation for you.

   The Array Manager installation displays messages indicating if you do not have the required drivers installed on the NetWare system. You may wish to observe the installation so as to take note of any required drives that need to be loaded. If your system has the required drivers, the installation completes with the following message:

   Install successful

   If your system does not have the required drivers, the installation completes with the following message:

   Install successful (Restart required)

   In this case, you must load the required drivers and restart the NetWare server. To restart the NetWare server, enter **restart server** at the NetWare system console command prompt

8. If you have not done so already, install the NWASPI.CDM driver. The NWASPI.CDM driver is required for Array Manager to provide non-RAID SCSI support. See Installing the NetWare NWASPI.CDM Driver for more information.

After installing the Array Manager managed system on a NetWare server, you can connect to the NetWare server from the Array Manager console. For more information on connecting to a managed system from the console, see "Connecting to a Remote System from the Console."

For information on starting the Array Manager service on NetWare, see "Starting the Array Manager Service for the First Time."

For information on enabling NetWare support for the PowerVault 660F, see "NetWare Support for the PowerVault 660F."

Unattended (Silent) Installation

You can install Array Manager on a Windows remote machine without requiring user input on the remote machine. You can use this procedure, for example, to install Array Manager on several different machines from a single location.

When performing the unattended installation, you run an install script that contains the steps required to complete the installation on the remote machine. There are two possibilities for using the install script:

- **Create your own script.** You can create your own install script by completing the installation process on a local machine while recording your installation selections in the script. You can then use this script to run the installation on a remote server. You would use this method if you no longer have access to the Array Manager install media. If you are creating your own script, you will perform the following procedures to create the script and then run the uninstall on a remote machine:

  a. Creating a Script for Unattended Installation

  b. Performing an Unattended Installation

- **Use a prerecorded script.** You can use one of the prerecorded install scripts provided with the Array Manager installation media. Array Manager provides the following install scripts:

  - Typical.iss — Performs a typical installation on Windows 2000.
  - Typw2k3.iss — Performs a typical installation on Windows Server 2003.
  - Node.iss — Installs the managed system (server) only. On Windows 2000, the Node.iss installs the server without replacing the Windows 2000 disk and volume management.
  - Console.iss — Installs the console (client) only.
  - Volume.iss — Installs Array Manager disk and volume management which replaces the Windows 2000 disk and volume management. The Volume.iss option only applies to Windows 2000.

  If you are using a prerecorded script, you will perform the Performing an Unattended Installation procedure to complete the installation on a remote machine.

Creating a Script for Unattended Installation

1. Locate the Array Manager installation media. The installation media may be on a CDROM or copied to the local machine.

2. Click the **Start** menu and select **Run...** to display the **Run** dialog box.

3. Enter **cmd** in the **Open** field to open a command shell.

4. Enter the following to install Array Manager on your local machine. This installation is required so that you can record the install options in a script.

   <*Setup.exe*> -r

   where:

   <*Setup.exe*> is the full directory path to the setup.exe script.

5. Complete the installation process by choosing the install options that are appropriate for installing on the remote server. As you complete the installation, the options you choose are recorded in the setup.iss script which is located in the directory specified by the %SYSTEMROOT% environment variable. To see which directory is specified by the %SYSTEMROOT% environment variable, enter the following:

   echo %SYSTEMROOT%

6. Save the setup.iss script in the %SYSTEMROOT% folder to a convenient location. You will specify this location when running the unattended installation. To complete the unattended installation, continue with the Performing an Unattended Installation procedure.

Performing an Unattended Installation

Use this procedure to perform an unattended installation. This procedure assumes that you either have access to the Array Manager install media, or that you have created your own install script.

1. Locate the setup.exe and the install scripts. You can find these scripts at the following locations:

   - **setup.exe** — This script is provided with the Array Manager install media. If you have already installed Array Manager on the local machine and no longer have the install media, you can find a copy of this script at the location specified by the following Registry entry:

   ```
   HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\
   Uninstall\{4536C827-30C2-11D4-A22F-00A0C9EA31FB}\UninstallString
   ```

   - **install script** — You can use the install scripts provided with the Array Manager install media or use a script that you created yourself using the Creating a Script for Unattended Installation procedure. Array Manager provides the following install scripts with the installation media:

     - Setup \ AMcheckSDK \ Typical.iss

- Setup\AMcheckSDK\Typw2k3.iss
- Setup\AMcheckSDK\Node.iss
- Setup\AMcheckSDK\Console.iss
- Setup\AMcheckSDK\Volume.iss

2. Click the **Start** menu and select **Run...** to display the **Run** dialog box.

3. Enter **cmd** in the **Open** field to open a command shell.

4. Log into the remote server on which you wish to install Array Manager by entering the following at the command prompt:

   telnet <*IP address*>

   where <*IP address*> is the IP address of the remote server.

5. Enter the following to initiate the unattended installation:

   <*Setup.exe*> -s -f1"<*InstallScript.iss*>"

   where:

   <*Setup.exe*> and <*InstallScript.iss*> are the full directory paths to the setup.exe and install script at the location you identified in step 1 of this procedure. The <*InstallScript.iss*> should be enclosed in quotation marks. There is no space between -f1 and "<*InstallScript.iss*>".

6. Reboot the system on which you have installed Array Manager.

Changing Installed Components

After you have installed Array Manager, you can modify, repair, or remove components of the installation through the Windows Control Panel. This section covers changes to all installed components except uninstalling Array Manager from the NetWare system, which is covered in the next section.

Note To make changes in the installation, you must be logged on with Administrator rights and you must exit the Array Manager application.

Note If you uninstall Array Manager or modify, repair, or remove installation components, you may notice that events generated before the uninstall, modify, repair, or remove no longer display an event description. This situation occurs when the component that generated the event is no longer installed on the system.

If you wish to preserve the events and their descriptions, you should back up or export the event log before you uninstall Array Manager, or modify, repair, or remove installation components.

To change the Array Manager installed components:

1. Open the Control Panel.
2. Double-click **Add/Remove Programs**.
3. Click **Dell OpenManage Array Manager** unless it is already selected.
4. Click the **Add/Remove** or **Change/Remove** button.
5. A dialog box appears with three options: **Modify**, **Repair**, **Remove**. Choose an option as described below.

   - **Modify**

     This option allows you to change the components that are installed on your system. Checked components will be installed. Unchecked components will not be installed and will be removed if they are currently installed.

Caution If you do not select a component that you currently have installed, this component will be removed.

   - **Repair**

     This option automatically detects which components you currently have installed on your system and reinstalls those components.

   - **Remove**

     This option completely removes Array Manager from your Windows system. For uninstalling Array Manager from a NetWare system, see Uninstalling Array Manager on NetWare for more information.

Note Files that were in use during the uninstall will not be deleted until you reboot the server.

Upgrading Array Manager

You may wish to upgrade an Array Manager installation to a newer version of Array Manager. When doing so, it is recommended that you uninstall the old version of Array Manager before installing the new one. See Uninstalling Array Manager for more information.

It is also recommended that you uninstall Array Manager before upgrading the Windows operating system. See Upgrading the Windows Operating System for more information.

Note If you uninstall Array Manager or modify, repair, or remove installation components, you may notice that events generated before the uninstall, modify, repair, or remove no longer display an event description. This situation occurs when the component that generated the event is no longer installed on the system. You can clear the event log to remove these events. See Clearing Events. If you wish to preserve the events and their descriptions, you should back up or export the event log before you uninstall Array Manager, or modify, repair, or remove installation components. See Backing Up and Restoring Events and Exporting Events to a Text File.

Upgrading the Windows Operating System

The Array Manager installation process installs files and makes registry entries on the server that are specific to the operating system. For this reason, changing the operating system can disable Array Manager. To avoid this problem, you should uninstall Array Manager before upgrading the Windows operating system. See Uninstalling Array Manager for more information.

After the upgrade, reinstall Array Manager using the Array Manager install media or by downloading Array Manager from the Dell support site at http://support.dell.com.

If you have already upgraded without uninstalling Array Manager and are now experiencing problems, see "Resolving Windows Upgrade Problems."

Uninstalling Array Manager

This section discusses uninstalling Array Manager.

Note If you uninstall Array Manager or modify, repair, or remove installation components, you may notice that events generated before the uninstall, modify, repair, or remove no longer display an event description. This situation occurs when the component that generated the event is no longer installed on the system. You can clear the event log to remove these events. See Clearing Events. If you wish

to preserve the events and their descriptions, you should back up or export the event log before you uninstall Array Manager, or modify, repair, or remove installation components. See Backing Up and Restoring Events and Exporting Events to a Text File.

Uninstalling Array Manager on Windows

There are two methods for uninstalling Array Manager on Windows.

- Use the Windows Add/Remove Programs utility. Refer to the Windows documentation for more information.
- Run the Array Manager setup.exe installation executable. If Array Manager is already installed on the system, the installation process provides an uninstall option.

Note Array Manager appears in the Add/Remove Programs list, whether it was installed using Array Manager or Server Administrator.

Uninstalling Array Manager on NetWare

Use the following procedure to uninstall Array Manager from a NetWare system.

1. Start the NetWare configuration utility by entering the following from the NetWare system console:

   nwconfig

2. Select **Product Options** and then **View/Configure/Remove installed products**. This action displays the **Currently Installed Products** screen.

   - If Array Manager is on the list, then Array Manager was installed directly on NetWare.
   - If Server Administrator is on the list, then Server Administrator was used to install Array Manager on NetWare.

3. Do one of the following:

   - If Array Manager is on the product list, select Array Manager, then press the Delete key.
   - If the product list does not display Array Manager but displays Server Administrator instead, then Array Manager was installed using the Server Administrator install process. In this case, you need to use the Server Administrator uninstall process to uninstall Array Manager.

Note Dell does not offer NetWare in Japan.

Unattended (Silent) Uninstall

If you need to uninstall Array Manager from more than one machine, you can do so using an unattended or "silent" uninstall. Unattended uninstall enables you to uninstall Array Manager from several remote machines at the same time without requiring user input.

To Run the Unattended Uninstall

1. Locate the Array Manager setup.exe and uninstall.iss scripts on a local machine that already has Array Manager installed. The setup.exe and uninstall.iss scripts are provided with the Array Manager install media. After installing Array Manager, the setup.exe and uninstall.iss scripts reside on the system at the location specified by the following Registry entry:

   ```
   HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\
   Uninstall\{4536C827-30C2-11D4-A22F-00A0C9EA31FB}\UninstallString
   ```

2. Click the **Start** menu and select **Run...** to display the **Run** dialog box.

3. Enter **cmd** in the **Open** field to open a command shell.

4. Log into the remote machine on which you wish to uninstall Array Manager by entering the following at the command prompt:

   telnet <*IP address*>

   where <*IP address*> is the IP address of the remote machine.

5. Enter the following to initiate the unattended uninstall:

   <*setup.exe*> -s /u -f1"<*uninstall.iss*>"

   where:

   <*setup.exe*> and <*uninstall.iss*> are the full directory paths to the setup.exe and uninstall.iss scripts at the location you identified in step 1 of this procedure. The <*uninstall.iss*> should be enclosed in quotation marks. There is no space between the -f1 and "<*uninstall.iss*>".

   The <*uninstall.iss*> script uninstalls Array Manager from the remote machine.

6. Reboot the system on which you have uninstalled Array Manager.

Starting Array Manager 3

In order to use Array Manager, you must have an open console on the management station and the Array Manager service must be running on the managed system. The following sections describe how to start the service and open the console.

- Starting the Array Manager Service for the First Time
- Starting and Stopping the Array Manager Service using the Array Manager Utilities
- Starting the Console

Starting the Array Manager Service for the First Time

This section describes how to start the Array Manager service for both Windows and NetWare systems.

On Windows

The Array Manager service starts automatically each time you boot the machine. No other action is necessary.

On NetWare

The Array Manager installation adds the Array Manager folder to the system search path. The installation also adds the Array Manager service to the **autoexec.ncf** file so that the service starts automatically each time you boot the machine.

If you want to start the Array Manager service without rebooting, however, use the following procedure to start the service manually:

1. Add the Array Manager folder to the search path by entering:

   Search Add SYS:*InstallationDirectory*

   where *InstallationDirectory* is the directory in which you installed Array Manager. The default installation directory is **SYS:\system\dell\arraymgr**.

2. Start the Array Manager service by entering:

arraymgr

For further information on the Array Manager service, see the section "Starting and Stopping the Array Manager Service" in the "Array Manager Utilities" chapter.

Starting the Console

The Array Manager console can be launched on a Windows system using either the **Start** menu or the command line.

See the following sections for starting the Array Manager console:

- "Starting Array Manager from the Start Menu" for information on starting Array Manager on a local Windows system.
- "Starting Array Manager from the Command Line" for information on using the Windows command line to launch Array Manager on the local Windows system.
- "Starting Array Manager with a Controller Highlighted" for information on launching Array Manager with a particular controller highlighted in the tree view. This procedure may be useful if you wish to quickly identify the status of a particular controller.
- "Starting Array Manager with a Remote Connection" for information on using the Windows command line to launch Array Manager and establish a connection to a remote system. This procedure may be useful for verifying the status of storage connected to a remote system.
- "Starting Array Manager from Computer Management" for information on starting Array Manager from the Computer Management window.
- "MMC Snap-in" for information on running Array Manager from the Microsoft Management Console (MMC).

Starting Array Manager from the Start Menu

You can start Array Manager on a Windows system using the **Start** menu. The Array Manager console launches with storage connected to the local system displayed in the tree view.

To start Array Manager from the Windows **Start** menu:

1. Click the **Start** menu.
2. Select **Programs | Dell OpenManage Applications | Array Manager | Array Manager Console**.

You may also wish to view the following related sections:

- "Starting Array Manager from the Command Line" for information on using the amconsole command to start Array Manager from the command line.
- "Starting Array Manager with a Controller Highlighted" if you wish a particular controller to be highlighted in the tree view. If the controller is attached to a remote system, see "Starting Array Manager with a Remote Connection."
- "Remote Management" for more information on connecting to remote systems.
- "Array Manager Utilities" for information on using the utilities to view or modify the Array Manager configuration.
- "MMC Snap-in" for information on installing and starting the MMC Snap-in.

Starting Array Manager from the Command Line

You can start Array Manager on a Windows system using the command line. The Array Manager console launches with storage connected to the local system displayed in the tree view.

To start Array Manager from a command line:

1. Open a DOS command shell.
2. Enter the following as a single command string:

   amconsole $user:*UserName* password:*Password*

   where:

   - *UserName* is the name of a user with administrator authority.
   - *Password* is the password of a user with administrator authority.

You may also wish to view the following related sections:

- "Starting Array Manager from the Start Menu" for information on starting Array Manager on a local system from the Windows Start menu.
- "Starting Array Manager with a Controller Highlighted" if you wish a particular controller to be highlighted in the tree view. If the controller is attached to a remote system, see "Starting Array Manager with a Remote Connection."
- "Remote Management" for more information on connecting to remote systems.
- "MMC Snap-in" for information on starting the MMC Snap-in.
- "amcli Command" for more information on using the amcli command.

Starting Array Manager with a Controller Highlighted

When starting Array Manager from the command line, you can specify that a particular controller be highlighted in the tree view. Highlighting a controller may be useful when the system is attached to multiple controllers and you only need the status of a particular one. For example, you may have received an error message regarding a controller. You might also be using another management application that has identified a problem with a controller. In these cases, you may wish to launch Array Manager with the controller highlighted so that you can quickly obtain more information regarding the controller's status.

To start Array Manager with a controller highlighted:

1. Open a DOS command shell.

2. Enter the following as a single command string:

   amconsole $device:"*ControllerName*" user:*UserName* password:*Password*

   or if the controller is attached to a remote system, then enter:

   amconsole $computer:*ComputerName* user:*UserName* password:*Password* domain:*DomainName* device:"*ControllerName*"

   where:

   - *ComputerName* is the name of the remote computer.
   - *ControllerName* is the controller that you wish to be highlighted in the Array Manager tree view.
   - *DomainName* is the domain to which the remote computer belongs. When connecting to a NetWare system, omit this parameter. It does not apply to NetWare.
   - *Password* is the password of a user with administrator authority. When connecting to a remote system with Array Manager, it is recommended that the password be the same on both the local and remote system. When connecting to a NetWare system, omit the *Password* parameter. NetWare will prompt you for the password.
   - *UserName* is the name of a user with administrator authority. When connecting to a remote system with Array Manager, it is recommended that the user name be the same on both the local and remote system. When connecting to a NetWare system, omit the *UserName* parameter. NetWare will prompt you for the user name.

Elements of the controller name should be separated by a space. When using spaces, enclose the item containing spaces in quotation marks. For example, to start Array Manager with a PERC 2 /Si controller selected in the tree view, enter:

amconsole $device:"perc 2/si controller 0" user:*UserName* password:*Password*

You may also wish to view the following related sections:

- "Starting Array Manager from the Start Menu" for information on starting Array Manager on a local system from the Windows Start menu.
- "Remote Management" for more information on connecting to remote systems.
- "MMC Snap-in" for information on starting the MMC Snap-in.
- "amcli Command" for more information on using the amcli command.

Starting Array Manager with a Remote Connection

When launching Array Manager with a remote connection, the Array Manager tree view displays the local system and the remote system with their attached storage.

To start Array Manager with a remote connection:

1. Open a DOS command shell.
2. Enter the following as a single command string:

   amconsole $computer:*ComputerName* user:*UserName* password:*Password* domain:*DomainName*

   where:

   - *ComputerName* is the name of the remote computer.
   - *DomainName* is the domain to which the remote computer belongs. When connecting to a NetWare system, omit this parameter. It does not apply to NetWare.
   - *Password* is the password of a user with administrator authority. When connecting to a remote system with Array Manager, it is recommended that the password be the same on both the local and remote system. When connecting to a NetWare system, omit the *Password* parameter. NetWare will prompt you for the password.
   - *UserName* is the name of a user with administrator authority. When connecting to a remote system with Array Manager, it is recommended that the user name be the same on both the local and remote system. When connecting to a NetWare system, omit the *UserName* parameter. NetWare will prompt you for the user name.

You may also wish to view the following related sections:

- "Starting Array Manager from the Start Menu" for information on starting Array Manager on a local system from the Windows Start menu.
- "Remote Management" for more information on connecting to remote systems.

- "MMC Snap-in" for information on starting the MMC Snap-in.
- "amcli Command" for more information on using the amcli command.

Starting Array Manager from Computer Management

In Windows 2000 and Windows Server 2003, Array Manager is available as one of the programs or utilities in the Computer Management window.

To access the Computer Management from the Windows Start menu:

1. Click **Start**
2. Select **Settings | Control Panel | Administrative Tools | Computer Management**.

Using the Console 4

The Array Manager console displays the physical and logical storage attached to the system. Using the console, you can identify the storage components and their relationships. For example, you can identify the system's storage controllers, the virtual disks on each controller, and the array disks included in each virtual disk.

You can manage the storage from the Array Manage console using convenient wizards and menus.

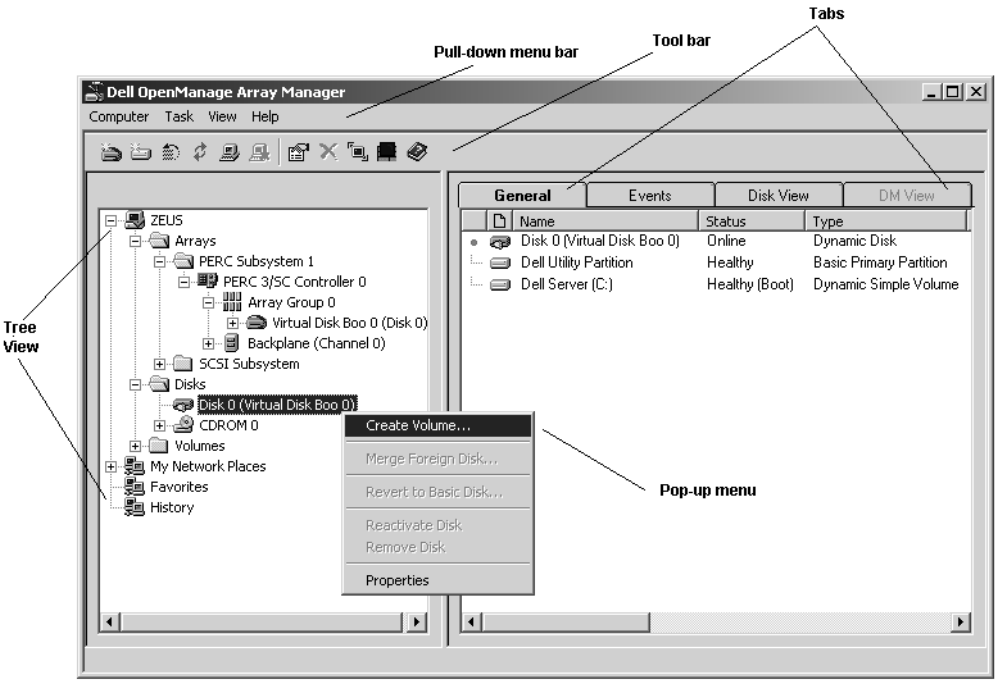
The console provides:

- A display of physical and logical storage devices attached to the system. These storage devices are displayed in the left pane of the GUI using an expandable tree view. See "Tree View" for more information.
- Pop-up menus to perform different functions. These menus are displayed by right-clicking on an object in the tree view.
- Pull-down menus and an icon bar for implementing functions.
- Wizards to assist you with common storage tasks. See "Wizards and Common Tasks" for more information.
- A property page to display information related to each storage device.
- Error icons for quick identification of an error condition. See "Error Icons" for more information.
- A property page to display alert notifications and system events. See "Events Tab" and "Event Monitoring" for more information.
- Tabs displayed in the right pane of the GUI that contain additional information on storage devices. See "Tabs" for more information.

Note For NetWare® users, the Array Manager console provides a graphical way to see other NetWare servers and storage objects. Because Array Manager does not support disk and volume management on NetWare servers, you will not see the Disks and Volumes storage objects in the console view of those servers.

Note Array Manager does not support disk and volume management on Windows Server 2003. For this reason, the Disks and Volumes storage objects are not displayed when running the console on Windows Server 2003. For information on accessing the native Windows Server 2003 Disk Management utility, see "Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux."

Array Manager Console



Tree View

The tree view displays the systems that Array Manager is currently connected to and the storage devices attached to those systems. Each system or storage device is represented as an icon or "object" in the tree view. These objects are displayed in a hierarchical order. A computer object representing a local or remote system is at the top of the hierarchy. By

clicking the button (+) to the left of each object, you can expand or contract the tree view to display or hide lower-level objects. When you launch Array Manager, the computer object is expanded by default to display the Arrays, Disks, and Volumes objects.

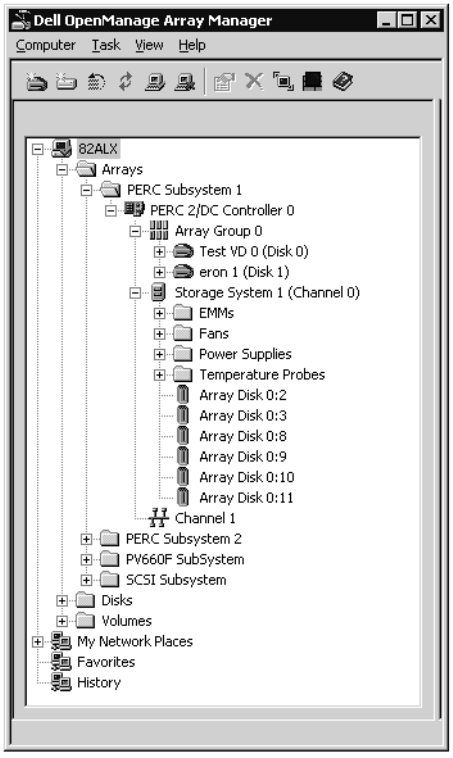
The tree view only displays those storage objects attached to your system. For example, if your system does not have an enclosure, then there is no enclosure object displayed in the tree view. In addition, different controllers support different features. For this reason, the storage objects in the tree view and their pop-up menu items can vary.

You can right-click any storage object to display its pop-up menu.

You can change the objects displayed in the tree view. For more information, see "Configuring the Console" or "Changing Tree View Storage Objects."

Note Array Manager does not support disk and volume management on Windows Server 2003. For this reason, the Disks and Volumes storage objects are not displayed when running the console on Windows Server 2003. For information on accessing the native Windows Server 2003 Disk Management utility, see "Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux."

Tree View



My Network Places

The **My Network Places** icon is displayed in the tree view. This icon functions the same as the **My Network Places** icon in the My Computer window of a Windows operating system.

History

The **History** icon is displayed in the tree view. Array Manager uses **History** to store a list of the computers to which you have recently connected. If you want to delete a computer from the History list, right-click the computer name and select **Remove Computer** from the pop-up menu.

Favorites

The **Favorites** icon is displayed in the tree view. You can right-click the **Favorites** icon to quickly connect to frequently used remote systems. The **Favorites** option lets you display and retrieve NetWare server names, as well as other computers.

Tabs

The right pane of the Array Manager console displays tabs. Each tab provides information on the configuration or status of your storage system. The tabs displayed in the console vary depending on which operating system you are running Array Manager. In addition, on some operating systems, a tab may be inactive (greyed out) if the operating system does not support the features displayed on that tab.

The tabs that may be displayed in the console are as follows:

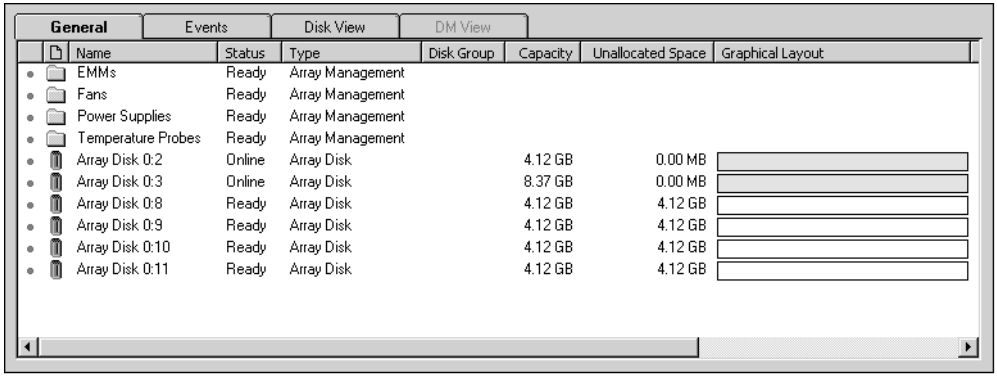
- General Tab
- Events Tab
- Disk View Tab
- DM View Tab

General Tab

The General tab provides information on the configuration and status of the storage object selected in the tree view. The column headers and information displayed on the General tab vary depending on the type of storage object selected.

Right-clicking items on the General tab displays a pop-up menu for implementing functions or configuring the tab view. For example, you can right-click an array disk and then select Blink to blink the LED display on the array disk. You can also change the sort order of the columns or hide a column by right-clicking the column header and making selections from the pop-up menu.

General Tab. This tab displays different information depending on which storage object is selected in the tree view.



Information displayed on the General tab includes the following:

- **Name** — The name of the object.
- **Status** — This can vary, depending on the object. Common Status conditions are Online, Healthy, and Resynching. If Status shows as Failed, refer to the "Array and Hard Disk Status Information" in the "Troubleshooting" chapter.
- **Layout** — Identifies the layout of the selected object. For example, if a volume is selected, the volume layout may indicate whether the volume is partitioned or not.
- **Type** — Identifies the object, such as Array Disk.
- **Disk Group** — Shows an entry for disks in a basic or dynamic group.
- **Capacity** — The maximum size of the disk.
- **Unallocated Space** — The amount of free disk space still available.
- **Graphical Layout** — How much of the disk is being used.
- **Progress** — The current progress (entage of completion) for tasks.
- **File System** — Indicates the type of file system. For example, a selected volume might have a FAT (File Allocation Table) or an NTFS (NT File System) file system.
- **Device** — The type of disk: SCSI, IDE, and so on. See "Display of IDE Devices" for more information on how Array Manager displays IDE devices.
- **Port** — Identifies the controller card. A SCSI port has zero or more target IDs, and a target ID has one or more LUNs.
- **LUN** — Logical Unit Number.
- **Target** — SCSI ID, which uniquely identifies the disk on the controller card.

- **Vendor** — Identifies vendor on hardware objects.

Events Tab

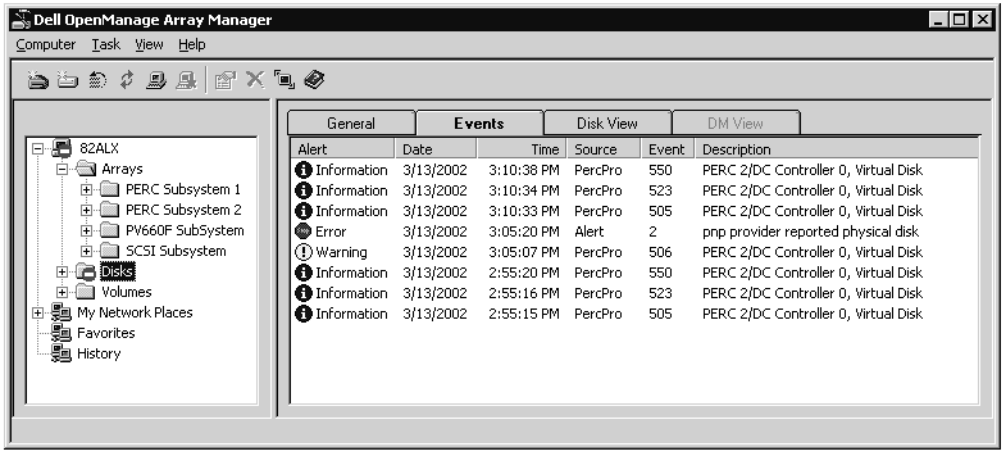
The Events tab displays events generated by controllers in response to system activity. These events are displayed on the Array Manager Events tab and added to the Windows application event log.

You can use these events for troubleshooting. For more information on the Events tab, see "Viewing Event Listings."

For more information on event monitoring and troubleshooting, see the following:

- Event Monitoring
- Troubleshooting

Events Tab

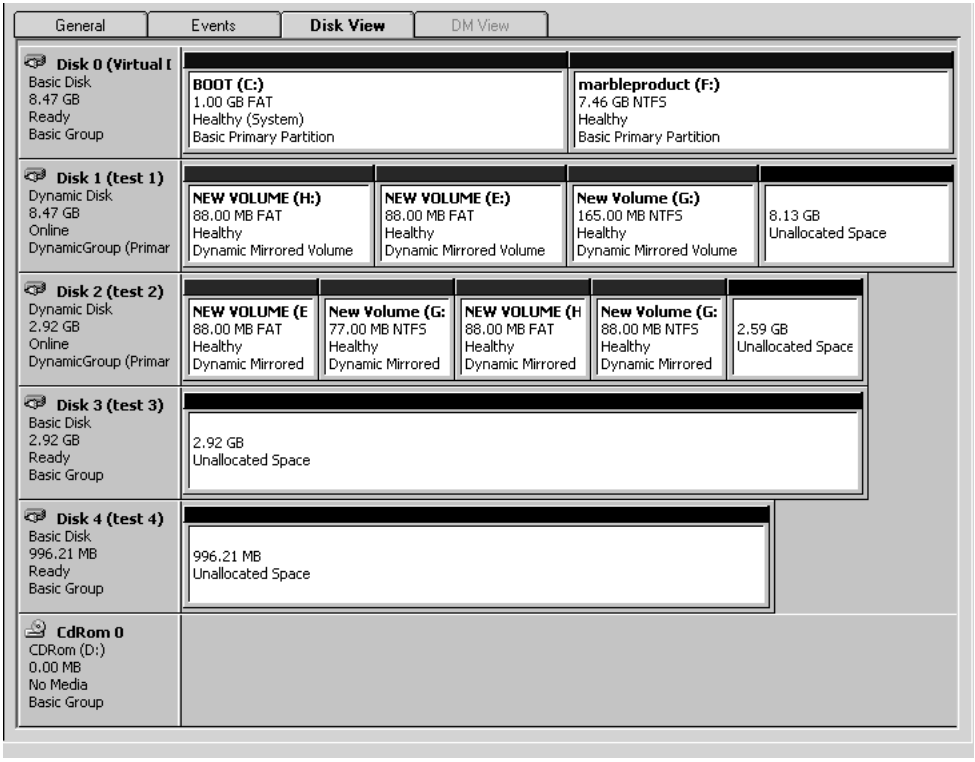


Disk View Tab

Similar to Windows Disk Management, the Disk View tab displays a graphical view of your system's disks, CDROMs, and other removable media. Right-clicking items on the Disk View tab displays a pop-up menu for implementing functions. For example, you can right-click a virtual disk and then select Create Volume to launch the Create Volume wizard from the Disk View tab.

Note On Windows Server 2003, the Disk View tab is inactive (greyed out).

Disk View Tab



DM View Tab

The DM View tab is displayed in the Array Manager console on Windows 2000. The DM View tab displays the information for the Windows 2000 Disk Management snap-in, which is also called LDM (Logical Disk Manager). Disk Management is the built-in disk and volume manager that comes with Windows 2000.

The DM View tab is inactive (dimmed) in an Array Manager console when it is connected to a local or remote Windows 2000 system. This is because Array Manager is an upgrade that replaces Disk Management during installation on a Windows 2000 system. For details, see the section "Replacing Windows 2000 Disk Management."

The reason that the DM View tab is displayed is because you can remotely connect to a server with Disk Management. Once you make the connection to the server with Disk Management, the DM View becomes active. Because Array Manager and Disk

Management are related products, you can connect to the server with Disk Management and view and manage its storage. Refer to the Windows 2000 Disk Management online help for details on how the function works.

Configuring the Console

To customize the Array Manager console display, you can add or remove categories, customize headers, resize columns, and change the colors in the graphical layout. Topics in this section include:

- Changing Tree View Storage Objects
- Customize Headers
- Set Graphical Layout Color Preferences
- Hide Unknown Windows Disks at Startup

Changing Tree View Storage Objects

The default Array Manager tree view shows three categories: Arrays, Disks, and Volumes. (NetWare and Windows Server 2003 users will see only Arrays.) You can add new categories to this tree view. For example, you may want Array Groups, Basic Groups, and Dynamic Groups to be displayed under the same object in the tree view. You can add a Groups category and specify that Array Groups, Basic Groups, and Dynamic Groups be included in this category.

To Create a Category

1. Right-click the computer icon and select **Add Categories** from the context menu.
2. Type a category name.
3. Use **Add** and **Remove** to set the categories to appear under the tree view. You can enter a unique category name or choose one from the listed items.
4. Click **OK** to finish or click **Cancel** to cancel any changes. You do not have to restart to complete this process.

   The new category appears in the left pane of the Array Manager console.

Customize Headers

This function allows you to choose the headers you want to see and place them in the order in which you would like to see them.

To Customize Headers

1. Highlight an object in the left pane of the console.
2. In the right pane, right-click in one of the column headings, such as **Name**.
3. Select **Customize header**.

   A dialog box comes up with two panes:

   - **Hide columns** on the left
   - **Visible columns and order** on the right

   Add and Remove buttons are provided to allow you to move the column headers provided by the program between each pane.
4. Use the Add or Remove buttons to move the column headers you want to the **Visible columns and order** pane. Once the desired headers are moved to the right pane, you can change their order with the Move Up and Move Down buttons.

Note You can also change the order of the column headers by dragging a header to the right or the left in the General tab view.

Set Graphical Layout Color Preferences

The color preference option allows you to select how you want to view particular items in the graphical layout, such as RAID-5 volumes, mirrored disks, and other managed objects.

To Change Colors in the Graphical Layout

1. Right-click the computer icon and choose **Settings**.
2. Select an object to change, such as Primary Partition. This places that item in the text area of the dialog box.
3. Use the drop-down Color menu on the right side to assign a color to that item.

4. Repeat these steps to change other objects.

5. Click **OK** to finish.

Rescan at Startup

At the bottom of the Quick Access window is a **Rescan at startup** check box. This check box is unchecked (deselected) by default. When checked (selected) the **Rescan at startup** check box indicates that Array Manager will perform a rescan when the Array Manager Console is launched. Only the first computer system in the Array Manager tree view is rescanned.

You can initiate a rescan on any computer system in the tree view by doing the following:

1. Select the computer system you want to rescan.

2. Select **Rescan** from the **View** pull-down menu.

Note When Array Manager is used to manage controllers that support Microsoft Cluster Server, it is possible that a resource change may have occurred while the Array Manager console was not running. It is therefore suggested that you check (select) the **Rescan at startup** check box on systems with cluster-enabled array controllers.

Hide Unknown Windows Disks at Startup

At the bottom of the Quick Access window is a **Hide unknown Windows disks at startup** check box. This check box is checked (selected) by default.

The array manager console displays the disks available to the Windows operating system for read and write purposes under the Disks object in the tree view. These disks can be both array disks and virtual disks. When a disk is in an unknown state, it is no longer available to the operating system.

By default, disks that are in an unknown state are not displayed in the console. If you wish the offline disks to be displayed in the console, deselect the **Hide unknown Windows disks at startup** check box.

When selecting or deselecting the **Hide unknown Windows disks at startup** check box, you need to restart the console in order for the new setting to take effect.

Note The **Hide unknown Windows disks at startup** feature is not available on NetWare and Windows Server 2003.

Display of IDE Devices

A limitation of the Windows operating system is that the only IDE devices recognized on the system as IDE devices are the two built-in IDE controllers and the four IDE drives that may be attached to these controllers. Any additional IDE controllers or IDE devices attached to a RAID controller are displayed by Windows as a SCSI device and not an IDE device. For this reason, you may notice a discrepancy in device type when viewing the Device property on the Array Manager General tab and the Device Type property displayed by right-clicking a disk object and selecting Properties. The General tab may display IDE as the device whereas the disk properties may display SCSI as the device type. This discrepancy occurs because the General tab receives information from the controller while the disk properties receive information from Windows.

Understanding Storage Concepts 5

Array Manager provides storage management using RAID (Redundant Array of Independent Disks) technology. Understanding storage management requires an understanding of RAID concepts, as well as some familiarity with how your system's RAID controllers and operating system view disk space.

What Is RAID?

RAID (Redundant Array of Independent Disks) is a technology for managing how data is stored on the physical or hard disks that reside in your system or are attached to it. A key aspect of RAID is the ability to span hard disks so that the combined capacity of multiple disks can be treated as a single, extended chunk of disk space. Another key aspect of RAID is the ability to maintain redundant data which can be used to restore data onto a disk in the event of a disk failure. RAID uses different techniques, such as striping, mirroring, and parity, to store and reconstruct data. There are different RAID levels that use different methods for storing and reconstructing data. The RAID levels have different characteristics in terms of read / write performance, data protection, and storage capacity. Not all RAID levels maintain redundant data, which means for some RAID levels lost data cannot be restored. Which RAID level you choose depends on whether your priority is performance, protection, or storage capacity.

Note The RAID Advisory Board (RAB) defines the specifications used to implement RAID. Although the RAID Advisory Board (RAB) defines the RAID levels, commercial implementation of RAID levels by different vendors may vary from the actual RAID specifications. An implementation used by a particular vendor may affect the read and write performance and the degree of data redundancy.

Hardware and Software RAID

RAID can be implemented with either hardware or software. A system using hardware RAID has a RAID controller that implements the RAID levels and processes data reads and writes to the array disks. When using software RAID, the operating system must implement the RAID levels. For this reason, using software RAID by itself can slow system performance. You can, however, use software RAID on top of hardware RAID

volumes to provide greater performance and variety in the configuration of RAID volumes. For example, you can mirror a pair of hardware RAID-5 volumes across two RAID controllers to provide RAID controller redundancy.

Note Careful consideration should be given to a decision between hardware RAID and software RAID. Software RAID generally has lower performance; and when it is used with SAN Fibre Channel storage, the volumes cannot be moved from server to server.

RAID Concepts

RAID uses particular techniques for writing data to disks. These techniques enable RAID to provide data redundancy or better performance. These techniques include:

- **Mirroring** — Duplicating data from one disk to another disk. Mirroring provides data redundancy by maintaining two copies of the same data on different disks. If one of the disks in the mirror fails, the system can continue to operate using the unaffected disk. Both sides of the mirror contain the same data at all times. Either drive can act as the operational drive. A mirrored RAID group is comparable in performance to a RAID-5 group in read operations but faster in write operations.
- **Striping** — Disk striping writes data across several disks. Typically, each "stripe" spans 3 or more disks but only consumes a portion of each disk. Each disk, therefore, may have several stripes. The amount of space consumed by a stripe is the same on each disk included in the stripe. The portion of a stripe that resides on a single disk is a "strip" or "chunk." Striping by itself does not provide data redundancy. Striping in combination with parity does provide data redundancy.
- **Stripe size or stripe depth** — The amount of space that each stripe consumes on a single disk.
- **Parity** — Parity refers to redundant data that is maintained using an algorithm in combination with striping. When one of the striped disks fails, the data can be reconstructed from the parity information using the algorithm.
- **Spanning or Concatenation** — Although concatenation is not considered a RAID level, Array Manager enables you to concatenate (or span) multiple disks so that the combined disk space can be treated as a single, large disk. RAID levels that use striping also employ the concept of spanning, since each stripe uses space from multiple disks. Spanned disks can also be mirrored.
- **Consistency Check** — Verifies the accuracy of redundant data in a RAID group. For example, in a system with parity, checking consistency computes the parity of the data drives and compares the results with the contents of the parity drive.
- **Hot Spare** — An extra, unused disk drive that is part of an array disk subsystem. A hot spare is always in standby mode. If a disk failure occurs, the hot spare replaces the failed drive without interrupting the system.

- **Disk Rebuild** — A rebuild is the recreating of data that existed on a drive before it failed. You can rebuild data only in arrays with data redundancy, such as RAID level 1. (See "Choosing RAID Levels and Concatenation" for more information about RAID levels.)
- **Hot Swapping** — A hot swap is the manual replacement of a defective physical disk unit while the computer is still running. When a new drive has been installed, you must issue a command to rebuild the drive. Before removing a drive, be sure to follow the steps required to remove the drive safely and avoid possible data loss. To remove a drive from a PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, and 4e/Di controller see "Offline Array Disk" and "Replacing a Failed Disk." To remove a disk from a PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controller, see "Prepare to Remove." To remove a disk from a PV660F storage system, see "Prepare to Remove."
- **Fault Tolerance** — Includes features like automatic failed drive detection, rebuild, and hot swapping.

RAID Levels

Each RAID level uses some combination of mirroring, striping, and parity to provide data redundancy or improved read and write performance. For specific information on each RAID level, see "Choosing RAID Levels and Concatenation."

Organizing Data Storage for Availability and Performance

RAID provides different methods or RAID levels for organizing the disk storage. Some RAID levels maintain redundant data so that you can restore data after a disk failure. Different RAID levels may also entail an increase or decrease in the system's I/O (read and write) performance.

Maintaining redundant data requires the use of additional array disks. As more disks become involved, the likelihood of a disk failure increases. Because of the differences in I/O performance and redundancy, one RAID level may be more appropriate than another based on the applications in the operating environment and the nature of the data being stored.

When choosing concatenation or a RAID level, the following performance and cost considerations apply:

- **Availability or fault-tolerance.** Availability or fault-tolerance refers to a system's ability to maintain operations and provide access to data even when one of its components has failed. In RAID volumes, availability or fault-tolerance is achieved by maintaining redundant data. Redundant data includes mirrors (duplicate data) and parity information (reconstructing data using an algorithm).

- **Performance.** Read and write performance can be increased or decreased depending on the RAID level you choose. Some RAID levels may be more appropriate for particular applications.
- **Cost efficiency.** Maintaining the redundant data or parity information associated with RAID volumes requires additional disk space. In situations where the data is temporary, easily reproduced, or non-essential, the increased cost of data redundancy may not be justified.
- **Mean Time Between Failure (MTBF).** Using additional disks to maintain data redundancy also increases the chance of disk failure at any given moment. Although this cannot be avoided in situations where redundant data is a requirement, it does have implications for the workload of your organization's system support staff.

For more information, see "Choosing RAID Levels and Concatenation."

Choosing RAID Levels and Concatenation

You can use RAID or concatenation to control data storage on multiple disks. Each RAID level or concatenation has different performance and data protection characteristics.

The following sections provide specific information on how each RAID level or concatenation store data as well as their performance and protection characteristics.

- Concatenation (Spanned Volume)
- RAID Level 0 (Striping)
- RAID Level 1 (Mirroring)
- RAID Level 5 (Striping with distributed parity)
- RAID Level 50 (Concatenated distributed parity)
- RAID Level 10 (Striping over mirror sets)
- RAID Level 1-Concatenated (Concatenated mirror)
- RAID Level 0+1 (Mirroring over stripe sets)
- Considerations for RAID 10 and 50 on PERC Subsystem 1 Controllers
- Comparing RAID Level Performance

Concatenation (Spanned Volume)

In Array Manager, concatenation refers to storing data on either one array disk or on disk space that spans multiple array disks. When spanning more than one disk, concatenation enables the operating system to view multiple array disks as a single disk.

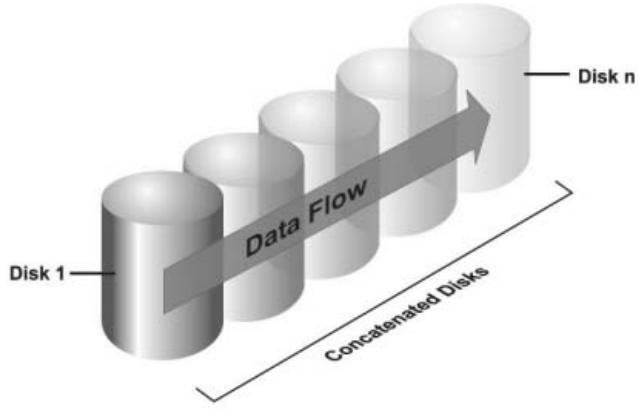
Data stored on a single disk can be considered a simple volume. This disk could also be defined as a virtual disk that comprises only a single array disk. Data that spans more than one array disk can be considered a spanned volume. Multiple concatenated disks can also be defined as a virtual disk that comprises more than one array disk.

In Array Manager, a dynamic volume that spans to separate areas of the same disk is also considered concatenated.

When an array disk in a concatenated or spanned volume fails, the entire volume becomes unavailable. Because the data is not redundant, it cannot be restored by rebuilding from a mirrored disk or parity information. Restoring from a backup is the only option.

Because concatenated volumes do not use disk space to maintain redundant data, they are more cost-efficient than volumes that use mirrors or parity information. A concatenated volume may be a good choice for data that is temporary, easily reproduced, or that does not justify the cost of data redundancy.

Concatenating Disks



- Concatenates *n* disks as one large virtual disk with a capacity of *n* disks.
- Data fills up the first disk before it is written to the second disk.
- No redundancy data is kept. When a disk fails, the large virtual disk fails.
- No performance gain.
- No redundancy.

Related Information:

- For the PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4/IM, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers, see the following:

- Supported RAID Levels (PERC Subsystem 1)
- Number of Array Disks per Virtual Disk
- Maximum Virtual Disks per Controller

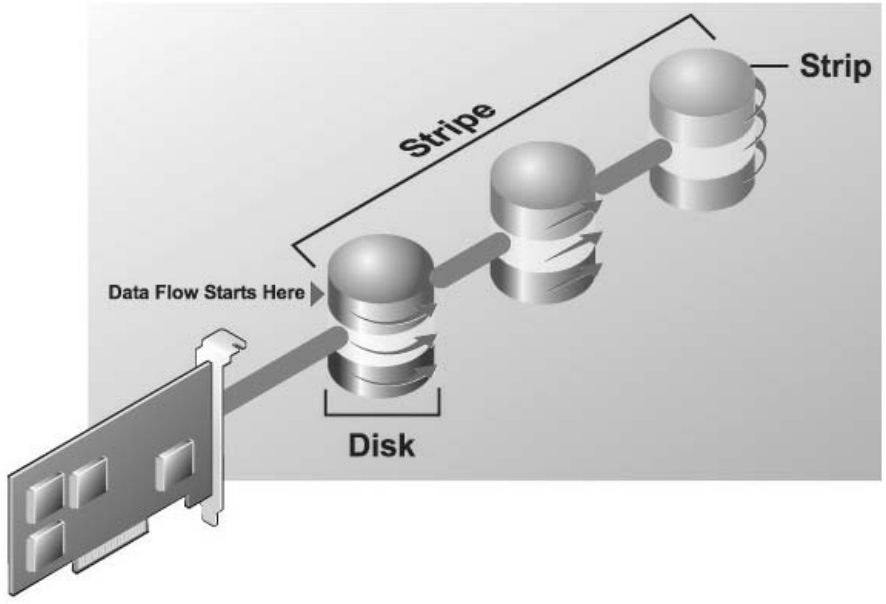
◆ For the PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers, see the following:

- Supported RAID Levels (PERC Subsystem 2)
- Number of Array Disks per Virtual Disk
- Maximum Virtual Disks per Controller

RAID Level 0 (Striping)

RAID 0 uses data striping, which is writing data in equal-sized segments across the array disks. RAID 0 does not provide data redundancy.

Striping Disks



RAID 0 Characteristics:

◆ Groups *n* disks as one large virtual disk with a capacity of *n* disks.

- Data is stored to the disks alternately.
- No redundancy data is kept. When a disk fails, the large virtual disk fails with no means of rebuilding the data.
- Better read and write performance.

Note The PERC 2/DC and PERC 2/SC controllers are limited to eight drives when using RAID 0, or RAID 5.

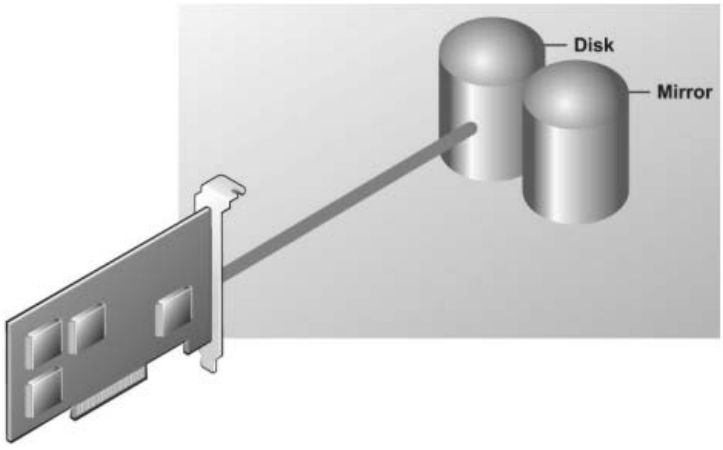
Related Information:

- For an overview of storage availability and performance considerations, see "Organizing Data Storage for Availability and Performance."
- For the PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4/IM, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers, see the following:
  - Supported RAID Levels (PERC Subsystem 1)
  - Number of Array Disks per Virtual Disk
  - Maximum Virtual Disks per Controller
- For the PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch and CERC SATA1.5/2s controllers, see the following:
  - Supported RAID Levels (PERC Subsystem 2)
  - Number of Array Disks per Virtual Disk
  - Maximum Virtual Disks per Controller

RAID Level 1 (Mirroring)

RAID 1 is the simplest form of maintaining redundant data. In RAID 1, data is mirrored or duplicated on one or more drives. If one drive fails, then the data can be rebuilt using the mirror.

Mirroring Disks



RAID 1 Characteristics:

- Groups *n* + *n* disks as one virtual disk with the capacity of *n* disks. The controllers currently supported by Array Manager allow the selection of two disks when creating a RAID 1. Because these disks are mirrored, the total storage capacity is equal to one disk.
- Data is replicated on the two disks.
- When a disk fails, the virtual disk still works. The data will be read from the failed disk's mirror.
- Better read performance, but slightly slower write performance.
- Redundancy for protection of data.
- RAID 1 is more "expensive" in terms of disk space since twice the number of disks are used than required to store the data without redundancy.

Related Information:

- For an overview of storage availability and performance considerations, see "Organizing Data Storage for Availability and Performance."
- For the PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4/IM, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers, see the following:
  - Supported RAID Levels (PERC Subsystem 1)
  - Number of Array Disks per Virtual Disk

- Maximum Virtual Disks per Controller

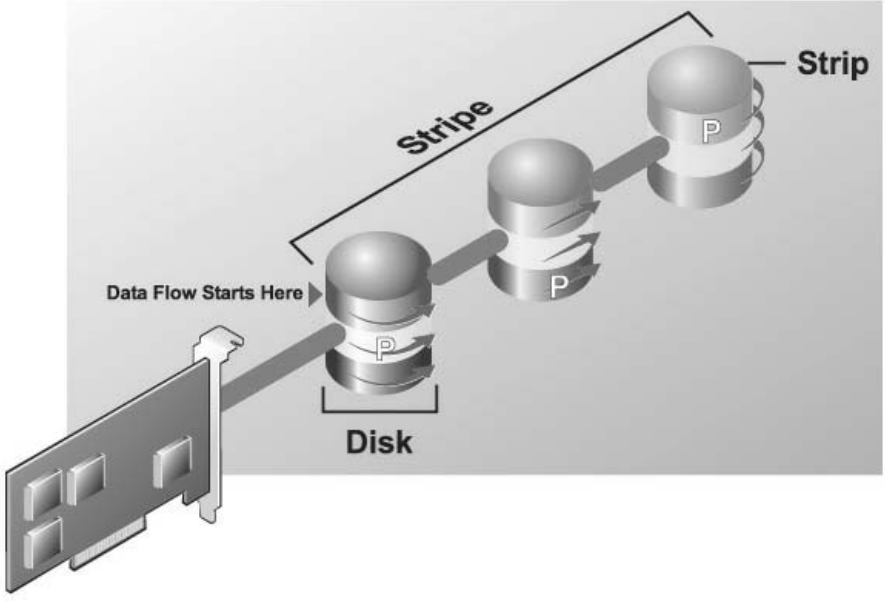
◆ For the PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers, see the following:

- Supported RAID Levels (PERC Subsystem 2)
- Number of Array Disks per Virtual Disk
- Maximum Virtual Disks per Controller

RAID Level 5 (Striping with distributed parity)

RAID 5 provides data redundancy by using data striping in combination with parity information. Rather than having all parity information reside on a single drive, the parity information is striped across all disks in the array.

Striping Disks with Distributed Parity Drive



RAID 5 Characteristics:

◆ Groups *n* disks as one large virtual disk with a capacity of (*n*-1) disks.

◆ Redundant information (parity) is alternately stored on all disks.

- When a disk fails, the virtual disk still works, but it is operating in a degraded state. The data is reconstructed from the surviving disks.
- Better read performance, but slower write performance.
- Redundancy for protection of data.

Note The PERC 2/DC and PERC 2/SC controllers are limited to eight drives when using RAID 0, or RAID 5.

Related Information:

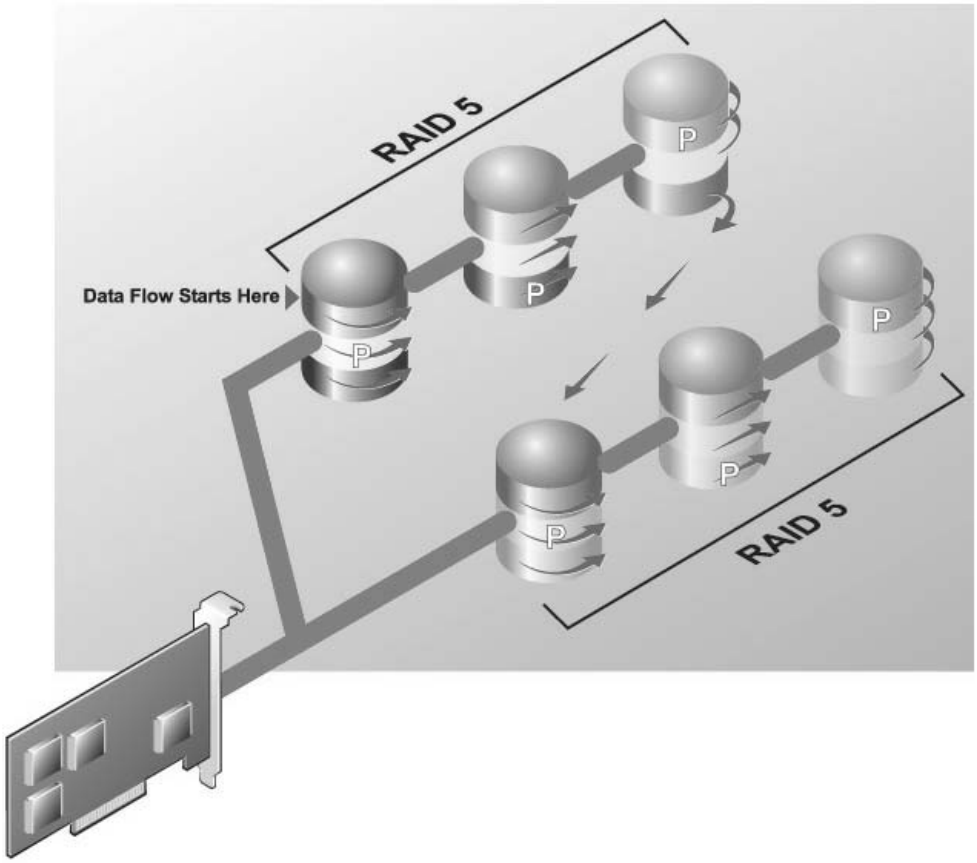
- For an overview of storage availability and performance considerations, see "Organizing Data Storage for Availability and Performance."
- For the PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4/IM, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers, see the following:
  - Supported RAID Levels (PERC Subsystem 1)
  - Number of Array Disks per Virtual Disk
  - Maximum Virtual Disks per Controller
- For the PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers, see the following:
  - Supported RAID Levels (PERC Subsystem 2)
  - Number of Array Disks per Virtual Disk
  - Maximum Virtual Disks per Controller

RAID Level 50 (Concatenated distributed parity)

RAID 50 is striping across more than one three-drive span. For example, a RAID 5 array that is implemented with three drives and then continues on with three more array drives would be a RAID 50 array.

It is possible to implement RAID 50 even when the hardware does not directly support it. In this case, you can implement more than one RAID 5 virtual disks and then convert the RAID 5 disks to dynamic disks. You can then create a dynamic volume that is spanned across all RAID 5 virtual disks.

RAID 50



RAID 50 Characteristics:

- Groups *n***s* disks as one large virtual disk with a capacity of *s**(*n*-1) disks, where *s* is the number of spans and *n* is the number of disks within each span.
- Redundant information (parity) is alternately stored on all disks of each RAID 5 span.
- Better read performance, but slower write performance.
- Requires proportionally as much parity information as standard RAID 5.
- Data is concatenated across all spans. RAID 50 is more expensive in terms of disk space.

Note On the PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers, there are special considerations when implementing RAID 50 on an array group that has disks of different sizes. See "Considerations for RAID 10 and 50 on PERC Subsystem 1 Controllers" for more information.

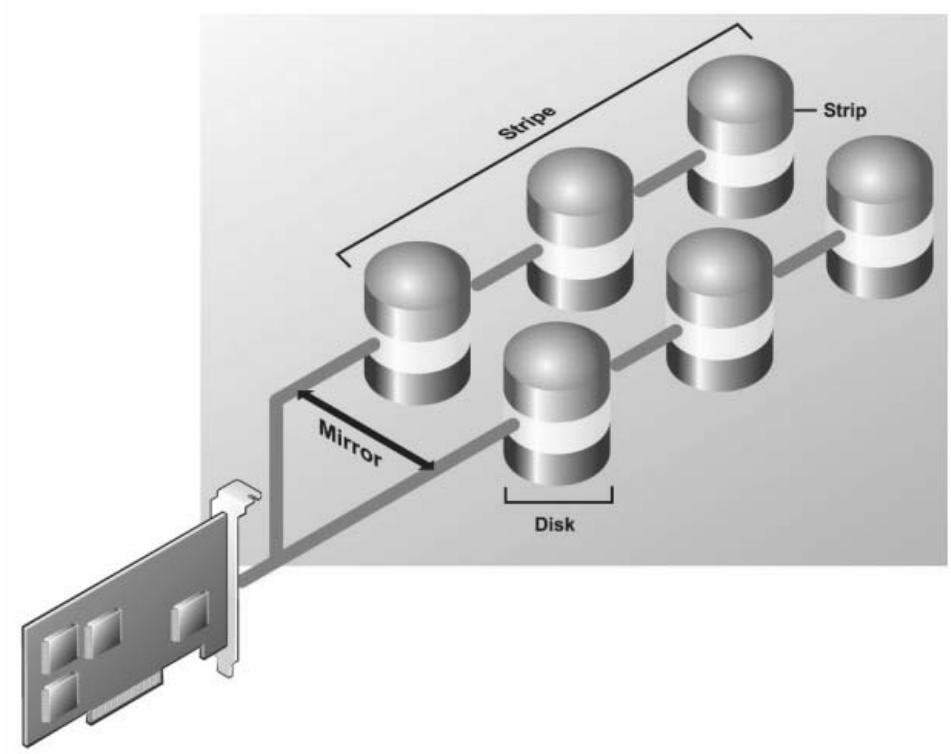
Related Information:

- For an overview of storage availability and performance considerations, see "Organizing Data Storage for Availability and Performance."
- For the PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4/IM, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers, see the following:
  - Supported RAID Levels (PERC Subsystem 1)
  - Number of Array Disks per Virtual Disk
  - Maximum Virtual Disks per Controller
- For the PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers, see the following:
  - Supported RAID Levels (PERC Subsystem 2)
  - Number of Array Disks per Virtual Disk
  - Maximum Virtual Disks per Controller

RAID Level 10 (Striping over mirror sets)

The RAID Advisory Board considers RAID Level 10 to be an implementation of RAID level 1. RAID 10 combines mirrored drives (RAID 1) with data striping (RAID 0). With RAID 10, data is striped across multiple drives. The set of striped drives is then mirrored onto another set of drives. RAID 10 can be considered a *stripe of mirrors*.

Striping over mirror sets



RAID 10 Characteristics:

- Groups *n* disks as one large virtual disk with a capacity of (*n* / 2) disks.
- Mirror images of the data are striped across sets of disk drives. This level provides redundancy through mirroring.
- When a disk fails, the virtual disk is still functional. The data will be read from the surviving mirrored disk.
- Improved read performance and write performance.
- Redundancy for protection of data.

Note On the PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers, there are special considerations when implementing RAID 10 on an array group that has disks of different sizes. See "Considerations for RAID 10 and 50 on PERC Subsystem 1 Controllers" for more information.

Related Information:

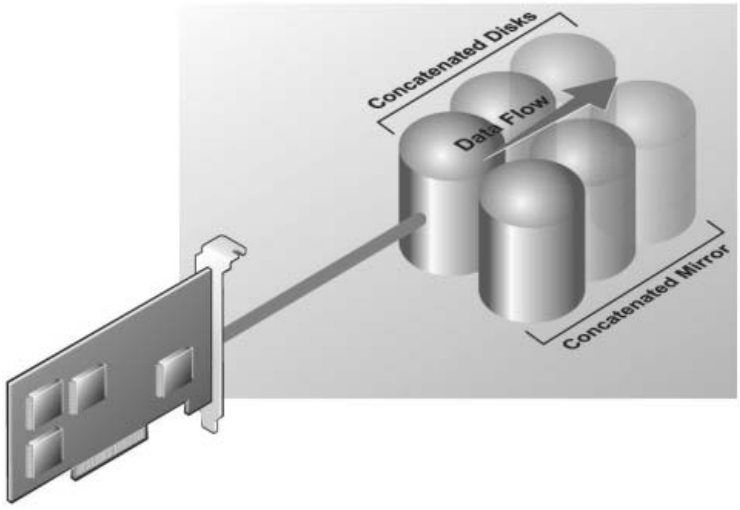
- For an overview of storage availability and performance considerations, see "Organizing Data Storage for Availability and Performance."
- For the PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4/IM, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers, see the following:
  - Supported RAID Levels (PERC Subsystem 1)
  - Number of Array Disks per Virtual Disk
  - Maximum Virtual Disks per Controller
- For the PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers, see the following:
  - Supported RAID Levels (PERC Subsystem 2)
  - Number of Array Disks per Virtual Disk
  - Maximum Virtual Disks per Controller

RAID Level 1-Concatenated (Concatenated mirror)

RAID-1 Concatenated is a RAID-1 array that spans across more than a single pair of array disks. This combines the advantages of concatenation with the redundancy of RAID-1. No striping is involved in this RAID type.

Also, RAID-1 Concatenated can be implemented on hardware that supports only RAID-1 by creating multiple RAID-1 virtual disks, upgrading the virtual disks to dynamic disks, and then using spanning to concatenate all of the RAID-1 virtual disks into one large dynamic volume.

RAID 1-Concatenated



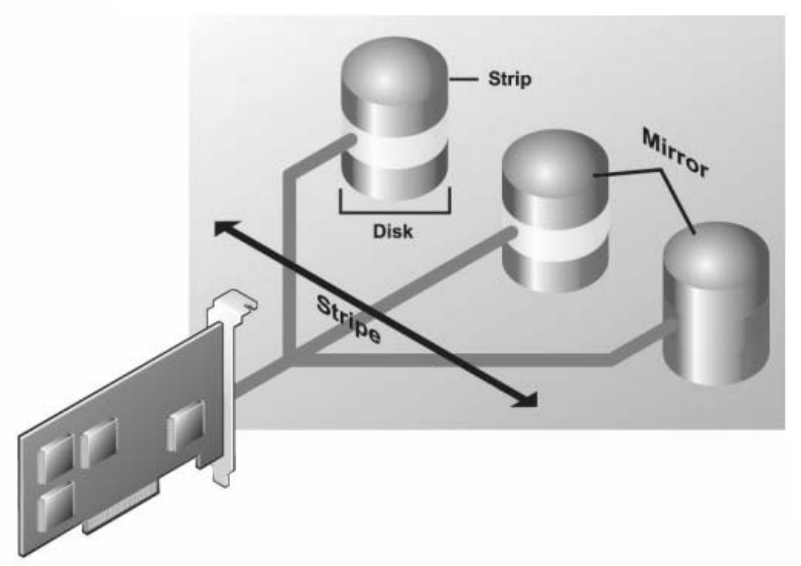
Related Information:

- For an overview of storage availability and performance considerations, see "Organizing Data Storage for Availability and Performance."
- For the PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4/IM, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers, see the following:
  - Supported RAID Levels (PERC Subsystem 1)
  - Number of Array Disks per Virtual Disk
  - Maximum Virtual Disks per Controller
- For the PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers, see the following:
  - Supported RAID Levels (PERC Subsystem 2)
  - Number of Array Disks per Virtual Disk
  - Maximum Virtual Disks per Controller

RAID Level 0+1 (Mirroring over stripe sets)

The RAID Advisory Board considers RAID Level 0+1 to be an implementation of RAID level 1. RAID 0+1 combines data striping (RAID 0) with mirrored drives (RAID 1). With RAID 0+1, data is striped across multiple drives and their mirrors. RAID 0+1 can be considered a *mirror of stripes*.

Mirroring Over Stripe Sets



Note In Array Manager, RAID 0+1 is only used with the PowerVault 660F storage system. This RAID level also allows you to use an odd number of disks.

RAID 0+1 Characteristics:

- Groups *n* disks as one large virtual disk with a capacity of (*n* / 2) disks.
- Mirror images of the data are striped across sets of disk drives. This level provides redundancy through mirroring.
- When a disk fails, the virtual disk is still functional. The data will be read from the surviving mirrored data stripes.
- Improved read performance and write performance.
- Redundancy for protection of data.

Related Information:

- For an overview of storage availability and performance considerations, see "Organizing Data Storage for Availability and Performance."
- For the PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4/IM, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers, see the following:
  - Supported RAID Levels (PERC Subsystem 1)
  - Number of Array Disks per Virtual Disk
  - Maximum Virtual Disks per Controller
- For the PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers, see the following:
  - Supported RAID Levels (PERC Subsystem 2)
  - Number of Array Disks per Virtual Disk
  - Maximum Virtual Disks per Controller

Considerations for RAID 10 and 50 on PERC Subsystem 1 Controllers

On the PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers, there are special considerations when implementing RAID 10 or RAID 50 on an array group that has disks of different sizes. When implementing RAID 10 or RAID 50, disk space is spanned to create the stripes and mirrors. The span size can vary to accommodate the different disk sizes. There is, however, the possibility that a portion of the largest disk in the array will be unusable, resulting in wasted disk space. For example, consider an array that has the following disks:

Disk A = 40 GB

Disk B = 40 GB

Disk C = 60 GB

Disk D = 80 GB

In this example, data will be spanned across all four disks until Disk A and Disk B and 40 GB on each of Disk C and D are completely full. Data will then be spanned across Disks C and D until Disk C is full. This leaves 20 GB of disk space remaining on Disk D. Data cannot be written to this disk space, as there is no corresponding disk space available in the array to create redundant data.

Comparing RAID Level Performance

The following table compares the performance characteristics associated with the more common RAID levels. This table provides general guidelines for choosing a RAID level. Keep in mind the needs of your particular environment when choosing a RAID level.

Note The following table does not show all RAID levels supported by Array Manager. For information on all RAID levels supported by Array Manager, see "Choosing RAID Levels and Concatenation."

RAID Level Performance Comparison

RAID Level	Data Availability	Read Performance	Write Performance	Rebuild Performance	Minimum Disks Required	Suggested Uses
RAID 0	None	Very Good	Very Good	N/A	N	Noncritical data
RAID 1	Excellent	Very Good	Good	Good	2N (N = 1)	Small databases, database logs, critical information
RAID 5	Good	Sequential reads: good. Transactional reads: Very good	Fair, unless using write-back cache	Poor	N + 1 (N = at least two disks)	Databases and other read-intensive transactional uses
RAID 10	Excellent	Very Good	Fair	Good	2N x X	Data-intensive environments (large records)
RAID 50	Excellent	Very Good	Fair	Fair	N + 2 (N = at least 4)	Medium-sized transactional or data-intensive uses

N = Number of array disks
X = Number of RAID sets

Wizards and Common Tasks 6

The Quick Access Window provides access to wizards that assist you in completing common storage tasks. The "Common Tasks" section provides links to commonly performed tasks and other useful information.

- Quick Access Wizards
- Common Tasks

Quick Access Wizards

You can complete the commonly performed storage tasks using the Quick Access window. Clicking an icon in this window launches a wizard to assist you in completing the task. In addition to clicking with the mouse, you can also use the Tab and Enter keys and the <Shift><X> key combination when selecting options on the Quick Access Window.

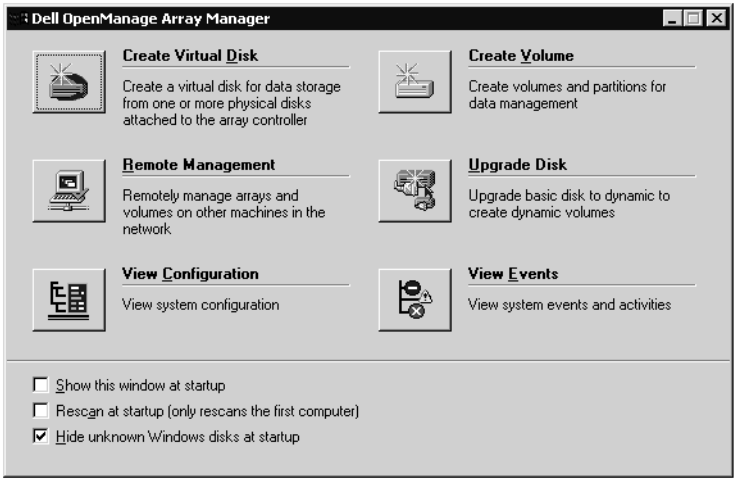
The Quick Access window is displayed when you start Array Manager. If you do not want the Quick Access window to be displayed when the console is launched, uncheck (deselect) the **Show this window at startup** check box located near the bottom of the Quick Access Window.

You can also open this window by clicking the **View** pull-down menu and selecting **Quick Access Window**.

You can use the Quick Access Window to specify that Array Manager perform a rescan when launched. See "Rescan at Startup" for more information.

Note When using the PowerVault 660F storage system, some options on the Quick Access Window may not function properly. In this case, highlight the appropriate object in the tree view hierarchy and select the option from the Task pull-down menu.

Quick Access Window



For more information on the Quick Access wizards, see the following:

- Create a Virtual Disk
- Create a Partition and Create a Dynamic Volume
- Connect to a Remote Computer
- Upgrade a Basic Disk to a Dynamic Disk
- View Configuration
- Viewing Event Listings

Create a Virtual Disk

The Create Virtual Disk Icon



The Create Virtual Disk wizard has an Express Mode and an Advanced Mode. The Express Mode enables you to quickly create a virtual disk using the available RAID levels and array disks. Depending on the available disk space and on controller considerations, the Express Mode calculates an appropriate virtual disk layout for you.

The Advanced Mode provides additional selections for creating the virtual disk. These selections include read, write, and cache policies. Advanced Mode also enables you to select the array disks and the controller channel that you want to use. When using Express Mode, these selections are made for you. Advanced Mode allows for a more detailed configuration and should be used by experienced administrators with a good knowledge of RAID levels and hardware.

Because of differences in controller features, the Create Virtual Disk wizard options may vary depending on the controller. See the following to create a virtual disk:

- "Creating Virtual Disks" for PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, and PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, CERC SATA1.5/2s controllers.
- "Create Virtual Disk (PV660F)" for the Dell PowerVault 660F storage system.

Create a New Volume

The Create Volume Icon



The Create Volume wizard contains information about creating both partitions and dynamic volumes. To create a partition, use the instructions in the "Create a Partition" topic in this chapter. To create a dynamic volume, use the instructions in the "Create a Dynamic Volume" topic later in this chapter.

Note Creating a volume is not supported on NetWare or Windows Server 2003. For more information on using the native Disk Management utility for Windows Server 2003, see "Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux." This section also describes considerations if you are planning on installing Linux on your system.

Create a Partition

The Create Volume Icon



You can create partitions only on basic disks. To check the disk type, click the disk and look at the **Type** column in the right pane.

1. If the Quick Access window is not already displayed, display this window by selecting the **Quick Access Window** option from the **View** pull-down menu.

2. Click the **Create Volume** button on the Quick Access window. The Create Volume wizard appears. You can read information about partitions and volumes and then click **Next** to continue.

3. You are asked to select whether to create a partition or a volume. Click **Next**.

   If the disk you have selected is a basic disk, the partition radio button is highlighted. You can create a partition only on a basic disk. To upgrade to a dynamic disk, see "Upgrade a Basic Disk to a Dynamic Disk" later in this chapter.

4. Select a disk for the partition.

5. The Create Volume wizard asks you to select a type and size of the partition. Select **Primary** or **Extended**. Click **Next** to continue.

6. Choose a free space for the partition. Click **Next** to continue.

7. Assign a drive letter to this volume, if you wish. Click **Next** to continue.

8. Format the volume using either NTFS, FAT, or FAT32. File and folder compression is supported only on NTFS volumes. Leave the **Allocation Unit Size** as **Default**.

Note It is strongly recommended that you leave the allocation unit size as **Default**. Selecting any other value requires an advanced knowledge of Windows file systems. You should thoroughly review the Microsoft Windows documentation before selecting any value other than **Default**.

---

9. Choose a file system and formatting option:

- Quick format — Formats the disk without checking for bad sectors.
- Enable file and folder compression — Can be used only if you choose NTFS format.

10. Click **Next** to continue, then **Finish**.

    If there is no free space on the selected disk, the free space selection is shaded. A message appears on the wizard that indicates a lack of available disk space.

Create a Dynamic Volume

The Create Volume Icon



You can create dynamic volumes only on dynamic disks. To check the disk type, click the disk and look at the **Type** column in the right pane.

There are particular considerations regarding dynamic disks and volumes on NetWare, Windows Server 2003, and Linux. See "Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux" for more information.

1. If the Quick Access window is not already displayed, display this window by selecting the **Quick Access Window** option from the **View** pull-down menu.

2. Click the **Create Volume** button on the Quick Access window. The Create Volume wizard appears. You can read information about partitions and volumes and then click **Next** to continue.

3. You are asked to select whether to create a partition or a volume. Make sure that the **Dynamic volume** button is highlighted.

   If you have selected a dynamic disk, the **Dynamic volume** button will already be highlighted. Click **Next** to continue.

4. You are asked to select the volume size, and the type of volume to create. Name the volume and provide a size in megabytes or gigabytes.

   If you are creating the volume on a Windows 2000 system, you will see a **Query Max Size** button. The **Query Max Size** button works differently depending on whether you had a disk selected when you chose the **Create Volume** command.

- If no disk was selected, clicking the **Query Max Size** button will display the maximum for all available disks in the Total volume size entry box.
- If a disk was selected, clicking the **Query Max Size** button will display the maximum size for the selected disk. However, if you click the button a second time, the size will be the maximum for all available disks.

Select one of the following volume types:

Check Box	Description
Concatenated	Simple or spanned volume (selected by default).
Striped	You must have unallocated space on at least two dynamic disks to create a striped volume. The number of columns is automatically selected as 2.
RAID-5	You must have at least three dynamic disks with sufficient unallocated space to create a RAID-5 volume. Note that you cannot create a RAID-5 volume if you are using a Windows 2000 Professional machine. For information on dynamic disks and volumes on a Windows Server 2003, Linux, or NetWare system, see "Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux."
Mirrored	You must have two dynamic disks with sufficient unallocated space to create a mirror. You can mirror only a simple or spanned volume. Any other selection will cause the mirror option on the menu to be inactive (dimmed). You cannot mirror a volume if you are using a Windows 2000 Professional machine. For information on dynamic disks and volumes on a Windows Server 2003, Linux, or NetWare system, see "Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux."

5. Click **Next** to continue, **Modify** to change, or **Back** to alter your selections.
6. After all selections are made, verify your settings and click **Next** to continue.
7. Do one of the following:
   - Choose a drive letter. Click **Next** to continue. Your next step will be to format the volume.

- If you do not want to assign a drive letter, check the appropriate box. This finishes the volume creation process. Click **Next** to continue, then **Finish.** The volume will appear without a drive letter and will need to be formatted.
- If you are using a Windows 2000 machine, you can also mount the volume at an empty NTFS folder and click **Next** to continue.

8. Format the volume, using either NTFS or FAT. File and folder compression is supported only on NTFS volumes. FAT32 is also available under Windows 2000. You can also choose a label, which will appear in the left pane of the Array Manager console.

9. You can choose a formatting option, if desired. Click the appropriate check box if you want to select the following. By default, neither box is checked.
   - Quick format — Formats the disk without checking for bad sectors.
   - Enable file and folder compression — Can be used only if you choose NTFS format.

10. Click **Next** to continue, then **Finish**.

Upgrade a Basic Disk to a Dynamic Disk

The Upgrade Dynamic Disk Icon



Use the following procedure to upgrade a basic disk to a dynamic disk.

There are particular considerations regarding dynamic disks and volumes on NetWare, Windows Server 2003, and Linux. See “Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux” for more information.

1. If the Quick Access window is not already displayed, display this window by selecting the **Quick Access Window** option from the **View** pull-down menu.

2. Click the **Upgrade Disk** button on the Quick Access window.

3. The Upgrade to Dynamic Disk wizard explains information about upgrading. Click **Next** to continue.

4. You are asked to select disks to upgrade. Add the basic disks you want to upgrade to the list of dynamic disks and click **Next** to continue, then **Finish**.

Note For more information on upgrading a basic disk to dynamic, see "Upgrading a Basic Disk to Dynamic" in the chapter on "Disks."

Connect to a Remote Computer

The Remote Management Icon



From the Array Manager console, you can use the Computer pull-down menu or My Network Places to connect to a remote system. For more information on remote connection, see "Connecting to a Remote System from the Console."

For information on the logon requirements for remote connection, see "Understanding the Conditions for Remote Connection."

Common Tasks

Using Array Manager, you can accomplish the following basic storage tasks:

- Create and configure virtual disks (RAID arrays). For more information see:
  - Creating Virtual Disks
  - Organizing Data Storage for Availability and Performance
- Manage RAID and non-RAID devices. For more information see:
  - Controllers
  - The Dell PowerVault 660F Storage System
  - Channels and Backplanes
  - Enclosures
  - Disks
  - Tape
  - CDROM

- Create partitions or dynamic volumes (Windows 2000). For more information see:
  - Create a Partition
  - Create a Dynamic Volume
- Format disks and perform other disk-related tasks (Windows 2000).
  - "Array and Hard Disk Menu Commands" or "Virtual Disks" for array disks and virtual disks on PERC Subsystem 1, PERC Subsystem 2, and SCSI Subsystem controllers.
  - "Array Disk Commands (PV660F)" or "Virtual Disk Commands (PV660F)" for the Dell PowerVault 660F Storage System.

Viewing Storage Status and Configuration 7

The Array Manager console displays the status and configuration of storage components. By selecting different tabs in the console, you can select different graphical displays of the system's storage. The console also displays error icons that indicate when a storage component has failed or is operating in a degraded state.

The Array Manager event log displays messages generated by the controllers. These event messages can be used to avoid or troubleshoot problems.

You can also right-click a storage object in the tree view to display a Properties dialog that contains status and other information for that object.

See the following for more information:

- Tabs
- Error Icons
- Viewing Event Listings
- Properties
- View Configuration

Error Icons

When a storage object has failed or is operating in a degraded state, an error icon is displayed over that object in the tree view or on the General tab to alert you to the error condition.

When a lower-level storage object has an error condition, that object and all objects above it in the tree view display the error icon. For example, if an array disk fails, the array disk icon in the tree view displays the Failed error icon. All objects above the array disk, including the top-level computer object, also display the Failed icon. This enables you to identify the failure of a storage object even when the tree view is collapsed and the failed object is hidden from view.

The error icons appear as follows.

No Disk Signature

If a disk requires a write signature, it displays the following icon.



Offline or Degraded

If a storage component is offline or operating in a degraded state, it displays the following icon.



Failed

If a storage object has failed, it displays the following icon.



Properties

You can right-click most storage objects in the tree view to display the properties for that object. For example, right-clicking an array disk object and selecting **Properties** from the pop-up menu displays the **Array Disk Properties** dialog. This dialog displays information about the array disk, such as its capacity and online or offline status.

Each storage object has its own properties dialog for displaying information appropriate for that object.

Note Certain storage objects, such as volumes and partitions, display the usual Windows properties dialog.

View Configuration

The Quick Access Windows has an option to display the storage configuration on your system's controllers, enclosures, and array groups.

To view the storage configuration:

1. If the Quick Access Window is not already displayed, select **Quick Access Window** from the **View** pull-down menu.

2. Click the **View Configuration** button on the **Quick Access Window**. This action displays the **View Hardware Configuration** dialog which displays your system's storage configuration.

Subsystems 8

This section applies to all subsystems supported by Array Manager.

A subsystem is a controller and all of the storage devices, such as disks and enclosures, that the controller manages. A subsystem may include a single controller, or several related controllers and their attached storage.

To view your storage subsystems in the Array Manager tree view, do the following:

1. Expand the **Arrays** object. This action enables you to see your storage subsystems. Array Manager supports the following subsystems:
   - PERC Subsystem 1
   - PERC Subsystem 2
   - SCSI Subsystem
   - PV660F Subsystem
2. Select a subsystem object and expand it until all of its lower-level objects are fully expanded. You can now see all of the controllers, disks, enclosures, and so forth included in the subsystem.

Note Because the controllers in each subsystem support different features, the objects in the expanded tree view may vary. For example, if a controller does not support a tape or CDROM device, then the Tape and CDROM objects are not displayed under that controller. For more information on the tree view, see "Tree View."

Subsystem Menu Commands

This section applies to the PERC Subsystem 1, PERC Subsystem 2, and the SCSI Subsystem.

The following sections describe pop-up menu commands for the PERC Subsystem 1, PERC Subsystem 2, and SCSI Subsystems. For information on the PV660F Subsystem, see "The Dell PowerVault 660F Storage System."

Right-click a subsystem in the Array Manager tree view to display its pop-up menu commands. Because the features for each subsystem can vary, your subsystem's pop-up menu may not display all of these commands.

- View Report
- Save
- Enable Smart Thermal Shutdown
- Disable Smart Thermal Shutdown
- Subsystem Properties

For information on "The Dell PowerVault 660F Storage System" pop-up menu commands, see "Subsystem Commands (PV660F)."

View Report

This section applies to the PERC Subsystem 1, PERC Subsystem 2, and the SCSI Subsystem.

This command enables you to view a report on the subsystem. The report indicates your system's controllers that belong to the subsystem and the virtual disks managed by those controllers. The report also provides additional information such as the virtual disk status, size and RAID type, as well as the array disk size, status, serial number and vendor.

To view a report on the subsystem:

1. Right-click the subsystem object in the Array Manager tree view.
2. Select **View Report** from the pop-up menu. This action displays the **Report** dialog which displays the status of the subsystem. To save the report to a file, click on the **Save** button.

For information on "The Dell PowerVault 660F Storage System" pop-up menu commands, see "Subsystem Commands (PV660F)."

Save

This section applies to the PERC Subsystem 1, PERC Subsystem 2, and the SCSI Subsystem.

This command enables you to save the information provided by the View Report command to a file.

To save information on the subsystem's controllers, virtual disks, and array disks to a file:

1. Right-click the subsystem object in the Array Manager tree view.
2. Select **Save** from the pop-up menu.

For information on "The Dell PowerVault 660F Storage System" pop-up menu commands, see "Subsystem Commands (PV660F)."

Enable Smart Thermal Shutdown

This section applies to a PERC Subsystem 1 or PERC subsystem 2 controller attached to an enclosure.

By default, the operating system and server shut down when the PV220S and PV221S enclosures reach a critical temperature of 0 or 50 degrees celsius. If you have implemented Channel Redundancy on the PV220S and PV221S enclosures, however, you can specify that only the enclosure and not the operating system and server be shut down when the enclosure reaches a critical temperature of 0 or 50 degrees celsius.

Note You must perform this procedure on each of the server's subsystems in the Array Manager tree view in order to implement smart thermal shutdown.

To specify that only the enclosure shut down at critical temperatures:

1. Right-click the **PERC Subsystem 1** object in the Array Manager tree view.
2. Select **Enable Smart Thermal Shutdown** from the pop-up menu.

For more information, see:

- Thermal Shutdown for the PV220S and PV221S Enclosures
- Disable Smart Thermal Shutdown

For information on "The Dell PowerVault 660F Storage System" pop-up menu commands, see "Subsystem Commands (PV660F)."

Disable Smart Thermal Shutdown

This section applies to a PERC Subsystem 1 or PERC subsystem 2 controller attached to an enclosure.

If you have previously selected the **Enable Smart Thermal Shutdown** option, you can deselect this option and return the system to its default setting by selecting the **Disable Smart Thermal Shutdown** option. The **Disable Smart Thermal Shutdown** setting indicates that the operating system and the server will shut down when the PV220S and PV221S enclosures reach a critical temperature of 0 or 50 degrees celsius.

To specify that the operating system and server shut down at critical temperatures:

1. Right-click the **PERC Subsystem 1** object in the Array Manager tree view.
2. Select **Disable Smart Thermal Shutdown** from the pop-up menu.

For more information, see:

- Thermal Shutdown for the PV220S and PV221S Enclosures
- Enable Smart Thermal Shutdown

For information on "The Dell PowerVault 660F Storage System" pop-up menu commands, see "Subsystem Commands (PV660F)."

Subsystem Properties

This section applies to the PERC Subsystem 1, PERC Subsystem 2, and the SCSI Subsystem.

This command enables you to display the subsystem properties.

To display the subsystem properties:

1. Right-click the subsystem object in the Array Manager tree view.
2. Select **Properties** from the pop-up menu. This action displays the **Subsystem Properties** dialog which displays the properties of the subsystem.

Subsystem properties may include:

- **Name** — This field displays the name of the subsystem object in the Array Manager tree view.
- **Vendor** — This field displays the subsystem vendor's name.

- **Status** — This field displays the current status of the subsystem. Possible values are:
  - **Ready** — The subsystem is functioning normally.
  - **Degraded** — The subsystem has suffered a failure of a component and is operating in a degraded state.
  - **Failed** — The subsystem has suffered a failure of one or more components and is no longer functioning.
- **Type** — This field indicates the type of storage component. For subsystem properties, the value in this field is "Subsystem."

For information on "The Dell PowerVault 660F Storage System" pop-up menu commands, see "Array Disk Commands (PV660F)."

Controllers 9

This section applies to all controllers supported by Array Manager.

Most operating systems (OS) do not read and write data directly from the disks, but instead send read and write instructions to a controller. The controller (or adapter) is the hardware in your system that interacts directly with the disks to write and retrieve data. A system can have more than one controller, with each controller attached to its own set of disks.

RAID controllers can span the boundaries of the disks so as to create an extended amount of storage space -- or a virtual disk -- using the capacity of more than one disk. The operating system views these virtual disks as if they are a single, large disk or drive.

The PERC Subsystem 1 and PERC Subsystem 2 controllers are RAID controllers. They have features for creating virtual disks. The SCSI Subsystem controllers are non-RAID SCSI controllers. They are not capable of creating virtual disks, but they do provide tape and CDROM management.

For information on the PV660F Subsystem, see "The Dell PowerVault 660F Storage System."

To view the controllers attached to your system:

1. Expand the **Arrays** object in the Array Manager tree view hierarchy.

2. Expand the subsystem object. This action displays the controllers included in the subsystem. Depending on which types of controllers your system has, there may be more than one subsystem object in the tree view.

Note Different controllers support different features. For this reason, the controller pop-up menu options and the storage objects displayed under a controller in the tree view can vary.

Understanding Read, Write, and Cache Policy

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

When creating a virtual disk, you specify the read, write, and cache policies for the virtual disk. The following describes how each of these policies functions.

Read Policy

The read policies indicate whether or not the controller should read sequential sectors of the logical drive when seeking data. The read policies are as follows.

Read policy for PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers

- **Read-Ahead.** When using read-ahead policy, the controller reads sequential sectors of the logical drive when seeking data. Read-ahead policy may improve system performance if the data is actually written to sequential sectors of the logical drive.
- **No-Read-Ahead.** Selecting no-read-ahead policy indicates that the controller should not use read-ahead policy.
- **Adaptive Read-Ahead.** When using adaptive read-ahead policy, the controller initiates read-ahead only if the two most recent read requests accessed sequential sectors of the disk. If subsequent read requests access random sectors of the disk, the controller reverts to no-read-ahead policy. The controller continues to evaluate whether read requests are accessing sequential sectors of the disk, and can initiate read-ahead if necessary.

Read policy for PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

- **Read Cache Enabled.** When the read cache is enabled, the controller reads the cache information to see if the requested data is available in the cache before retrieving the data from the disk. Reading the cache information first can provide faster read performance because the data (if available in the cache) can more quickly be retrieved from the cache than from the disk.
- **Read Cache Disabled.** When the read cache is disabled, the controller retrieves data directly from the disk and not from the cache.

Write Policy

The write policies specify whether the controller sends a write-request completion signal as soon as the data is in the cache or after it has been written to disk. The write policies are as follows.

Write policy for PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers

Note The **Write-Back Caching** option is not available for controllers that do not have a battery.

- **Write-Back Caching.** When using write-back caching, the controller sends a write-request completion signal as soon as the data is in the controller cache but has not yet been written to disk. Write-back caching may provide improved performance since subsequent read requests can more quickly retrieve data from the controller cache than they could from the disk. Write-back caching also entails a data security risk, however, since a system failure could prevent the data from being written to disk even though the controller has sent a write-request completion signal. In this case, data may be lost. Other applications may also experience problems when taking actions that assume the data is available on the disk.
- **Write-Through Caching.** When using write-through caching, the controller sends a write-request completion signal only after the data is written to the disk. Write-through caching provides better data security than write-back caching, since the system assumes the data is available only after it has been safely written to the disk.

Note Write-through caching is the default write policy setting when cluster mode is enabled. In cluster mode, the PERC 3/DC, 4/DC, and 4e/DC controllers only allow write-through caching.

Write policy for PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

- **Write Cache Enabled.** When the write cache is enabled, the controller writes data to the write cache before writing data to the array disk. Because it takes less time to write data to the write cache than it does to a disk, enabling the write cache can improve system performance. Once data is written to the write cache, the system is free to continue with other operations. The controller, in the meantime, completes the write operation by writing the data from the write cache to the array disk. The **Write Cache Enabled** option is only available if the controller has a functional battery. The presence of a functional battery ensures that data can be written from the write cache to the array disk even in the case of a power outage.

- **Write Cache Disabled.** This is the only available option if the controller does not have a functional battery.

Cache Policy

The Direct I/O and Cache I/O cache policies apply to reads on a specific logical drive. These settings do not affect the read-ahead policy. The cache policies are as follows.

Cache policies apply to the PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers.

- **Cache I/O.** Specifies that all reads are buffered in cache memory.
- **Direct I/O.** Specifies that reads are not buffered in cache memory. When using direct I/O, data is transferred to the controller cache and the host system simultaneously during a read request. If a subsequent read request requires data from the same data block, it can be read directly from the controller cache. The direct I/O setting does not override the cache policy settings. Direct I/O is also the default setting.

Understanding Hot Spares

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

A hot spare is an unused backup array disk that can be used to rebuild data from a redundant virtual disk. Hot spares remain in standby mode. When an array disk that is used in a redundant virtual disk fails, the assigned hot spare is activated to replace the failed array disk without interrupting the system or requiring your intervention. If a virtual disk using the failed array disk is not redundant, then the data is permanently lost without any method (unless you have a backup) to restore the data.

Hot spare implementation is somewhat different for controllers in PERC Subsystem 1 and PERC Subsystem 2.

- **PERC Subsystem 1:** For the PERC Subsystem 1 controllers, assigning a hot spare is equivalent to assigning an array disk to replace another array disk. If more than one redundant virtual disk resides on the array disk, then all redundant portions of the array disk are rebuilt. Note however, that you need to delete any non-redundant virtual disks (such as RAID 0) that reside on the array disk before rebuilding the array disk.

- **PERC Subsystem 2:** For PERC Subsystem 2 controllers, a hot spare is assigned to a virtual disk. When an array disk fails, only the portion of the array disk containing the virtual disk is rebuilt onto the hot spare. Data or space on the array disk not included in the virtual disk will not be rebuilt.

The following sections provide more information regarding hot spare considerations for PERC Subsystem 1 and PERC Subsystem 2 controllers:

- Considerations for Hot Spares (PERC Subsystem 1)
- Considerations for Hot Spares (PERC Subsystem 2)

The following sections describe procedures for assigning a hot spare:

- Assign Global Hot Spare. This procedure assigns a hot spare that is a backup for any physical disk in the array that fails.
- Configure Dedicated Hot Spare. This procedure assigns a hot spare that is a backup for a particular virtual disk.

Considerations for Hot Spares (PERC Subsystem 1)

This section applies to PERC 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers

When creating a virtual disk, the array disks included in the virtual disk can be different sizes. When assigning a hot spare on a PERC Subsystem 1 controller, the hot spare only needs to be the same size (or larger) as the smallest array disk included in the virtual disk.

This is because when using a PERC Subsystem 1 controller, you can assign array disks of different sizes to a virtual disk. Once you have fully consumed a smaller array disk with a virtual disk, however, any portion of larger array disks that are not consumed by the virtual disk become unusable. Therefore, there is no data on the unused portion of a larger disk that needs to be rebuilt. A redundant virtual disk will also be either striped or mirrored in equal portions across its member array disks. The amount of data requiring a rebuild will therefore not be larger than the smallest array disk.

A dedicated hot spare can only be assigned to a single virtual disk. A global hot spare is assigned to all redundant virtual disks on the controller. A global hot spare must be the same size (or larger) as the smallest array disk included in the largest virtual disk on the controller.

Once you have assigned a global hot spare, any new virtual disks created on the controller will not be protected by the hot spare if the partition size of the disk is larger than the global hot spare. In this case, you can unassign the global hot spare after creating a new virtual disk and then assign a new and larger hot spare to cover all redundant virtual disks on the controller.

Note The Configure Dedicated Hot Spare command is not available on the PERC 2/SC and PERC 2/DC controllers.

Dedicated Hot Spare Considerations for RAID 10 and RAID 50 on PERC Subsystem 1 Controllers

If you have created a RAID 10 or RAID 50 virtual disk that does not fully consume its member array disks, then you will not be able to assign a dedicated hot spare to the RAID 10 or RAID 50 virtual disk. Array Manager does not allow you to create RAID 10 and RAID 50 virtual disks from partial array disks. You will therefore not encounter this situation if you are using Array Manager to create your virtual disks. If you use the controller BIOS, however, you may be able to create a RAID 10 or RAID 50 virtual disk that uses partial array disks. In this situation, you will not be able to assign a dedicated hot spare to these RAID 10 or RAID 50 virtual disks.

Array Disk State, Event Messages and Hot Spares on PERC Subsystem 1 Controllers

If you have a hot spare assigned to a virtual disk and an array disk in the virtual disk fails, the failed array disk may change from Online state to Ready state without displaying a Failed state. This occurs when the hot spare is activated before the array disk is able to report the Failed state. Because the Failed state is not reported, the "array disk failed" event 500 is not generated.

Once the hot spare is activated it changes to the Rebuilding state. If you review the event log and identify a "rebuilding" event such as 526 or 527, you can assume that an array disk has failed.

Considerations for Hot Spares (PERC Subsystem 2)

This section applies to PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

A dedicated hot spare on PERC Subsystem 2 controllers rebuilds the portion of a redundant virtual disk that resides on the failed array disk. On PERC Subsystem 2 controllers, individual array disks may be included in more than one virtual disk. (Assigning a portion of an array disk to a virtual disk does not preclude the remaining portion of the array disk from being used by other virtual disks.)

You can assign the same dedicated hot spare to more than one virtual disk. In this case, the hot spare attempts to rebuild all portions of redundant virtual disks that reside on a failed array disk. To increase the likelihood that the hot spare is able to rebuild all virtual disks, you should do the following:

1. Create virtual disks that share the same set of array disks.

2. Only assign dedicated hot spares to those virtual disks that share the same set of array disks.

3. Assign a hot spare that is big enough to rebuild the total size of all virtual disks to which the hot spare is assigned. For example, if you have a 20 MB, a 30 MB, and a 50 MB virtual disk all sharing the same set of array disks, then the hot spare needs to be 100 MB or larger.

Once the hot spare is activated to rebuild a particular virtual disk, it is no longer available for rebuilding other virtual disks should an additional array disk fail. For this reason, when a hot spare is activated it is automatically unassigned from the remaining virtual disks. To maintain data protection, you will need to add a new hot spare and assign it to the other virtual disks.

Note The Configure Dedicated Hot Spare command is not available on the CERC SATA1.5/2s controller.

PERC Subsystem 1 Controllers

- Dell PowerEdge™ Expandable RAID Controller 2 Single Channel (PERC 2/SC)
- Dell PowerEdge Expandable RAID Controller 2 Dual Channel (PERC 2/DC)
- Dell PowerEdge Expandable RAID Controller 3 Single Channel (PERC 3/SC)
- Dell PowerEdge Expandable RAID Controller 3 Dual Channel Lite (PERC 3/DCL)
- Dell PowerEdge Expandable RAID Controller 3 Dual Channel (PERC 3/DC)
- Dell PowerEdge Expandable RAID Controller 3 Quad Channel (PERC 3/QC)
- Dell PowerEdge Expandable RAID Controller 4 Single Channel (PERC 4/SC)
- Dell PowerEdge Expandable RAID Controller 4 Dual Channel (PERC 4/DC)
- Dell PowerEdge Expandable RAID Controller Express Dual Channel (PERC 4e/DC)
- Dell PowerEdge Expandable RAID Controller 4 Dual Integrated Channel (PERC 4/Di)
- Dell PowerEdge Expandable RAID Controller 4 Integrated Mirroring (PERC 4/IM)
- Dell PowerEdge Expandable RAID Controller 4 Express Single Integrated Channel (PERC 4e/Si)
- Dell PowerEdge Expandable RAID Controller 4 Express Dual Integrated Channel (PERC 4e/Di)

- Dell PowerEdge Cost Effective RAID Controller Advanced Technology Attachment 100/4ch (CERC ATA100/4ch)

Description of PERC Subsystem 1 Controllers

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4/IM, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers

The PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC 4/Di, 4e/Si, 4e/Di, and controllers are high-performance intelligent PCI-to-SCSI host adapters with RAID control capabilities. The PERC 4e/DC (dual channel) controller provides PCI-express architecture.

The CERC ATA100/4ch controller is a PCI-to-IDE host adapter with RAID control capabilities. Each channel on the CERC ATA100/4ch supports 100 MB data transfer. Each channel on this controller supports a single device.

Each channel on the PERC 2/SC and PERC 2/DC controllers supports SCSI data transfer rates of up to 80 MB per second per channel and up to 160 MB per second per channel for the PERC 3/SC, 3/DCL, 3/DC, and 3/QC controllers. The channels on these controllers also support up to 15 devices.

Each channel on the PERC 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, controllers supports SCSI data transfer rates of up to 320 MB.

The PERC 4/IM controller supports SCSI data transfer rates of up to 320 MB per second. Using the PERC 4/IM controller BIOS, you can create an integrated mirror. See "Integrated Mirroring and the PERC 4/IM Controller" for more information.

For more information about specific controller features, refer to the hardware manual that came with the controller.

The following features are available on the PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers.

Note Supported features may vary from controller to controller.

- **Hot spares.** The PERC 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA/1004ch controllers support global and dedicated hot spares. The PERC 2/SC and 2/DC controllers support global hot spares. A hot spare is a backup for a disk that fails. Global hot spares are backups for any disk that fails in an array. A dedicated hot spare is a backup for an array disk that is included in a specific virtual disk. For more information, see the "Understanding Hot Spares" section.
- **Online capacity expansion.** Online capacity expansion (OLCE) enables you to expand the capacity of a virtual disk while it remains online by adding additional disks to the virtual disk. For more information, see the "Add Member" section.

- **FlexRAID Virtual Sizing.** This feature is available on the PERC 2/SC and PERC 2/DC controllers. FlexRAID virtual sizing enables you to add and remove array disks of different sizes from a virtual disk without rebooting the system. For more information, see the "FlexRAID Virtual Sizing" section.
- **RAID migration.** After creating a virtual disk, you can change the RAID level and stripe size of the virtual disk. For more information, see the "Transform Type" section.
- **Array and volume migration.** The PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers enable you to migrate array disks, virtual disks, and volumes from one controller to another. For more information, see the "Array and Volume Migration" section.
- **Cache policies.** The 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, and 4e/Di controllers provide different caching policies to meet the needs of data security and system performance. For more information, see the "Change Policy" section.
- **Check consistency.** The check consistency feature determines the integrity of mirrored data and parity information. When necessary, this feature rebuilds the parity information. For more information, see the "Check Consistency" section.
- **Cluster Support.** Array Manager supports PERC 3/DC, 4/DC, and 4e/DC controllers that have Microsoft Cluster Server (MSCS) enabled. A *cluster* refers to two or more servers that are connected so that their resources can be shared and accessed as if the clustered servers were a single machine. Clusters provide increased availability because when one server in the cluster experiences downtime, another server can take over the processing and workload requests of the failed server. For more information, see the "Cluster Support" section.

Note The PERC 4/IM controller provides integrated mirroring. The ability to create an integrated mirror with Array Manager is not supported. You can, however, create the integrated mirror using the controller BIOS. If you do so, the integrated mirror will be displayed as a virtual disk in the Array Manager console. See "Integrated Mirroring and the PERC 4/IM Controller" for more information.

- **Background Initialization.** The background initialization of a redundant virtual disk prepares the virtual disk to maintain redundant data and improves write performance. See "Background Initialization on PERC Subsystem 1 Controllers" for more information.

Supported RAID Levels (PERC Subsystem 1)

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4/IM, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers

The following table indicates the RAID levels supported by the PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4/IM, 4e/Si, 4e/Di, and CERC ATA100/4ch Controllers. Refer to the "comments" column for information on how a controller implements a particular RAID level.

For information on controller features and virtual disks, see "Number of Array Disks per Virtual Disk" and "Maximum Virtual Disks per Controller."

Supported RAID Levels for PERC Subsystem 1 Controllers

Concatenation or RAID Level	Controllers	Comments
RAID-0	PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch	
RAID-1	PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/IM, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch	
RAID-5	PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch	

Supported RAID Levels for PERC Subsystem 1 Controllers		
Concatenation or RAID Level	**Controllers**	**Comments**
RAID 10	PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, and CERC SATA1.5/6ch	
RAID 50	PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, and 4e/Di	RAID 50 support on the PERC 4/Di, 4e/Si, and 4e/Di controllers is implemented as striping across multiple RAID 5 volumes.

Cluster Support

This section applies to 3/DC, 4/DC, and 4e/DC controllers

Array Manager supports PERC 3/DC, 4/DC, and 4e/DC controllers that have Microsoft Cluster Server (MSCS) enabled.

Note Array Manager does not set resource ownership in a Microsoft Windows cluster configuration.

Note Dell does not provide cluster support on NetWare 6.x.

Updated Information

This section applies to 3/DC, 4/DC, and 4e/DC controllers

After a failover of cluster resources, it is necessary to perform a rescan operation in order for Array Manager to display the most up-to-date information about shared resources. You can specify that a rescan automatically be performed when the Array Manager console is launched by selecting the **Rescan at startup** check box on the Quick Access window.

If a failover occurs while the Array Manager console is displayed, you can perform a rescan by selecting **Rescan** from the **View** pull-down menu or using the Rescan button on the Array Manager toolbar.

Creating and Deleting Virtual Disks on Cluster-enabled Controllers

This section applies to 3/DC, 4/DC, and 4e/DC controllers

If you are using a PERC 3/DC, 4/DC, or 4e/DC controller in a cluster configuration, you must shut down and power off the other systems in the cluster before creating or deleting the virtual disk. The following procedure describes the sequence of actions required to create or delete a virtual disk from a cluster-enabled controller. For the purposes of this procedure, the system on which you are deleting the virtual disk is identified as "system A" and the other system in the cluster is identified as "system B."

1. Stop the clustering services on system B.
2. Shut down and power off system B.
3. Create or delete the virtual disk on system A using the procedure described in the "To Create a Virtual Disk"section or the "Deleting Virtual Disks"section. If you are creating a virtual disk on a cluster-enabled controller, you should review the "Virtual Disk Considerations for PERC Subsystem 1 Controllers" section before continuing.
4. Reboot system A.
5. Restart system B.

Automatic Rescan

This section applies to 3/DC, 4/DC and 4e/DC controllers

There is a **Rescan at startup** check box at the bottom of the Quick Access window that is designed for use with cluster-enabled controllers. When the **Rescan at startup** check box is selected, Array Manager automatically performs a rescan when the Array Manager console is launched. It is recommended that this check box be selected on a system with a cluster-enabled PERC 3/DC, 4/DC, 4e/DC controller.

You can display the Quick Access Window by selecting **Quick Access Window** from the **View** pull-down menu.

Downloading Firmware and Cluster Controllers

This section applies to 3/DC, 4/DC, and 4e/DC controllers

When downloading firmware on a clustered controller, it is recommended that you shut down the other systems in the cluster first. When restarted, the other systems in the cluster should be able to see the firmware updates that you have applied.

If, however, you have downloaded firmware on a controller without first shutting down the other systems in the cluster, you may find that the other systems cannot see the firmware update until you restart the disk management service on those systems. For example, if you download firmware onto system A, and system B cannot see the firmware update, then restart the disk management service on system B.

Note If you downloaded firmware using an application other than Array Manager, rescan the controller to see the updated firmware version in the Array Manager console.

Integrated Mirroring and the PERC 4/IM Controller

This section applies to PERC 4/IM controller

The PERC 4/IM controller enables you to mirror a disk drive that resides internally in the server. This feature can be used to mirror a system's boot drive, ensuring that the system remains running in the event of a boot drive failure. The PERC 4/IM controller firmware maintains the mirrored data on both drives so that the system's CPU is not burdened with the extra processing required to maintain the mirrored data.

When implementing mirroring on a PERC 4/IM controller, you use the controller BIOS to create a virtual disk from hard disks. A hard disk is a physical disk that is directly visible to the operating system. Unlike creating a virtual disk on other controllers, the PERC 4/IM controller is able to implement a mirror for a hard disk that already contains data. The data is then copied to the mirror. Any data previously residing on the mirror is overwritten.

Once you have created the integrated mirror using the controller BIOS, the operating system sees the mirror as a virtual disk and a virtual disk object for the mirror is displayed in the Array Manager tree view. Expanding the virtual disks object displays the disks included in the mirror. These disks are now labeled "array disks" since they are part of a virtual disk and no longer individually visible to the operating system.

Background Initialization on PERC Subsystem 1 Controllers

On PERC 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers, background initialization of a redundant virtual disk begins automatically within 0 to 5 minutes after the virtual disk is created. The background initialization of a redundant virtual disk prepares the virtual disk to maintain redundant data and improves write performance. For example, after the background initialization of a RAID 5 virtual disk completes, the parity information has been initialized. After the background initialization of a RAID 1 virtual disk completes, the array disks are mirrored.

The background initialization process helps the controller identify and correct problems that may occur with the redundant data at a later time. In this regard, the background initialization process is similar to a check consistency.

The background initialization should be allowed to run to completion. If cancelled, the background initialization will automatically restart within 0 to 5 minutes. Other processes such as read and write operations are possible while the background initialization is running. Other processes, such as creating a virtual disk, cannot be run concurrently with a background initialization. These processes cause the background initialization to cancel.

FlexRAID Virtual Sizing

This section applies to PERC 2/SC and 2/DC controllers

FlexRAID virtual sizing is a feature of PERC 2/SC and PERC 2/DC controllers. When enabled, the operating system sees your virtual disk as an arbitrarily large disk, typically 1024 GB, instead of its actual size. This lets you add and remove array disks of different sizes to your configuration without having to reboot the system.

Note This feature is not recommended for use on a Windows 2000 or Windows Server 2003 server.

This section on FlexRAID virtual sizing has the following topics:

- Enabling FlexRAID Virtual Sizing
- Retained FlexRAID Virtual Sizing Characteristics
- FlexRAID and Dynamic Volumes

Enabling FlexRAID Virtual Sizing

This section applies to PERC 2/SC and 2/DC controllers

This feature is specific to each virtual disk. You create the virtual disk first in Array Manager. Then you must exit Array Manager, reboot the system to bring up the PERC 2/SC or PERC 2/DC controller BIOS utility, and then enable virtual sizing on that virtual disk through a command in the utility.

To Enable FlexRAID Virtual Sizing

1. After booting the system, when you see "PowerEdge Expandable RAID Controller" displayed in the BIOS, press **Ctrl+M** to enter the BIOS configuration utility.
2. Select **Objects**.
3. Select **Logical Drives**.
4. Select any logical drive, such as **Logical Drive 0**. Logical drives and virtual disks are the same.
5. Select **View/Update Parameters**.
6. Go to **FlexRAID Virtual Sizing = ...** and select **Enabled**.
7. Exit the BIOS configuration utility and reboot.

Retained FlexRAID Virtual Sizing Characteristics

This section applies to PERC 2/SC and 2/DC controllers

In Array Manager, if you delete a virtual disk with FlexRAID virtual sizing enabled and then create another virtual disk in its place, that disk will retain the deleted disk's FlexRAID virtual sizing characteristics. For example, if you just deleted six virtual disks with virtual sizing enabled on them and create eight new ones, you should still see virtual sizing on the first six of them, but not on the last two. To avoid having retained virtual sizing characteristics from deleted virtual disks, reboot before creating new virtual disks.

FlexRAID and Dynamic Volumes

This section applies to PERC 2/SC and 2/DC controllers

FlexRAID virtual sizing does not work with dynamic volumes. On Windows 2000, you will not be able to upgrade a basic disk to dynamic if FlexRAID virtual sizing is enabled. You do not need FlexRAID virtual sizing on a Windows 2000 machine because you can add storage and create and modify dynamic volumes without having to reboot.

Controller Rebuild Rate

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers

During a rebuild, the complete contents of an array disk are reconstructed. The rebuild rate, configurable between 0% and 100%, represents the percentage of the system resources dedicated to rebuilding failed array disks. At 0%, the rebuild will have the lowest priority for the controller, will take the most time to complete, and will be the setting with the least impact to system performance. At 100%, the rebuild will be at the highest priority for the controller, will minimize the rebuild time, and will be the setting with the most impact to system performance.

The setting can be modified while a rebuild is already in progress.

Drive Status Changes

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers

Because of hardware restrictions, the PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers will not detect a drive status change until I/O is attempted. For example, when an unconfigured drive is removed, the controller will not detect the change until a manual rescan is done or other I/O operations are attempted.

S.M.A.R.T. Alerts

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers

The PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers will not report S.M.A.R.T. alerts for unassigned or hot spare drives.

PERC Subsystem 1 Controller Issues

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers

This section on controller issues covers the following topics:

- Array Manager and PERC Console Compatibility Issues
- Controller Rebuild Rate
- Drive Status Changes
- S.M.A.R.T. Alerts

Array Manager and PERC Console Compatibility Issues

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers

Both Array Manager and PERC Console programs can be installed on the same server. However, only one program can be used to configure PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers at any given time. To maintain the integrity of storage device configuration, simultaneous use of Array Manager and PERC Console is not supported. It is strongly recommended that Array Manager and PERC Console not be used at the same time on the same server.

PERC Subsystem 2 Controllers

The following controllers are supported by the PERC subsystem 2:

- Dell PowerEdge™ Expandable RAID Controller 2 (PERC 2)
- Dell PowerEdge Expandable RAID Controller 2 single integrated channel (PERC 2/Si)

- Dell PowerEdge Expandable RAID Controller 3 single integrated channel (PERC 3/Si)
- Dell PowerEdge Expandable RAID Controller 3 dual integrated channels (PERC 3/Di)
- Dell PowerEdge Cost Effective RAID Controller Serial Advanced Technology Attachment 6 channel (CERC SATA1.5/6ch)
- Dell PowerEdge Cost Effective RAID Controller Serial Advanced Technology Attachment (CERC SATA1.5/2s)

Description of PERC Subsystem 2 Controllers

This section applies to PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

The PERC 2 RAID controller contains four channels. PERC 2/Si and PERC 3/Si are single integrated channel controllers. The PERC 3/Di is a dual integrated channel controller. The CERC SATA1.5/6ch and CERC SATA1.5/2s controllers are PCI RAID controllers with Serial Advanced Technology Attachment.

Supported RAID Levels (PERC Subsystem 2)

This section applies to PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

The PERC 2, PERC 2/Si, PERC 3/Si, and PERC 3/Di controllers support the following RAID levels:

- Concatenated
- RAID-0
- RAID-1
- RAID-5
- RAID-10
- RAID-1 Concatenated

The CERC SATA1.5/2s controller supports the following RAID levels:

- RAID-0
- RAID-1

For RAID support information on the CERC SATA / 6ch controller, see RAID Support on the CERC SATA1.5 / 6ch.

For more information on storage concepts used in this section, see "What Is RAID?" and "Choosing RAID Levels and Concatenation."

For information on controller features and virtual disks, see "Number of Array Disks per Virtual Disk" and "Maximum Virtual Disks per Controller."

RAID Support on the CERC SATA1.5/6ch

The CERC SATA1.5 / 6ch controller supports the following RAID levels:

- Concatenated (supports creating a concatenated virtual disk using a single array disk.)
- RAID-0
- RAID-1
- RAID-5
- RAID-10

The CERC SATA1.5 / 6ch controller does not support a concatenated virtual disk using more than a single array disk or a RAID 1-concatenated virtual disk. If you have a virtual disk using an unsupported RAID level on this controller, you should do the following:

1. Back up the virtual disk's data to tape.
2. Delete the unsupported virtual disk.
3. Create a virtual disk using a supported RAID level.
4. Restore the data from backup.

Managing the Controller Battery

This section applies to PERC 2 and 3/Di controllers

The battery of the PERC 2 controller preserves the contents of its nonvolatile cache memory in the event of power loss. The battery provides a minimum of 72 hours protection for DIMMs up to 64 MB and 48 hours protection for 128-MB DIMMs. However, the duration of protection, or holdover time, depends on the battery's charge level. This section has two topics:

- Battery Recharging

- Battery Reconditioning

Note The PERC 2/Si and PERC 3/Si RAID controllers do not contain a battery and therefore do not support battery management.

Battery Recharging

This section applies to PERC 2 and 3/Di controllers

The battery is fully discharged when you first install the controller. The controller detects this condition and automatically begins charging the battery. This charging process takes approximately three hours. The controller can be used during this time; however, the battery is unable to meet the specified holdover time until it is fully charged. The battery is still able to handle brief power losses during the initial charge cycle.

Battery Reconditioning

This section applies to PERC 2 and 3/Di controllers

After you have installed the controller, it notifies you every six months that a battery reconditioning cycle is required. This reconditioning cycle requires a full discharge and recharge of the battery. It ensures that the battery's capacity is being measured correctly and that the battery's full holdover time is maintained.

To Recondition the Battery

1. Right-click a controller.
2. Click **Properties**.
3. Click the **Recondition** button.

The reconditioning cycle takes approximately 8 to 10 hours. During this time, the battery's holdover time is reduced to zero (during discharge) and then restored (during full charge). To ensure that no data is lost from nonvolatile cache memory, the controller automatically disables cache (both read and write) during a reconditioning cycle. When the cache is disabled, performance is degraded for the reconditioning period. Once the reconditioning cycle is complete, the cache is restored to its original settings.

It is highly recommended that battery reconditioning be done during off-peak times.

Note On some servers, the PERC 3/Di controller has a lithium ion battery which does not require reconditioning. In this case, the Recondition button is inactive (greyed out). There are other situations that cause the Recondition button to be inactive. For example, if the battery is currently in a recondition cycle or is not installed or otherwise missing, the Recondition button is inactive.

Array Manager and FAST Compatibility Issues

This section applies to PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

Array Manager and FAST cannot both be installed on the same machine. If the Array Manager installation process detects that FAST is already installed, the Array Manager installation will not complete.

SCSI Subsystem Controllers

The following controllers are supported by the SCSI subsystem:

- Adaptec Ultra SCSI Controllers (40 MB device speed). These controllers include the 3944UW, 2944UW, AIC-7860, and AIC-7880 controllers. The AIC-7860 and AIC-7880 controllers are an ASIC embedded on the motherboard.
- Adaptec Ultra2 SCSI Controllers (80 MB device speed). These controllers include the 2940U2W and AIC-7890 controllers. The AIC-7890 controllers are an ASIC embedded on the motherboard.
- Adaptec Ultra160 SCSI Controllers (160 MB device speed). These controllers include the 39160, 29160LP, AIC-7899, and AIC-7892 controllers. The AIC-7892 and AIC-7899 controllers are an ASIC embedded on the motherboard.
- LSI 1020 and LSI 1030 Controllers (320 MB device speed). These controllers provides Ultra320 SCSI capability.

Description of SCSI Subsystem Controllers

This section applies to Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers

The SCSI subsystem controllers are non-RAID controllers that support SCSI devices. Because these controllers are non-RAID, they do not support virtual disks. You can manage these non-RAID controllers and their attached SCSI devices with Array Manager.

The Ultra SCSI controllers supported by Array Manager are a high voltage differential (HVD) SCSI drop-in controller that fits in a PCI slot on the server or is embedded in the server. This controller can be used to attach HVD devices such as an HVD tape drive.

The Ultra2 SCSI and Ultra160 SCSI controllers are low voltage differential (LVD) drop-in controllers that fit in a PCI slot on the server or are embedded in the controller. These controllers can be used to attach external enclosures and other LVD SCSI devices.

The LSI 1020 and LSI 1030 controllers provide Ultra320 SCSI capability. The LSI 1020 is a PCI-X to Single Channel Ultra320 SCSI controller. The LSI 1030 is a PCI-X to Dual Channel Ultra320 SCSI controller.

Array Manager also supports RAID controllers on which you can create virtual disks. For information on these controllers, see the following:

- PERC Subsystem 1 Controllers (SCSI controllers)
- PERC Subsystem 2 Controllers (SCSI controllers)
- The Dell PowerVault 660F Storage System (fiber channel controller)

Controller Menu Commands

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, 4/IM, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, CERC SATA1.5/2s, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers

The following sections describe pop-up menu commands for the PERC Subsystem 1, PERC Subsystem 2, and SCSI Subsystem controllers. For information on the PV660F controller commands, see "The Dell PowerVault 660F Storage System."

Right-click a controller in the Array Manager tree view to display its pop-up menu commands. Because the features for each controller can vary, your controller's pop-up menu may not display all of these commands.

- Rescan Controller
- Enable Alarm
- Disable Alarm
- Quiet Alarm
- Rebuild Rate
- Initialize Array Disks
- Export Log File

- Controller Properties

Rescan Controller

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, CERC SATA1.5/2s, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers

The **Rescan** command scans the controller channels to verify the currently connected devices or to recognize new devices that have been added to the channels. When you do a rescan on a controller object, all channels or backplanes attached to the controller are rescanned. If you want to rescan only a single channel or backplane, then do the rescan on the channel or backplane object.

Use this command to display new objects in the Array Manager tree view and status changes.

The following are examples of common situations where you may need to rescan the controller:

- To see new disks that were attached to the controller after a configuration was set
- So the operating system can see a new virtual disk as if it were a physical disk
- So that a new virtual disk is displayed under the Disks object in the tree view
- So that a virtual disk can use the additional space allocated to it with the Add Member command
- To update the status of an offline disk
- To display updated information in a clustered configuration after a failover of cluster resources
- To see the updated firmware version in the Array Manager console if you downloaded firmware using an application other than Array Manager.

To Rescan a Controller, Channel, or Backplane

1. Right-click a controller, channel, or backplane.
2. Click **Rescan**.

After the rescan is finished, the console is refreshed and any configuration changes appear under the Array Group object and under the channel or backplane object.

For PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

The removal of an array disk within a virtual disk is reflected immediately in the console, whereas adding an array disk to a virtual disk is reflected only by performing a rescan. Array disks that are not part of virtual disks will not be polled for status changes.

Note On PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers, displaying a status change of a virtual disk or one of its member array disks requires that you do I/O on the controller on which the virtual disk resides.

Enable Alarm

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/DC, 4e/DC, and CERC ATA100/4ch controllers

Use this function to enable the controller alarm setting. When enabled, the alarm sounds in the event of a device failure.

To Enable Alarm

1. Right-click a controller.
2. Click **Enable Alarm**.

Disable Alarm

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/DC, 4e/DC, and CERC ATA100/4ch controllers

Use this function to disable the alarm. If disabled, the alarm does not sound in the event of a device failure.

To Disable Alarm

1. Right-click a controller.
2. Click **Disable Alarm**.

Quiet Alarm

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/DC, 4e/DC, and CERC ATA100/4ch controllers

Use this function to quiet the alarm once it is sounding. Once quieted, the alarm is still enabled in the event of a future device failure.

To Quiet Alarm

1. Right-click a controller.
2. Click **Quiet Alarm**.

Rebuild Rate

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers

Use this command to change the rebuild rate setting. The rebuild rate refers to how much of the system's resources are dedicated to rebuilding a failed array disk when a rebuild is necessary. A rebuild rate of 100% means that the system resources are completely dedicated to rebuilding the failed array disk. You can improve system performance during a rebuild by lowering the rebuild rate. This means that the rebuild will take longer to complete, however.

To Rebuild a Failed Array Disk

1. Right-click a controller.
2. Click **Rebuild Rate**. This action displays the Rebuild Rate dialog.
3. Adjust the setting on the **Rebuild Rate** dialog as appropriate.

Initialize Array Disks

This section applies to PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

Use this function to initialize one or more array disks simultaneously. Array disks must be initialized before they can be used by Array Manager.

To Initialize Array Disks

1. Right-click a controller.
2. Click **Initialize Array Disks**.
3. Click **OK** after selecting which array disks you want to initialize.

Export Log File

This section applies to PERC 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

This command extracts a log from the controller(s) and places it in the Array Manager installation directory. When you export the log file from the controller, the controller's event information is cleared. The log gives detailed information on the controller activities and can be useful for troubleshooting. The log's file name is the word "log" and the month and day, such as "log0921" or "log_0921" for a log created on September 21.

To Export the Log File from the Controller

1. Right-click a controller.
2. Click **Export Log File**.

Note CERC SATA1.5/2s does not support the Export Log Files function on NetWare.

Controller Properties

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, 4/IM, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, CERC SATA1.5/2s, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers

To display the controller properties, do the following:

1. Right-click the controller object in the Array Manager tree view.
2. Select **Properties** from the pop-up menu.

The controller properties can vary depending on the model of the controller. Controller properties may include:

- **Name** — This field displays the name of the controller. The name is comprised of the controller model followed by a channel number.
- **Vendor** — This field displays the controller vendor's name.
- **Status** — This field displays the current status of the controller. Possible values are:
  - **Ready** — The controller is functioning normally.
  - **Degraded** — The controller has suffered a failure of a component and is operating in a degraded state.
  - **Failed** — The controller has suffered a failure of one or more components and is no longer functioning.
  - **Communication Lost** — Array Manager is unable to communicate with the controller using SES commands.
- **Type** — This field indicates the type of storage component. For controller properties, the value in this field is "Controller."
- **Firmware Version** — This field indicates the version of the controller's firmware.
- **Memory Size** — This field indicates the total memory available on the controller.
- **Rebuild Rate** — This field indicates the rebuild rate. The rebuild rate is the percentage of the system's resources dedicated to rebuilding a failed disk when a rebuild is necessary. For more information, see "Rebuild Rate."
- **Battery Status** — This field displays the current status of the battery. Possible values are:
  - **OK** — The battery is functioning normally.
  - **None** — There is no battery or the battery has been disconnected.
  - **Needs Replacement** — The battery's recharge limit has been reached and the battery needs to be replaced.
- **Charge Count** — Indicates the number of times the controller's battery has been recharged.
- **Recondition** — Activates the battery recondition cycle for the PERC 2 and PERC 3/Di controllers. For more information on reconditioning the battery on PERC 2 and PERC 3/Di controllers, see "Managing the Controller Battery."

For PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers

The battery has a recharge limit of 1100. Once this limit is reached, the battery status changes to Needs Replacement.

For PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

Properties for each of the PERC 2, PERC 2/Si, PERC 3/Si, PERC 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers appear differently. When a controller is opened in read-only mode, the console displays **Read only** in the name that appears in the left pane. If the controller is already in use, it can be opened only in read-only mode.

Note A PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, or CERC SATA1.5/2s controller will be opened in **Read Only** mode if it is already in use.

Array Groups 10

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, 4/IM, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s, controllers

An array group is a set of hard disks connected to a RAID controller from which virtual disks may be created.

In the Array Manager tree view, an Array Group object can be expanded to display virtual disks and array disks that can be used in a virtual disk.

To view array groups, expand the Arrays, subsystem, and controller objects in the tree view.

Array Group Menu Commands

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, 4/IM, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

The following sections describe pop-up menu commands for the PERC Subsystem 1 and PERC Subsystem 2 array groups. For information on the PV660F array group commands, see "The Dell PowerVault 660F Storage System."

Right-click an array group in the Array Manager tree view to display its pop-up menu commands. Because the features for each controller can vary, the array group pop-up menu may not display all of these commands.

- Create Virtual Disk
- Reset Configuration
- Show / Hide Array Disks
- Array Group Properties

Create Virtual Disk

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

For more information about creating virtual disks on PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers, see the section "Creating Virtual Disks".

Note Array Manager does not support creating an integrated mirror (virtual disk) on the PERC 4/IM controller. You can, however, create an integrated mirror using the controller BIOS. See "Integrated Mirroring and the PERC 4/IM Controller" for more information.

Reset Configuration

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

Use **Reset Configuration** to erase all information on the controller, so that you can perform a fresh configuration.

This operation will destroy any existing virtual disks on the controller. A warning is displayed before the operation is performed. You can cancel the execution at this step.

To Reset a Configuration

1. Right-click an array group and click **Reset Configuration**.
2. You will now need to reconfigure your controllers and virtual disks.

Caution Resetting a configuration permanently destroys all data on these virtual disks. If the system or boot partition resides on these virtual disks, it will be destroyed.

Show/Hide Array Disks

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

When a controller has array disks, the array disks attached to the controller can be displayed in the right pane of the Array Manager console. You can also specify that the array disks be hidden from view in the right pane.

To show or hide the display of array disks in the right pane:

1. Expand the controller object in the tree view. This action displays the lower-level objects contained in the controller.

2. Right-click the array group and click **Show Array Disks** or **Hide Array Disks** as appropriate.

When you right-click an array group object in the tree view, the pop-up menu displays either the **Show Array Disks** or **Hide Array Disks** option. Both options are not displayed on this menu at the same time. If the controller does not have array disks attached to it, then the **Show Array Disks** and **Hide Array Disks** options are not displayed.

Array Group Properties

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

This command displays the array group properties.

1. Right-click an array group.

2. Click **Properties**. This action displays the Array Group Properties dialog.

The array group properties include:

- **Name** — This field displays the name assigned to the array group.
- **Vendor** — This field displays the array group vendor's name.
- **Status** — This field displays the current status of the array group. Possible values are:
  - **Ready** — The array group is functioning normally.
  - **Degraded** — The array group has suffered a failure of a component and is operating in a degraded state.

- **Failed** — The array group has suffered a failure of one or more components and is no longer functioning.

◆ **Type** — This field indicates the type of storage component. For array group properties, the value in this field is "Array Group."

Channels and Backplanes 11

This section applies to all controllers supported by Array Manager.

A controller can have channels, backplanes, or enclosures. A channel can be connected with a SCSI cable to a SCSI device. If the channel is either IDE or SATA, then the channel can be connected directly to a device.

The devices connected to a channel can be either internal or external. Backplanes are internal devices that allow hot swappable devices to be connected. Enclosures are external devices that also allow hot swappable devices to be connected.

In the Array Manager tree view, a Channel or Backplane object can be expanded to display the disks attached to the channel or backplane.

To view channels or backplanes, expand the Arrays, subsystem, and controller objects in the tree view.

Channel Redundancy

When creating a virtual disk, it is possible to use disks attached to different channels to implement channel redundancy. This configuration might be used for disks that reside in enclosures subject to thermal shutdown. See the following for more information:

- Thermal Shutdown
- Channel Redundancy and Thermal Shutdown

Note Channel redundancy only applies to controllers that have more than one channel and that attach to an external disk enclosure.

Channel/Backplane Menu Commands

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, PERC 4/IM, 4e/Si, 4e/Di, CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, and CERC SATA1.5/6ch, CERC SATA1.5/2s, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers

The following sections describe pop-up menu commands for the PERC Subsystem 1, PERC Subsystem 2, and SCSI Subsystem channels and backplanes. For information on the PV660F channel and backplane commands, see "The Dell PowerVault 660F Storage System."

Right-click a channel or backplane in the Array Manager tree view to display its pop-up menu commands.

- Rescan Channel or Backplane
- Channel and Backplane Properties

Rescan Channel or Backplane

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, PERC 4/IM, 4e/Si, 4e/Di, CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, CERC SATA1.5/2s, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers

This operation can be used to rescan the channel to verify the connected devices or to check for devices that have been added. Performing a rescan on a channel or backplane object is similar to performing a rescan on the controller. The difference is that only the channel or backplane is rescanned when performing a rescan on the channel or backplane. For more information on rescanning, refer to "Rescan Controller."

Channel and Backplane Properties

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, PERC 4/IM, 4e/Si, 4e/Di, CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, CERC SATA1.5/2s, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers

This command displays channel or backplane properties.

To View Channel or Backplane Properties

1. Right-click the channel or backplane for which you want to view the properties.
2. Click **Properties**. This action displays the properties dialog.

Because the features for each controller can vary, your controller may not display all of these properties.

The channel or backplane properties may include:

- **Name** — This field displays the name assigned to the channel.
- **Vendor** — This field displays the channel vendor's name.
- **Status** — This field displays the current status of the channel. Possible values are:
  - **Ready** — The channel is functioning normally.
  - **Degraded** — The channel has suffered a failure of a component and is operating in a degraded state.
  - **Failed** — The channel has suffered a failure of one or more components and is no longer functioning.
- **Type** — This field indicates the type of storage component. For channel properties, the value in this field is "Channel."
- **Channel Termination** — This field indicates the size of the controller channel. Possible values are:
  - **Narrow** — Indicates an 8 bit channel.
  - **Wide** — Indicates a 16 bit channel.
- **SCSI Data Rate** — This field displays the SCSI speed.
- **Initiator ID** — Indicates the SCSI ID for the channel. The default value is 7. You can change the default value in the BIOS. In cluster mode, the value for Initiator ID is usually 6 or 7.
- **Mode** — Indicates whether the channel is operating in RAID or SCSI mode.

Enclosures 12

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, PERC 2, 2/Si, 3/Si, 3/Di, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers.

Array Manager enables management of various enclosures and their components. These enclosures include:

- Dell PowerVault 200S (PV200S)
- Dell PowerVault 201S (PV201S)
- Dell PowerVault 210S (PV210S)
- Dell PowerVault 211S (PV211S)
- Dell PowerVault 220S (PV220S)
- Dell PowerVault 221S (PV221S)

In addition to managing the array disks contained in the enclosure, you can monitor the status of the enclosure's fans, power supply, and temperature probes. You can view these components in the Array Manager tree view.

Array Manager displays the properties of the enclosure's fans, power supply, and temperature probes. Array Manager also notifies you of enclosure status changes through events that are displayed on the Events tab and logged in the Windows event log.

The following sections provide more information on the enclosure components managed by Array Manager:

- Enclosure Components in the Tree View
- Thermal Shutdown (Only applies to the PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Di, PERC 2, and 3/Di controllers.)
- Enclosure Menu Commands
- Enclosure Procedures
- Enclosure Properties
- Enclosure Array/Hard Disks

- Enclosure Fans
- Enclosure Power Supplies
- Enclosure Temperature Probes
- Enclosure Management Modules (EMMs)
- Adding and Removing Fibre Channel RAID and Enclosure Management
- Enclosure Management for the PV660F Storage System

Enclosure Components in the Tree View

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, PERC 2, 2/Si, 3/Si, 3/Di, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers.

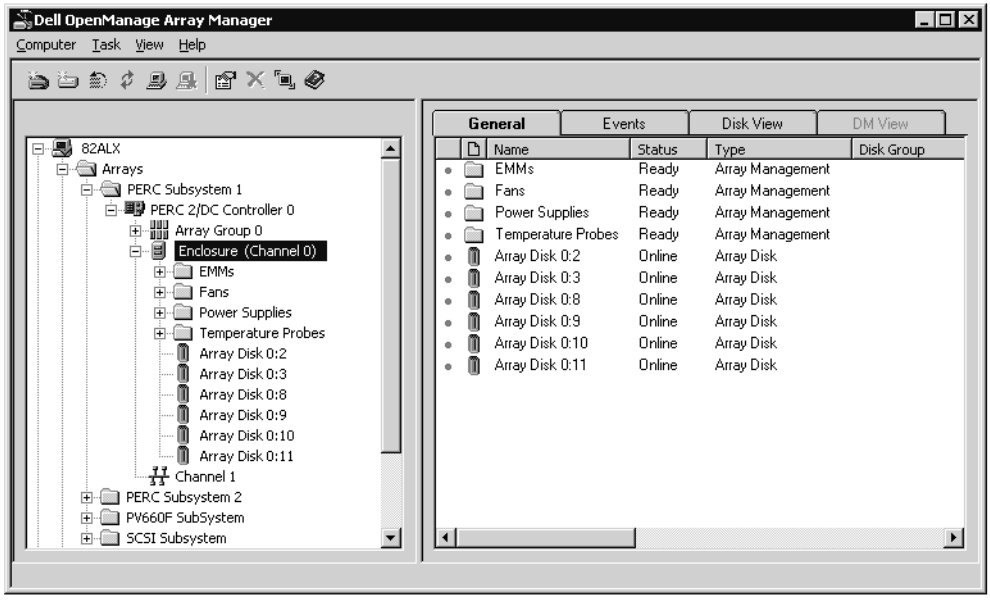
When you expand the storage subsystem object in the Array Manager tree view, you see the controllers attached to the storage subsystem. The controller object expands to display the controller's channels. These channel objects are numbered in the Array Manager tree view to match the channel number that appears on the controller hardware. (See the controller's documentation for information on identifying the channels on the controller hardware.) For example, if the controller has two channels numbered 0 and 1, then the channel objects in the Array Manager tree view are named Channel 0 and Channel 1.

If the controller has a channel connected to an enclosure, then this connection is reflected in the channel name displayed in the tree view. For example, if Channel 1 is connected to an enclosure, then the channel object's name in the tree view is Enclosure (Channel 1).

Expanding the enclosure channel object displays the enclosure's fans, power supply, and temperature probes. Each of these objects expands to display the individual fans and so forth provided by the enclosure.

In the following graphic, the enclosure object is expanded to show the EMMs, fans, power supplies, temperature probes, and array disks included in the enclosure. Because the enclosure is connected to channel 2 of the controller, the enclosure's default name is Enclosure (Channel 2).

Enclosure Management Storage Objects



For more information on the enclosure's components, commands, and the array disks included in the enclosure, see the following sections:

- Enclosure Array / Hard Disks
- Enclosure Fans
- Enclosure Power Supplies
- Enclosure Temperature Probes
- Enclosure Management Modules (EMMs)
- Enclosure Menu Commands

Thermal Shutdown

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Di, PERC 2, and 3/Di controllers.

Enclosure management provides a feature which automatically shuts down the operating system, the server, and the enclosure when the enclosure's temperature reaches dangerous extremes. The temperature when shutdown occurs is determined by the enclosure temperature probe's Minimum Error Threshold and the Maximum Error Threshold. These thresholds are default settings that cannot be changed.

Thermal Shutdown for the PV20xS and PV21xS Enclosures

For the PV20xS and PV21xS enclosures, the operating system shuts down and the server powers off when the enclosure reaches 0 degrees celsius or 50 degrees celsius. After the server powers off, the enclosure also powers off.

On NetWare, Array Manager requires that the Dell OpenManage Server Agent (OMSA) 4.0.0 or later be installed on the managed system to provide thermal shutdown for the PV20xS and PV21xS enclosures. If the Server Agent is not installed, then only the operating system shuts down when the enclosure reaches 0 degrees celsius or 50 degrees celsius.

Thermal Shutdown for the PV220S and PV221S Enclosures

For the PV220S and PV221S enclosures, the operating system and server shut down at 0 degrees celsius or 50 degrees celsius if you have not implemented Channel Redundancy and Thermal Shutdown. If you have implemented channel redundancy (and also selected the **Enable Smart Thermal Shutdown** option), then only the enclosure shuts down at 0 degrees celsius or 50 degrees celsius. This will not result in data loss because redundant data resides on the enclosure attached to the other channel.

The PV220S and PV221S enclosures automatically shut down when their temperature reaches 55 degrees celsius. This shutdown occurs whether or not you have implemented channel redundancy or have Array Manager installed.

Note On NetWare, the Dell OpenManage Server Agent (OMSA) 4.0.0 or later should be installed on the managed system in order to provide a graceful shutdown of the operating system without the risk of data loss.

For more information on the **Enable Smart Thermal Shutdown** option, see "Enable Smart Thermal Shutdown."

Channel Redundancy and Thermal Shutdown

It is possible to create a virtual disk that uses array disks that are attached to different controller channels. The array disks may reside in an external enclosure or the backplane (internal enclosure). If the virtual disk is maintaining redundant data on different channels, then the virtual disk is channel redundant. Channel redundancy means that if one of the channels fails, data will not be lost because redundant data resides on another channel.

Channel redundancy might also be used for disks that reside in enclosures subject to thermal shutdown. Should the enclosure attached to one of the channels shut down, redundant data is maintained on the other channel.

Channel redundancy is implemented by selecting array disks on different channels when using the Advanced Mode of the Create Virtual Disk wizard.

Note Channel redundancy only applies to controllers that have more than one channel and that attach to an external disk enclosure.

Considerations for Implementing Channel Redundancy on PERC 3/DCL, 3/DC, 3/QC, 4/DC, 4e/DC, 4/Di, and 4e/Di controllers.

- It is recommended that you use a RAID 10 or RAID 50 when implementing channel redundancy on the PERC 2/DC, 3/DCL, 3/DC, 3/QC, 4/Di, and 4e/Di controllers.
- If one of the channels (or enclosure) fails, you will no longer be able to maintain I/O operations on any of the channels included in the channel redundant configuration even though data has not been lost. To restore I/O, do one of the following:
  - Fix the failed channel and reboot the system.
  - Reboot the system. This will restore I/O on the channels that have not suffered a failure. The virtual disks on these channels will be in a degraded state.

Creating a Channel-redundant Virtual Disk on PERC 2/DC, 3/DCL, 3/DC, 3/QC, 4/DC, 4e/DC, 4/Di, 4e/Di, PERC 2, and 3/Di controllers.

The following instructions describe creating a virtual disk that uses channel redundancy.

1. If the Quick Access window is not already displayed, display this window by selecting the **Quick Access Window** option from the **View** pull-down menu.
2. Click the **Create Virtual Disk** button on the Quick Access Window.
3. Click the **Advanced Mode** button.

4. Specify the virtual disk properties. These include the name, RAID type, stripe size, and read, write, and cache policy for the virtual disk. See "Virtual Disk Properties" for more information.

   For PERC 2/DC, 3/DCL, 3/DC, 3/QC, 4/DC, 4e/DC, 4/Di, and 4e/Di controllers: It is recommended that you only use RAID 10 or RAID 50.

   For PERC 2 and 3/Di controllers: It is recommended that you only use RAID 10 or Concatenation of RAID-1.

5. Click **Uncheck All**. This action deselects all array disks displayed in the Available Disks portion of the wizard. Deselecting the disks enables you to make your own custom selection.

6. Depending on the controller type:

   For PERC 2/DC, 3/DCL, 3/DC, 3/QC, 4/DC, 4e/DC, 4/Di, and 4e/Di controllers:

   Use the tabs in the Available Disks portion of the wizard to select array disks on different channels. There are specific RAID level and configuration requirements for implementing channel redundancy. You must select the same number of array disks on each channel that you use. Refer to the "Number of Array Disks per Virtual Disk" section for information on the number of array disks that can be used for different RAID levels. Refer to the "Supported RAID Levels (PERC Subsystem 1)" section for information on controller-specific implementations of the RAID levels.

   RAID 10

   a. Select one array disk on each of two channels. (In other words, each of the two channels will have a single disk selected.)

   b. Click **Add Span**.

   c. Select an additional disk on each of the two channels and click **Add Span**. You have now selected the minimum number of disks for a RAID 10.

   d. Repeat step c until you have the desired number of disks.

   e. Click **Finish**.

   RAID 50

   a. Select one array disk on each of three channels. (In other words, each of the three channels will have a single disk selected.)

   b. Click **Add Span**.

c. Select an additional disk on each of the three channels and click **Add Span**. You have now selected the minimum number of disks for a RAID 50. Continue selecting a disk on each channel until you have selected the desired number of disks.

d. Repeat step c until you have the desired number of disks.

e. Click **Finish**.

For PERC 2 and 3/Di controllers:

RAID 10 or Concatenation of RAID-1

a. Select one array disk on each of two channels. (In other words, each of the two channels will have a single disk selected.)

b. Click **Add Span**.

c. Select an additional disk on each of the two channels and click **Add Span**. You have now selected the minimum number of disks for a RAID 10 or Concatenation of RAID-1.

d. Repeat step c until you have selected the desired number of disks. You must select the same number of array disks on each channel that you use.

e. Click **Finish**.

Related Information

For information related to the enclosure's temperature, see the following:

- Enclosure Temperature Probes
- Set Thresholds for Temperature
- Checking the Enclosure's Temperature
- Enable Smart Thermal Shutdown

Enclosure Array/Hard Disks

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, PERC 2, 2/Si, 3/Si, 3/Di, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers.

The enclosure's array disks or hard disks are displayed under the enclosure object in the tree view. Selecting a disk in the tree view displays the disk's status information in the Tabs.

If you right-click an array disk in the tree view, the Array Disks context menu is displayed.

Enclosure Fans

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, PERC 2, 2/Si, 3/Si, 3/Di, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers.

The fans that are installed in the enclosure are displayed under the Fans object in the tree view. You can select the Fans object to display the individual fans and their status information in the right pane. You can also expand the Fans object to display the individual fans in the left pane.

If you right-click the Fans object or an individual fan, a context menu is displayed with a Properties option. For more information on fan properties, see "Fan Properties."

Enclosure Power Supplies

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, PERC 2, 2/Si, 3/Si, 3/Di, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers.

The power supplies that are installed in the enclosure are displayed under the Power Supplies object in the tree view. You can select the Power Supplies object to display the individual power supplies and their status information in the right pane. You can also expand the Power Supplies object to display the individual power supplies in the left pane.

If you right-click the Power Supplies object or an individual power supply, a context menu is displayed with a Properties option. For more information on power supply properties, see "Power Supply Properties."

Enclosure Temperature Probes

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Di, PERC 2, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers.

The temperature probes that are installed in the enclosure are displayed under the Temperature Probes object in the tree view. You can select the Temperature Probes object to display the individual temperature probes and their status information in the right pane. The status information in the includes the current temperature in celsius and the warning and error thresholds for the temperature probe. The error threshold has a default value that cannot be changed. You can set the warning threshold, however. See "Set Thresholds for Temperature" for information on setting the warning threshold.

Right-clicking the temperature probes object in the left pane displays a context menu with a Properties option. You can also expand the temperature probe object to display the individual temperature probes in the tree view. Right-clicking an individual temperature probe also displays a context menu with a Properties option. This option enables you to set the minimum and maximum warning threshold for the temperature probe.

For information related to the enclosure's temperature, see the following:

- Checking the Enclosure's Temperature
- Set Thresholds for Temperature
- Temperature Properties

Enclosure Management Modules (EMMs)

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, PERC 2, 2/Si, 3/Si, 3/Di, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers.

The Enclosure Management Modules (EMMs) that are installed in the enclosure are displayed under the EMMs object in the tree view. You can select the EMMs object to display the individual EMM modules and their status information in the right pane. You can also expand the EMMs object to display the individual EMM modules in the left pane.

The enclosure may contain one or more EMMs. The EMM modules monitor components of the enclosure. These components include:

- Fans
- Power supplies
- Temperature probes

- The insertion or removal of a drive
- The LEDs on the enclosure

When the enclosure's alarm is enabled, the EMM activates the alarm when certain conditions occur. For more information on enabling the alarm and the conditions that activate the alarm, see "Enable Alarm." For more information on EMMs, see the enclosure hardware documentation.

If you right-click the EMMs object or an individual enclosure management module, a context menu is displayed with a Properties option. You can view the properties of each individual EMM module to verify that all EMM modules are at the same firmware level. (All EMM modules in the enclosure should have the same version of firmware.) For more information on enclosure management module properties, see "EMM Properties."

For information on downloading firmware for an EMM module, see "Download Firmware."

Enclosure Menu Commands

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, PERC 2, 2/Si, 3/Si, 3/Di, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers.

This section describes the commands associated with the enclosure and its fans, power supplies, and temperature probes.

Right-clicking an enclosure object in the tree view displays a context menu with the enclosure commands. Right-clicking the enclosure's fans, power supplies, and temperature probes also displays context menu for these components.

The commands on the enclosure's context menu can vary depending on the model of the enclosure. Enclosure context menu commands may include:

- Rescan — Checks whether any new array disks and other components such as fans or temperature probes have been added to the enclosure.
- Enable Alarm — Enables an audible alarm that sounds whenever the fault LED lights.
- Disable Alarm — Turns off the audible alarm settings. If the alarm is already sounding, you can turn it off with this command.
- Set Tag Data — Allows you to enter or change asset information for the enclosure.
- Download Firmware — Allows you to download firmware to the enclosure.
- Enclosure Properties — Displays enclosure properties.

- Set Thresholds for Temperature — Sets the minimum and maximum values for the temperature warning threshold. This command is not located on the enclosure's context menu. It is located on the context menu for the individual temperature probes.

The context menu commands for fans, power supplies, and temperature are:

- EMM Properties — Displays enclosure management module (EMM) properties.
- Fan Properties — Displays fan properties.
- Power Supply Properties — Displays power supply properties.
- Temperature Properties — Displays temperature properties. To set the temperature threshold, expand the temperature probes object to display the individual temperature probes in the tree view. The properties for an individual temperature probe enable you to set the warning thresholds.

Rescan

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, PERC 2, 2/Si, 3/Si, 3/Di, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers.

Use this command to rescan the enclosure. The Rescan command identifies any new array disks, fans, temperature probes, and so forth that have been added to the enclosure.

To rescan the enclosure:

1. Right-click the enclosure object in the tree view.
2. Select **Rescan** from the context menu.

Note If you downloaded firmware using an application other than Array Manager, rescan the controller to see the updated firmware version in the Array Manager console.

Enable Alarm

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4e/Si, 4e/Di, PERC 2, 2/Si, 3/Si, 3/Di, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers.

Use this command to enable the alarm. When enabled, the audible alarm sounds whenever the fault LED lights. This may occur to signal events such as:

- The enclosure temperature has exceeded the warning threshold
- A power supply, fan, or enclosure management module (EMM) has failed
- The split bus module is not installed.

To enable the enclosure alarm:

1. Right-click the enclosure object in the tree view.
2. Select **Enable Alarm** from the context menu.

Disable Alarm

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4e/Si, 4e/Di, PERC 2, 2/Si, 3/Si, 3/Di, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers.

Use this command to disable the enclosure alarm. When disabled, the alarm does not sound when the enclosure exceeds a warning threshold for temperature or experiences other error conditions such as a failed fan, power supply, or controller. If the alarm is already sounding, you can turn it off with this function.

To disable the alarm:

1. Right-click the enclosure object in the tree view.
2. Select **Disable Alarm** from the context menu.

Set Tag Data

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, PERC 2, 2/Si, 3/Si, 3/Di, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers.

Use this command to specify the asset tag information and the enclosure's name. When you specify the enclosure's name, the new name is displayed for the enclosure object in the tree view. You can also use this procedure to view the service tag number for the enclosure.

To specify the asset information or view the enclosure's service tag number:

1. Right-click the enclosure object in the tree view.

2. Select **Set Tag Data** from the context menu. This action displays the **Set Tag Data** dialog. This dialog has the following fields: two entry boxes which can be modified: **Asset Tag** and **Asset Name**.

   - **Service Tag** — Displays the enclosure's service tag number. This number is required when talking to Dell support. You could also use this number to identify a specific enclosure in your environment by matching this number with the number on the service tag affixed to the enclosure. For example, if you have multiple enclosures in your environment, and one of them experiences a failure, you can use the service tag number to identify which enclosure has failed. You cannot modify this field.
   - **Asset Tag** — Displays the asset tag information for the enclosure. You can modify this field to display information (such as an inventory number) useful for your environment.
   - **Asset Name** — Displays the name assigned to the enclosure. You can modify this field. The name specified in the **Asset Name** field is also the name for the enclosure displayed in the Array Manager tree view.

3. Specify the new values as required for the **Asset Tag** and the **Asset Name**.

4. Click **Apply** and then **Ok** to make the changes. Click **Cancel** to exit the function without accepting the changes.

Download Firmware

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, PERC 2, 2/Si, 3/Si, 3/Di, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers.

Use this command to download firmware to the enclosure. When you use this command in a joined bus configuration, the firmware is downloaded to both of the EMM modules in the enclosure. If you are using split bus or cluster configuration, then you need to download the firmware to each channel to ensure that the firmware on both EMM modules is at the same level.

This command assumes that the firmware is available on your system. If the firmware is not available on your system, it can be downloaded from the Dell support site at support.dell.com.

If the enclosure is attached to a clustered controller, it is required that you shut down other systems in the cluster before downloading the firmware. See "Downloading Firmware and Cluster Controllers" for more information.

On the PV220S and PV221S enclosures, you should stop I/O before downloading the firmware.

Note If you downloaded firmware using an application other than Array Manager, rescan the controller to see the updated firmware version in the Array Manager console.

Note The Download Firmware function is only available on the PV220S and PV221S enclosures.

To download firmware to the enclosure:

1. Right-click the enclosure object in the tree view.
2. Select **Download Firmware** from the context menu. This action displays the **Firmware Download** dialog.
3. Enter the path to the firmware in the **Open** field. You can also use the **Browse** button to locate the firmware.
4. Click **Download Firmware** once the correct path to the firmware is displayed in the **Open** field. Note that if the unit is in split backplane mode and the two sides are connected to different servers, then the error message "EMM in Unknown state" might appear. You should ignore this message. The state of the EMM will return to normal after the firmware download has completed.

Error Messages Associated with Downloading Firmware

Downloading firmware on PV220S and PV221S enclosures may generate the following error messages:

- **Unsafe removal of device.** Windows generates this message because it views downloading firmware as removing a device (the old firmware).
- **EMM in Unknown state.** If the enclosure is in split backplane mode and the two sides are connected to different servers, then this error message may display. You can ignore this message. The state of the EMM will return to normal after the firmware download has completed.

Set Thresholds for Temperature

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Di, PERC 2, 3/Di, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers.

Use this command to set the warning threshold for the enclosure's temperature probes.

1. Expand the enclosure object in the tree view. This action displays the lower-level objects contained in the enclosure.

2. Expand the **Temperature Probes** object in the tree view. This action displays the individual temperature probes.

3. Right-click the temperature probe for which you are setting the warning threshold and select **Properties** from the context menu. This action displays the **Temperature Probe Properties** dialog.

   The **Temperature Probe Properties** dialog displays the current status of the temperature probe as well as the default values in celsius for the minimum and maximum error thresholds. The default values for the error thresholds indicate that the enclosure has reached a dangerous temperature that may cause damage. These default values cannot be changed.

4. Enter the desired value in celsius for the minimum and maximum warning thresholds in the **Minimum Warning Threshold** and the **Maximum Warning Threshold** fields. The warning thresholds indicate that the enclosure is approaching an unacceptably warm or cool temperature.

   If you wish to apply your changes to all temperature probes in the enclosure, then select the **Apply these values to all temperature probes in the enclosure** check box.

5. Click **Apply** and then **Ok** to make the changes. Click **Cancel** to exit the function without accepting the changes.

For information related to the enclosure's temperature, see the following:

- Enclosure Temperature Probes
- Temperature Properties
- Checking the Enclosure's Temperature

Enclosure Procedures

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, PERC 2, 2/Si, 3/Si, 3/Di, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers.

This section describes procedures commonly used in association with enclosures. For additional information on enclosure commands and properties, see "Enclosure Menu Commands" and "Enclosure Properties."

The enclosure procedures are as follows:

- Identifying an Open Channel for the Enclosure
- Making a New Hard Drive Visible to Array Manager (does not apply to the Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers)
- Viewing Enclosure Status Information and Event Messages
- Checking the Enclosure's Temperature
- Verifying the EMM Firmware Version
- **Identifying the enclosure's service tag number** — The command that enables you to view the enclosure's service tag number from Array Manager is the Set Tag Data command. See "Set Tag Data" for more information.
- **Preparing a hard drive for removal** — The command that enables you to prepare a hard drive for removal is an array disk command. See "Prepare to Remove."

  The Prepare to Remove command does not apply to the Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers.
- **Troubleshooting** — For general information on troubleshooting procedures, see "Troubleshooting."
- **Removing the wrong drive** — You can avoid removing the wrong drive by blinking the LED display of the drive you intend to remove. See "Blink."

  If you have already removed the wrong drive, see "Recovering from Removing the Wrong Drive."

Identifying an Open Channel for the Enclosure

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, PERC 2, 2/Si, 3/Si, 3/Di, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers.

If you have not yet attached the enclosure to an open channel, you may need to identify a channel on the controller that can be used for this purpose. To identify an open channel:

1. Expand the Array Manager tree view until the controller object is expanded. When the controller object is expanded, the tree view displays the controller's channels. These channels are numbered starting from zero.

2. Identify a channel that is not attached to storage. When a controller's channel is attached to an enclosure or an array, the channel object can be expanded to display the disks included in the array or the enclosure. A channel object that cannot be expanded in the Array Manager tree view is an open channel not currently attached to storage. Array Manager displays a number for each channel. These numbers correspond to the channel numbers on the controller hardware. You can use these numbers to identify which open channel displayed in the tree view is the open channel on the controller hardware.

Making a New Hard Drive Visible to Array Manager

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, PERC 2, 2/Si, 3/Si, and 3/Di controllers.

After adding a new hard drive or array disk, you may find that the hard drive is not visible in the Array Manager tree view. To make the new hard drive visible in the tree view:

1. Right-click the computer system object in the tree view.

2. Select **Rescan** from the context menu.

Viewing Enclosure Status Information and Event Messages

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, PERC 2, 2/Si, 3/Si, 3/Di, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers.

You can view the current status of an enclosure and its components. To view status information, select the object for which you wish to see the status in the tree view. In the right pane, the **Status** field indicates the current status. The possible values are:

- **Ready** — Indicates that the enclosure or its component is functioning normally.
- **Offline** — Indicates that the enclosure or its component is not accessible. The enclosure or its component may be removed.
- **Failed** — Indicates that the enclosure or its component has experienced a failure and is no longer functioning.

For information on enclosure events displayed in the Array Manager event log, see "Events Generated by the PERC 2/SC, 3/SC, 2/DC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, 4/IM, 4e/DC, 4e/Si, 4e/Di, CERC ATA100/4ch and the PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s Controllers."

Checking the Enclosure's Temperature

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Di, PERC 2, 3/Di, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers.

To check the enclosure's temperature, select a temperature probe object in the tree view. The temperature reported by the temperature probe is displayed in celsius in the **Current Value** field in the right pane.

For information related to the enclosure's temperature, see the following:

- Enclosure Temperature Probes
- Set Thresholds for Temperature
- Temperature Properties

Verifying the EMM Firmware Version

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, PERC 2, 2/Si, 3/Si, 3/Di, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers.

Note If you downloaded firmware using an application other than Array Manager, rescan the controller to see the updated firmware version in the Array Manager console.

The firmware for all Enclosure Management Modules (EMMs) in PV220S and PV221S enclosures should be the same version. To verify which firmware version a particular EMM module in an PV220S and PV221S enclosure has, do the following.

1. Expand the Array Manager tree view until all EMM modules installed in the enclosure are displayed. This means expanding the EMMs object, since the individual EMM modules are at the level below the EMMs object in the tree view.

2. Right-click the EMM object for which you wish to verify the firmware version and select Properties. The EMM Properties dialog displays. This dialog indicates the EMM module's firmware version.

For information on downloading firmware for an EMM module, see "Download Firmware."

Enclosure Properties

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, PERC 2, 2/Si, 3/Si, 3/Di, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers.

To display the enclosure properties, do the following:

1. Expand the controller object in the Array Manager tree view.

2. Right-click the channel that is attached to the enclosure.

3. Select **Properties** from the context menu.

The enclosure properties can vary depending on the model of the enclosure. Enclosure properties may include:

- **Name** — This field displays the name assigned to the enclosure. The name is comprised of the enclosure's name followed by the controller channel number to which the enclosure is connected. You can modify the default enclosure name using the Set Tag Data command. The enclosure's name is also displayed in the tree view. For more information on the naming convention for enclosure objects in the tree view, see "Enclosure Components in the Tree View**."**
- **Vendor** — This field displays the enclosure vendor's name.
- **Status** — This field displays the current status of the enclosure. Possible values are:
  - **Ready** — The enclosure is functioning normally.
  - **Degraded** — The enclosure has suffered a failure of one of its component and is operating in a degraded state.
  - **Failed** — The enclosure has suffered a failure of one or more of its components and is no longer functioning.
  - **Communication Lost** — Array Manager is unable to communicate with the enclosure using SES commands. This status is displayed when the enclosure does not respond to a status query from Array Manager for any reason. For example, disconnecting the SCSI cable would cause this status to be displayed. When communication is lost, Array Manager reports the associated drive status as indicated by the controller. In the case of a disconnected SCSI cable, the drive status might be "offline" or "removed" depending on the controller. Array Manager logs an event indicating that communication with the enclosure has been lost as well as events associated with drive status changes.
- **Type** — This field indicates the type of storage component. For enclosure properties, the value in this field is "Enclosure."
- **Channel Termination** — This field indicates the size of the channel to which the enclosure is attached. Possible values are:
  - **Narrow** — Indicates an 8 bit channel.
  - **Wide** — Indicates a 16 bit channel.
- **SCSI Data Rate** — This field displays the SCSI speed.
- **Initiator ID** — Indicates the SCSI ID for the channel. The default value is 7. You can change the default value in the BIOS. In cluster mode, the value for Initiator ID is usually 6 or 7.
- **Mode** — Indicates whether the channel is operating in RAID or SCSI mode.
- **Channel** — This field displays the number of the controller channel connected to the enclosure. This number matches the channel number on the controller hardware.
- **Product ID** — This field displays the product ID of the enclosure.
- **Enclosure ID** — This field displays the enclosure processor's SCSI ID.

- **Enclosure Configuration** — This field indicates whether the mode in which the enclosure is operating. Possible values are:
  - **Joined Bus** — Indicates that the enclosure is operating in joined-bus mode.
  - **Split Bus** — Indicates that the enclosure is operating in split-bus mode.
  - **Clustered** — Indicates that the enclosure is operating in cluster mode.

  For more information on joined, split, and cluster modes, see the enclosure hardware documentation. For information on how to cable the enclosure to accommodate these different modes, see the enclosure's hardware documentation.
- **Application Version** — This field indicates the version of the upgradable firmware.
- **Kernel Version** — This field indicates the version of the non-upgradable firmware.
- **Split Bus Part Number** — This field displays the part number of the enclosure's split bus module.
- **Backplane Part Number** — This field displays the part number of the enclosure's backplane.
- **Serial Number** — Indicates the enclosure's serial number.

You can also display the properties for the following enclosure components:

- EMM Properties
- Fan Properties
- Power Supply Properties
- Temperature Properties

EMM Properties

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, PERC 2, 2/Si, 3/Si, 3/Di, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers.

To display the EMM properties, right-click the EMM object in the tree view and select Properties from the context menu.

The EMM properties can vary depending on whether you select the EMMs folder in the tree view or expand the EMMs folder and select an individual EMM. The EMM properties can also vary depending on which enclosure you are using.

EMM properties may include:

- **Name** — This field indicates the name of the EMM object as displayed in the tree view.

- **Vendor** — This field displays the EMM module's name.
- **Status** — This field displays the current status of the EMM. Possible values are:
  - **Ready** — The EMM is functioning normally.
  - **Degraded** — The EMM has reached a warning threshold.
  - **Failed** — The EMM has suffered a failure and is no longer functioning.
  - **Communication Lost** — Array Manager is unable to communicate with the enclosure using SES commands. This status is displayed when the enclosure does not respond to a status query from Array Manager for any reason. For example, disconnecting the SCSI cable would cause this status to be displayed.
- **Type** — When the EMMs folder is selected, the value in this field is Array Management which means the folder contains items that can be managed with Array Manager. When an EMM module is selected, this field indicates that the enclosure component is an EMM module.
- **Firmware** — This field indicates the version of the firmware loaded on the EMM module. All EMM modules in the enclosure should have firmware that is at the same level. You can verify the firmware version by opening the Properties of each individual EMM module and comparing the values in the Firmware field. For information on downloading firmware to a module, see "Download Firmware."
- **Part Number** — This field displays the part number of the EMM module.
- **Max SCSI Data Rate** — The maximum SCSI speed that the EMM supports.

Fan Properties

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, PERC 2, 2/Si, 3/Si, 3/Di, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers.

To display the fan properties, right-click the fan object in the tree view and select Properties from the context menu.

The fan properties can vary depending on whether you select the Fans folder in the tree view or expand the Fans folder and select an individual fan.

Fan properties may include:

- **Name** — This field indicates the name of the fan object as displayed in the tree view.
- **Vendor** — This field displays the fan vendor's name.
- **Status** — This field displays the current status of the fan. Possible values are:
  - **Ready** — The fan is functioning normally.

- **Failed** — The fan has suffered a failure and is no longer functioning.
- **Offline** — The fan has been removed.
- **Communication Lost** — Array Manager is unable to communicate with the enclosure using SES commands. This status is displayed when the enclosure does not respond to a status query from Array Manager for any reason. For example, disconnecting the SCSI cable would cause this status to be displayed.

◆ **Type** — When the Fans folder is selected, the value in this field is Array Management which means the folder contains items that can be managed with Array Manager. When an individual fan is selected, this field indicates that the enclosure component is a fan.

◆ **Speed** — This field indicates the fan speed. Possible values are Fast, Medium, Slow, and Stopped.

Power Supply Properties

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, PERC 2, 2/Si, 3/Si, 3/Di, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers.

To display the power supply properties, right-click the power supply object in the tree view and select Properties from the context menu.

The power supply properties can vary depending on whether you select the Power Supplies folder in the tree view or expand the Power Supplies folder and select an individual Power Supply. Power Supply properties may include:

◆ **Name** — This field indicates the name of the power supply object as displayed in the tree view.

◆ **Vendor** — This field displays the power supply vendor's name.

◆ **Status** — This field displays the current status of the power supply. Possible values are:

- **Ready** — The power supply is functioning normally.
- **Failed** — The power supply has suffered a failure and is no longer functioning.
- **Offline** — The power supply has been removed. When a power supply is removed, the power supply's object in the Array Manager tree view is no longer displayed.

- **Communication Lost** — Array Manager is unable to communicate with the enclosure using SES commands. This status is displayed when the enclosure does not respond to a status query from Array Manager for any reason. For example, disconnecting the SCSI cable would cause this status to be displayed.

- **Type** — When the Power Supplies folder is selected, the value in this field is Array Management which means the folder contains items that can be managed with Array Manager. When an individual power supply is selected, this field indicates that the enclosure component is a power supply.

Temperature Properties

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Di, PERC 2, 3/Di, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers.

To display general properties for all of the temperature probes, right-click the top-level temperature probe object (the level directly under the controller object) and select Properties from the context menu. To set the temperature thresholds, see "Set Thresholds for Temperature."

The temperature probe properties can vary depending on whether you select the Temperature Probes folder in the tree view or expand the Temperature Probes folder and select an individual temperature probe. Temperature probe properties may include:

- **Name** — This field indicates the name of the temperature probe object as displayed in the tree view.

- **Vendor** — This field displays the temperature probe vendor's name.

- **Status** — This field displays the current status of the temperature probe. Possible values are:

  - **Ready** — The temperature probe is functioning normally.

  - **Degraded** — The temperature probe has reached a warning temperature threshold.

  - **Failed** — The temperature probe has suffered a failure and is no longer functioning.

  - **Communication Lost** — Array Manager is unable to communicate with the enclosure using SES commands. This status is displayed when the enclosure does not respond to a status query from Array Manager for any reason. For example, disconnecting the SCSI cable would cause this status to be displayed.

- **Type** — When the Temperature Probes folder is selected, the value in this field is Array Management which means the folder contains items that can be managed with Array Manager. When an individual temperature probe is selected, this field indicates that the enclosure component is a temperature probe.
- **Current Value** — This field indicates the current temperature of the enclosure as reported by the temperature probe.
- **Minimum Error Threshold** — This field displays the temperature currently set as the minimum error temperature. The minimum error temperature indicates that the enclosure has become so severely cooled that loss of data or other failure may occur. When this temperature is reached, the operating system that writes data to the enclosure automatically shuts down. This field cannot be modified.
- **Minimum Warning Threshold** — This field displays the temperature currently set as the minimum warning temperature. The minimum warning temperature indicates that the enclosure is approaching an unacceptably cool temperature. Action should be taken to restore the enclosure to a normal temperature. Otherwise, data loss or other failure may occur.
- **Maximum Warning Threshold** — This field displays the temperature currently set as the maximum warning temperature. The maximum warning temperature indicates that the enclosure is approaching an unacceptably warm temperature. Action should be taken to restore the enclosure to a normal temperature. Otherwise, data loss or other failure may occur.
- **Maximum Error Threshold** — This field displays the temperature currently set as the maximum error temperature. The maximum error temperature indicates that the enclosure has become so overheated that loss of data or other failure may occur. When this temperature is reached, the operating system that writes data to the enclosure automatically shuts down. This field cannot be modified.
- **Apply these values to all temperature probes in this enclosure** — When this check box is selected, Array Manager applies your changes to all temperature probes in the enclosure.
- **Reset Defaults** — Click this button to reset the **Minimum Warning Threshold** and the **Maximum Warning Threshold** fields to their default values.

For information related to the enclosure's temperature, see the following:

- Enclosure Temperature Probes
- Set Thresholds for Temperature
- Checking the Enclosure's Temperature

Virtual Disks 13

This section applies to PERC 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

In order to implement RAID functions, RAID controllers must create a virtual disk. A virtual disk refers to storage created by a RAID controller from one or more array disks. Although a virtual disk may be created from several array disks, it is seen by the operating system as a single disk. Depending on the RAID level used, the virtual disk may retain redundant data in case of a disk failure or have particular performance attributes. See "Understanding Storage Concepts" for more information.

In the Array Manager tree view, a virtual disk object can be expanded to display the array disks included in the virtual disk.

To view virtual disks, expand the Arrays, subsystem, controller, and Array Groups objects in the tree view.

Note Virtual disks can only be created on a RAID controller.

Protecting Data with a Hot Spare

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

When you create a redundant virtual disk using a RAID controller, you have the opportunity to maintain system operations even when a disk fails. To do so, you would assign a hot spare to the virtual disk. When a disk fails, the redundant data is rebuilt onto the hot spare without interrupting system operations. More information regarding hot spares and the procedures for implementing them is included in the controller chapter. See the following for more information:

- Understanding Hot Spares
- Assign Global Hot Spare

- Configure Dedicated Hot Spare

Channel Redundant Virtual Disks

When creating a virtual disk, it is possible to use disks attached to different channels to implement channel redundancy. This configuration might be used for disks that reside in enclosures subject to thermal shutdown. See the following for more information:

- Thermal Shutdown
- Channel Redundancy and Thermal Shutdown

Note Channel redundancy only applies to controllers that have more than one channel and that attach to an external disk enclosure.

Creating Virtual Disks

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

In order to implement RAID functions, you need to create a virtual disk. A virtual disk refers to storage created by a RAID controller from one or more array disks. Although a virtual disk may be created from several array disks, it is seen by the operating system as a single disk. See "What Is RAID?" for more information.

Depending on your controller, refer to the following sections:

- For PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch, see "Virtual Disk Considerations for PERC Subsystem 1 Controllers," and then continue with "To Create a Virtual Disk."
- For PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers, see "Virtual Disk Considerations for PERC Subsystem 2 Controllers," and then continue with "To Create a Virtual Disk."
- For a PERC 4/IM controller, see "Integrated Mirroring and the PERC 4/IM Controller."
- For the Dell PowerVault 660F Storage System, see "Create Virtual Disk (PV660F)."

You may also wish to refer to the following related sections:

- "What Is RAID?" in the "Understanding Storage Concepts" chapter.
- "Channel Redundancy" for information on creating a channel-redundant virtual disk

- Deleting Virtual Disks
- Reconfiguring / Migrating Virtual Disks
- Array and Hard Disk Menu Commands

Virtual Disk Considerations for PERC Subsystem 1 Controllers

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers

In addition to the considerations described in this section, you should also be aware of the controller limitations described in "Number of Array Disks per Virtual Disk."

Be aware that the following considerations apply when creating virtual disks:

- **Creating virtual disks on PERC 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA/4ch controllers.** When you create a virtual disk, you specify which array disks are to be included in the virtual disk. The virtual disk you create spans the specified array disks. Depending on the size of the virtual disk, the virtual disk may not consume all of the space on the array disks. Any leftover space on the array disks cannot be used for a second virtual disk unless the array disks are of equal size. In addition, when the array disks are of equal size and you can therefore use the leftover space for a second virtual disk, this new virtual disk cannot expand to include any array disks not included in the original virtual disk.
- **Space allocation when deleting and creating virtual disks on PERC 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers.** When you delete a virtual disk, you free up or make available space on the array disks that were being used by the deleted virtual disk. If you have created several virtual disks on an array group, then deleting virtual disks can result in pockets of free space residing in various locations on the array disks. When you create a new virtual disk, the controller must decide which free space on the array disks to allocate to the new virtual disk. The PERC 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers look for the largest area of free space and allocate this space to the new virtual disk.
- **Deleting virtual disks on PERC 2/SC and 2/DC controllers.** Due to hardware constraints, virtual disks on PERC 2/SC and 2/DC controllers must be deleted in reverse order of their creation. In other words, you must delete the most recently created virtual disk first and continue to delete in reverse order.
- **SCSI limitation of 2TB.** Virtual disks created on a PERC 2/SC, 3/SC, 2/DC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, or CERC ATA100/4ch controller cannot be created from array disks with an aggregate size greater than 2TB.

This is a limitation of the controller implementation. For example, you cannot select more than 30 array disks that are 73GB in size, regardless of the size of the resulting virtual disk. When attempting to select more than 30 disks of this size, a pop-up message is displayed that indicates that the 2TB limit has been reached, and that you should select a smaller number of array disks. The 2TB limit is an industry-wide SCSI limitation.

- **Expanding virtual disks.** You can only use the Add Member command to expand a virtual disk that uses the full capacity of its member array disks. For more information, see "Add Member."
- **Virtual disk names not stored on controller.** The names of the virtual disks that you create are not stored on the controller. This means that if you reboot using a different operating system, the new operating system may rename the virtual disk using its own naming conventions.
- **Creating and deleting virtual disks on cluster-enabled controllers.** There are particular considerations for creating or deleting a virtual disk from a cluster-enabled controller. Review the "Creating and Deleting Virtual Disks on Cluster-enabled Controllers" section before attempting to create the virtual disk.
- **Requirements for using the entire capacity of an array disk on PERC 2/SC and 2/DC controllers.** With PERC 2/SC and 2/DC controllers, it is recommended that whenever you create a virtual disk, you use the entire capacity of its member array disks. If you create a virtual disk that does not use all of the capacity of its member array disks, then the next virtual disk created must contain the remaining unused capacity of the array disks from which the previous virtual disk was created. When using RAID 50, you can only create virtual disks that use the maximum capacity of the member array disks.
- **Implementing channel redundancy.** A virtual disk is channel-redundant when it maintains redundant data on more than one channel. If one of the channels fails, data will not be lost because redundant data resides on another channel. For more information, see "Channel Redundancy."
- **Integrated Mirroring and the PERC 4/IM Controller.** Array Manager does not support creating an integrated mirror (virtual disk) on the PERC 4/IM controller. You can, however, create an integrated mirror using the controller BIOS. See "Integrated Mirroring and the PERC 4/IM Controller" for more information.

Virtual Disk Considerations for PERC Subsystem 2 Controllers

This section applies to PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

In addition to the considerations described in this section, you should also be aware of the controller limitations described in "Number of Array Disks per Virtual Disk."

Be aware that the following considerations apply when creating virtual disks:

- **Limitation on number of virtual disks on each array disk.** On the PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers, each array disk can be included in no more than ten virtual disks.
- **Hot spare implementation.** In order for a hot spare to effectively rebuild a virtual disk, you need to understand the requirements for creating a virtual disk and assigning hot spares. Review the "Understanding Hot Spares" and "Considerations for Hot Spares (PERC Subsystem 2)" sections before creating a virtual disk on a PERC Subsystem 2 controller.

Number of Array Disks per Virtual Disk

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

There are limitations on the number of array disks that can be included in the virtual disk. These limitations depend on the controller. When creating a virtual disk, controllers support a certain number of stripes and spans (methods for combining the storage on array disks). Because the number of total stripes and spans is limited, the number of array disks that can be used is also limited. The limitations on stripes and spans affect the possibilities for concatenation and RAID levels as follows:

- Maximum number of spans affects concatenation, RAID Level 1-Concatenated, RAID 10, and RAID 50.
- Maximum number of stripes affects RAID 0, RAID 5, and RAID 50.
- Number of array disks in a mirror is always 2. This affects RAID 1. RAID Level 1-Concatenated, and RAID 10.

In the case of RAID 50, you can use a greater number of array disks than is possible for the other RAID levels. The number of channels on the controller, however, imposes limitations on how many array disks can be included in a virtual disk when using RAID 50. This is because only a limited number of array disks can be physically attached to the controller.

PERC Subsystem 1: Array Disks per Virtual Disk

The following table describes the specific limitations for the PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers.

Number of Array Disks per Virtual Disk			
Concatenation or RAID Level	**PERC 2/SC and 2/DC**	**PERC 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di,**	**CERC ATA100/4ch**
Concatenation	1 to 8	2 to 8	1 to 40
RAID 0	1 to 8	1 to 32	1 to 40
RAID 1	2	2	2
RAID 5	3 to 8	3 to 32	3 or 40
RAID 10	An even number from 2 to 16	An even number from 2 to 16	4
RAID 50	Between 6 and the number of array disks attached to the controller. For example, if 14 array disks are attached to a 2/SC controller, then the allowable number of array disks included in a virtual disk on this controller is between 6 and 14. A prime number of array disks is not allowed in a RAID 50 array for these controllers. For example, the following prime numbers are excluded: 7, 11, 13, 17, 19, 23, 29, 31, 37, 41, 43, 47, 51, and 53.		Not supported

PERC Subsystem 2: Array Disks per Virtual Disk

The following table describes the specific limitations for the PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers.

Number of Array Disks per Virtual Disk			
Concatenation or RAID Level	**PERC 2, 2/Si, 3/Si, and 3/Di**	**CERC SATA1.5/6ch**	**CERC SATA1.5/2s**
Concatenation	1 to 32	1	1
RAID 0	1 to 32	1 to 32	2
RAID 1	2	2	2
RAID 5	3 to 16	3 to 16	N / A
RAID Level 1-Concatenated (Concatenated Mirror)	An even number from 2 to 32	N / A	N / A
RAID 10	2 to 32	2 to 32	N / A

For more information on storage concepts used in this section, see "What Is RAID?" and "Choosing RAID Levels and Concatenation."

Note In addition to this online help, you should review the hardware documentation that is provided with the controllers. Reviewing the hardware documentation along with this online help may provide a fuller understanding of the controller limitations.

Maximum Virtual Disks per Controller

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, CERC SATA1.5/2s controllers

Controllers have limitations on the number of virtual disks that they can support.

Maximum Virtual Disks on PERC Subsystem 1 Controllers

You can create a maximum of 8 virtual disks on one PERC 2/SC or PERC 2/DC controller and a maximum of 40 virtual disks on one CERC ATA100/4ch controller. On the PERC 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, and 4e/Di, controllers, you can create a maximum of 40 virtual disks.

Maximum Virtual Disks on PERC Subsystem 2 Controllers

You can create a maximum of 24 virtual disks on a PERC 2, PERC 2/Si, PERC 3/Si, PERC 3/Di, CERC SATA1.5/6ch, or CERC SATA1.5/2s controller.

To Create a Virtual Disk

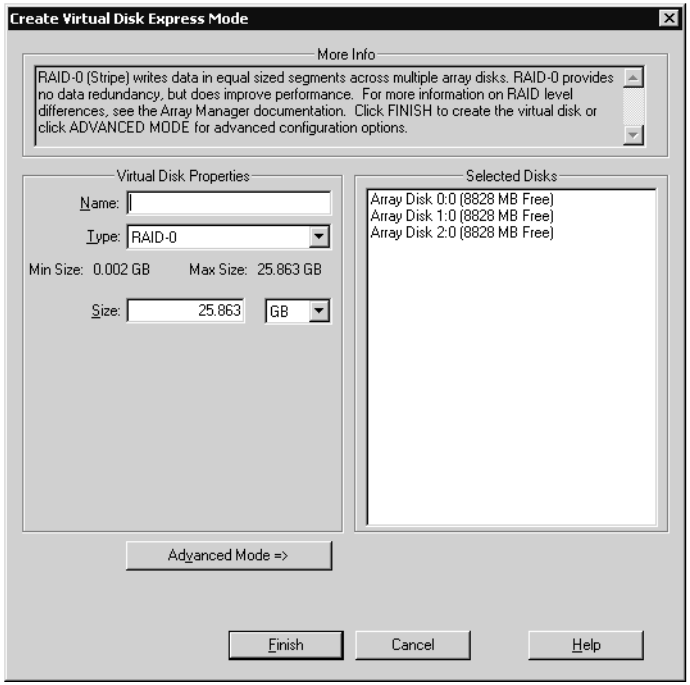
This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

1. Before creating a virtual disk, you should be familiar with the information in the following sections:
   - Virtual Disk Considerations for PERC Subsystem 1 Controllers
   - Virtual Disk Considerations for PERC Subsystem 2 Controllers
   - Number of Array Disks per Virtual Disk
   - Maximum Virtual Disks per Controller
2. If the Quick Access window is not already displayed, display this window by selecting the **Quick Access Window** option from the **View** pull-down menu.
3. Click the **Create Virtual Disk** button on the Quick Access Window. The Create Virtual Disk wizard appears. The wizard has an Express Mode and an Advanced Mode.

   #### The Create Virtual Disk Wizard – Express Mode

   Express Mode calculates an appropriate virtual disk layout based on the available space and controller considerations. Using Express Mode, you can quickly create a virtual disk using recommended selections.

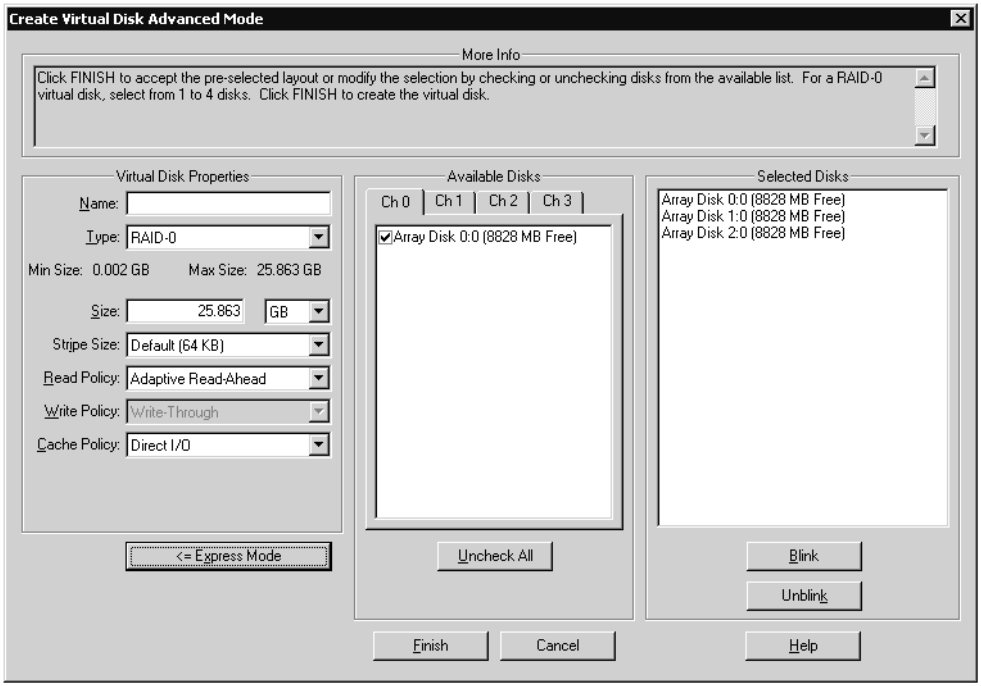
Create Virtual Disk Express Mode



The Create Virtual Disk Wizard – Advanced Mode

Advanced Mode allows you to specify the read, write, and cache policy for the virtual disk. You can also select the array disks and the controller channel to be used. You need a good knowledge of RAID levels and hardware to use this mode.

Create Virtual Disk Advanced Mode. The Read, Write, and Cache Policy options vary depending on the controller.



4. Make your selections using the Express Mode or Advanced Mode options. Click the **Advanced Mode** or **Express Mode** button to switch from one mode to another. The following describes each selection.

   Virtual Disk Properties

   - **Name** — Enter a name for the virtual disk. This name will be displayed in the Array Manager tree view. Note that Array Manager automatically numbers the virtual disks in the order that they are created. You do not need to include the number as part of the virtual disk name.

Note When creating a concatenated virtual disk on the CERC SATA1.5/2s controller, you cannot specify the virtual disk name.

   - **Type** — Specify the RAID type or level. The RAID levels that can be used depend on the number of available array disks and on the controller. If you are using Express Mode, Array Manager displays the possible RAID levels depending on the available array disks. If you are using Advanced Mode, Array Manager displays the possible RAID levels depending on which array disks you selected.

In Advanced Mode, if the RAID level you want to use is not displayed, you may need to change your array disk selection. Information on the selected RAID level is displayed at the top of the wizard. For additional information on RAID levels, see "Choosing RAID Levels and Concatenation" and "What Is RAID?" in the "Understanding Storage Concepts" chapter.

- **Size** — Specify the size that the virtual disk should be in either megabytes or gigabytes. The minimum and maximum size possible is displayed above the Size field. When using Express Mode, the **Query Best Fit Disk** button may display after you change the size of the virtual disk. Click the **Query Best Fit Disk** button to select array disks that will be most closely used to their full capacity by the new virtual disk. The **Query Best Fit Disk** button enables more efficient use of disk space. This option is only displayed in Express Mode when there are different possibilities for array disk selection.

Note This note applies to PERC 2/SC, PERC 2/DC, PERC 3/SC, PERC 3/DC, PERC 3/DCL, PERC 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers. When you create a RAID 10 or RAID 50 virtual disk, the size option is redundant, as the maximum size will be created. When you click **Finish**, the Size text box displays the maximum size.

- **Stripe Size** — (Advanced Mode). Specify the stripe size for the virtual disk. The stripe size is the amount of array disk space used on each array disk included in the stripe.
- **Read Policy** — (Advanced Mode). Specify the SCSI read policy for the logical drive. The Read Policy options vary depending on which controller you have. See the following for more information:

  Read policy for PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers

  Read policy for PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers
- **Write Policy** — (Advanced Mode). Specify the SCSI write policy for the logical drive. The Write Policy options vary depending on which controller you have. See the following for more information:

  Write policy for PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers

  Write policy for PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

Note Write-through caching is the default write policy setting when cluster mode is enabled. In cluster mode, the PERC 3/DC controllers only allow write-through caching.

- **Cache Policy** — (Advanced Mode). Specify the SCSI cache policy for the logical drive. Cache policies apply to the PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers. See "Cache Policy" for more information.

Selected Disks

When using Express Mode, the wizard selects disks for you depending on the RAID level you selected and on the controller. When using Advanced Mode, the wizard displays the disks that you selected from the Available Disks section of Advanced Mode. Use the **Blink** and **Unblink** buttons to identify which disks are selected by blinking their LED display.

Available Disks

When using Advanced Mode, you can select the disks to be used from the Available Disks tabs. These tabs indicate the disks attached to each controller channel that are available depending on the RAID level you selected and the controller.

5. Click **Finish** to create the virtual disk.

   If the newly created virtual disk does not appear under Disks in the tree view, then perform a rescan by selecting **Rescan** from the View pull-down menu. **Rescan** forces the Array Manager software to locate and query all the objects in the system. This can take several minutes if there are numerous disks on the system.

   When initializing is complete, the virtual disk appears in the tree view of the Array Manager console. Each virtual disk is automatically mapped to a Windows disk. For example, if Virtual disk 0 is mapped to Windows Disk 0, the Array Manager console will display Virtual disk 0 as "Virtual Disk 0 (Disk 0)," and the Windows disk 0 as "Disk 0 (Virtual Disk 0)."

   If a virtual disk is not mapped to a Windows disk, reboot the system.

Note On the PERC 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers, initialization begins automatically after you have created a redundant virtual disk.

6. Add a write signature to the virtual disk.

   When you create a virtual disk, a write signature is not automatically written to the disk. Windows requires a write signature on virtual disks. For instructions on writing a disk signature, see the section "Write a Disk Signature" in the "Disks" chapter.

Reconfiguring/Migrating Virtual Disks

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

You can reconfigure or migrate a virtual disk in order to increase the disk's capacity or change its RAID level. The procedure for reconfiguring a virtual disk varies depending on which controller you have.

- Reconfiguring Virtual Disks on PERC Subsystem 1 Controllers
- Reconfiguring Virtual Disks on PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s Controllers

Reconfiguring Virtual Disks on PERC Subsystem 1 Controllers

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers

The PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers have pop-up menu commands that enable you to expand the virtual disk, change the RAID level, and change the cache policy. Use the following commands to accomplish these tasks.

- Add Member — Expands virtual disk capacity by adding one or more array disks; also allows changing a virtual disk's RAID level.
- Transform Type — Changes the RAID level of a virtual disk.
- Change Policy — Changes the virtual disk's cache policy.

Reconfiguring Virtual Disks on PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s Controllers

This section applies to PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

Use the Virtual Disk Reconfiguration wizard to change attributes of a virtual disk on a PERC 2, PERC 2/Si, PERC 3/Si, PERC 3/Di, CERC SATA1.5/6ch, or CERC SATA1.5/2s controller. You can:

- Migrate from one RAID type to another without adding space.

- Add space without changing the current RAID type.
- Migrate and add space.
- Specify a new stripe size.
- Move the virtual disk to other array disks.

To Reconfigure a Virtual Disk

Use this procedure to reconfigure a virtual disk on a PERC 2, PERC 2/Si, PERC 3/Si, PERC 3/Di, CERC SATA1.5/6ch, or CERC SATA1.5/2s controller:

1. Right-click a virtual disk to reconfigure and click **Reconfigure**.
2. The Virtual Disk Reconfiguration window appears. Make the appropriate changes for the desired configuration.
3. Click **OK** to continue or **Cancel** to quit the operation.

Virtual Disk Menu Commands

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4/IM, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers

The following sections describe pop-up menu commands for the PERC Subsystem 1 and PERC Subsystem 2 virtual disks. For information on the PV660F virtual disk commands, see "The Dell PowerVault 660F Storage System."

Right-click a virtual disk in the Array Manager tree view to display its pop-up menu commands. Because the features for each virtual disk can vary, the virtual disk pop-up menu may not display all of these commands.

- Add Member
- Transform Type
- Deleting Virtual Disks
- Change Policy
- Rename Virtual Disk
- Cancel Background Initialization
- Configure Dedicated Hot Spare
- Remove a Dedicated Hot Spare

- Check Consistency
- Pause Consistency Check
- Resume Consistency Check
- Initialize Virtual Disks
- Format
- Restore Dead Disk Segments
- Split Mirror
- Unmirror
- Blink Virtual Disk
- Unblink Virtual Disk
- Virtual Disk Properties

Add Member

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers

Use this function to add array disks to a virtual disk. You can also change the virtual disk's RAID level.

Note If you add a member to a virtual disk that shares array disk storage with another virtual disk, you may not be able to create or delete virtual disks on the controller. Whenever you create a virtual disk that you plan to expand later with the **Add Member** command, it is recommended that you use the entire capacity of its member array disks.

Note The **Add Member** command is only available for RAID 0, RAID 1, and RAID 5.

To Add Array Disks to a Virtual Disk Using Add Member

1. Right-click a virtual disk.
2. Click **Add Member**. A Virtual Disk Add Member dialog box appears. The available disks are listed in the left pane. You can choose appropriate ones to add by selecting them and using the **Add Disk** button to move them to the right pane.
3. Choose the RAID level in the pull-down list **Virtual Disk Type**.

4. Click **OK** to continue or **Cancel** to cancel the operation.

   Add Member shows progress information in the right pane. To view the progress, click the parent of the virtual disk. The status of the virtual disk will be **Reconstructing,** and progress information will display until the **Add Member** operation is finished. At the end of the operation, the **Type** category shows the changed RAID level, if the RAID level was changed.

5. Right-click the computer object in the Array Manager tree view and select **Rescan**. A rescan is required before the virtual disk can use the extra space. On a Windows system, you also need to reboot the system in order for Windows to have access to the additional storage.

Transform Type

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers

Use the **Transform Type** function to change the RAID level or stripe size of a virtual disk without adding or removing disks. The following is a list of supported RAID level migrations:

- RAID-1 to RAID-0
- RAID-5 to RAID-0

To Transform the Virtual Disk Type

1. Right-click the virtual disk for which you want to transform the type. The Virtual Disk Transform Type dialog displays.

2. Select the RAID type.

3. Click **OK** to continue or **Cancel** to quit the operation.

4. A warning dialog appears asking you whether to continue. Click **Yes** to continue.

5. **Transform Type** displays progress information in the right pane. To view the progress, click the parent of the virtual disk. The status of the virtual disk will be **Reconstructing,** and progress information will display until the operation is finished. At the end of the operation, the **Type** category shows the changed RAID level.

Deleting Virtual Disks

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

Deleting a virtual disk permanently deletes all information contained within all volumes residing on that virtual disk.

Caution All data (including file systems, volumes, Windows disks) on the virtual disk is lost when the disk is deleted.

Because of PERC 2/SC and 2/DC controller hardware constraints, virtual disks must be deleted in the reverse order of their creation on these controllers. In other words, the last disk that you created must be deleted first.

There are particular considerations for deleting a virtual disk from a cluster-enabled PERC 3/DC controller. Review the "Creating and Deleting Virtual Disks on Cluster-enabled Controllers" section before attempting to delete the virtual disk.

To Delete a Virtual Disk

1. Right-click the virtual disk to be deleted.
2. Click **Delete**. A confirmation dialog box appears. Click **OK** to continue. The virtual disk disappears from the left pane.

   If the disk you are deleting was not the last one created, a message may display that asks you to delete the most recently created virtual disk first.
3. Reboot your system. This step is optional but highly recommended. You should reboot your system after deleting a virtual disk and before creating new ones. This ensures that the operating system recognizes the disk configuration correctly.

Note If you delete a virtual disk and immediately create a new virtual disk with all the same characteristics as the one that was deleted, the controller will recognize the data as if the first virtual disk were never deleted. In this situation, if you don't want the old data after recreating a new virtual disk, reinitialize the virtual disk.

Change Policy

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

Use this function to change the read or cache policies of a virtual disk.

To Change the Cache Policy of a Virtual Disk

1. Right-click the virtual disk and select **Change Policy** to display the Virtual Disk Change Policy dialog box.

2. Use the drop-down menus to specify the read and cache polices. For more information on the read and write policies, see "Understanding Read, Write, and Cache Policy."

Note In cluster mode, the PERC 3/DC controller only allows write-through caching.

If you wish to verify your changes, right-click the virtual disk, select **Properties,** and check the virtual disk properties.

Rename Virtual Disk

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4/IM, 4e/Si, 4e/Di, CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, and CERC SATA1.5/6ch controllers

Use this command to rename a virtual disk. Renaming a virtual disk enables you to change the virtual disk name displayed in the Array Manager tree view. The numbering format for the virtual disk remains unchanged.

Depending on the controller you have, there are different considerations regarding the controller BIOS:

- On PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4/IM, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers, changing the virtual disk name with Array Manager does not change the name in the BIOS.
- On PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers, changing the virtual disk name with Array Manager also changes the name in the BIOS. If you do not specify a name for a virtual disk (either when you create the virtual disk or by using the Rename command) then the name for the virtual disk in the BIOS will be "Virtual Disk".

1. Expand the Array Manager tree view until the virtual disk you wish to rename is displayed.

2. Right-click the virtual disk and select **Rename** from the pop-up menu. This action displays the Rename Virtual Disk dialog.

3. Type the new name for the virtual disk and click **Ok**. The virtual disk name cannot contain more than 32 characters. Only alphanumeric English characters can be used when renaming a virtual disk.

Note The **Rename** command is not available for concatenated virtual disks on a CERC SATA1.5/2s controller.

Cancel Background Initialization

This section applies to PERC 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers

On PERC 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers, background initialization of redundant virtual disks begins automatically after the virtual disk is created.

It is important that the background initialization be allowed to run to completion. You can, however, cancel the background initialization. When you do so, the controller will restart the background initialization within 0 to 5 minutes. See "Background Initialization on PERC Subsystem 1 Controllers" for more information.

To cancel background initialization, right-click the redundant virtual disk and select **Cancel Background Initialization**.

Configure Dedicated Hot Spare

This section applies to PERC 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, ERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

Use this command to assign and configure a dedicated hot spare for a redundant virtual disk. If one of the array disks belonging to a redundant virtual disk fails or is removed, then the virtual disk becomes degraded. A degraded virtual disk that has a properly configured hot spare automatically rebuilds onto that hot spare.

You can configure a dedicated hot spare for redundant virtual disks using RAID 1, 10, 5, 50, and RAID-1 concatenated.

Hot spares created with the **Configure Dedicated Hot Spare** command are dedicated to a virtual disk. They will only be used if an array disk in the virtual disk fails. When you dedicate a hot spare for a virtual disk, you may notice that other virtual disks are also assigned the same hot spare. This is because the dedicated hot spare applies to all virtual disks that use the same array disks as the virtual disk for which you are configuring the hot spare. The hot spare will rebuild the data for all redundant virtual disks that are using the failed array disk.

To remove the dedicated hot spare and return the array disk to its normal status, use the Remove a Dedicated Hot Spare procedure.

For more information on hot spares including size requirements, see "Understanding Hot Spares." For considerations regarding RAID 10 and RAID 50 virtual disks created using the controller BIOS, see "Dedicated Hot Spare Considerations for RAID 10 and RAID 50 on PERC Subsystem 1 Controllers."

Note Configuring a dedicated hot spare for a virtual disk is highly recommended.

To Configure a Dedicated Hot Spare

1. Right-click a redundant virtual disk and click **Configure Dedicated Hot Spare...**. This action displays the available array disks that can be configured as hot spares.

2. Assign the array disk you want to configure as a hot spare by highlighting the disk in the left panel and clicking the **Add Disk>>** button to move the disk to the right panel of the **Virtual Disk Dedicated Hot Spare Configuration** dialog box.

3. Click **OK**.

Note The **Configure Dedicated Hot Spare** command is not available on the PERC 2/SC, PERC 2/DC, and CERC SATA1.5/2s controllers.

Remove a Dedicated Hot Spare

This section applies to PERC 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch, controllers

Use this operation to remove a dedicated hot spare from a virtual disk. This operation returns an array disk that was previously configured as a dedicated hot spare to its normal status.

Note When removing a hot spare that was created with the **Configure Dedicated Hot Spare** menu option, be aware that multiple virtual disks may be using the hot spare. This is because the hot spare applies to all virtual disks that use the same array disks as the virtual disk for which the hot spare was created. Removing a dedicated hot spare removes the hot spare from all of these virtual disks.

To Remove a Dedicated Hot Spare

1. Right-click the redundant virtual disk for which you are removing the dedicated hot spare and click **Configure Dedicated Hot Spare...**.
2. Specify the disk or disks that you are removing as hot spares by highlighting the disk in the right panel and clicking the **<<Remove Disk** button to move the disk to the left panel of the **Virtual Disk Dedicated Hot Spare Configuration** dialog box.
3. Click **OK**.

Check Consistency

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

This operation verifies the accuracy of the redundant data in the virtual disk and rebuilds parity if necessary. You also use this function if your disk is in a degraded state. Running **Check Consistency** may be able to restore your disk to **Ready** status.

You cannot check the consistency of a non-redundant RAID type.

To Check Consistency and Restore Redundancy

1. Right-click the virtual disk to be checked.
2. Click **Check Consistency**.
3. The **Check Consistency** operation displays progress information in the right pane. To view the progress, click the parent of the virtual disk. The status of the virtual disk will be **Resynching,** and progress information will display until the operation is finished.

Note On PERC 2, 2/Si, 3/Si, 3/Di, and CERC SATA1.5/6ch controllers, you can pause the consistency check once it is in progress with the Pause Consistency Check option.

Pause Consistency Check

This section applies to PERC 2, 2/Si, 3/Si, 3/Di, and CERC SATA1.5/6ch controllers

If you have initiated a consistency check using the Check Consistency option, you can suspend or pause the consistency check with the **Pause Consistency Check** option. Once paused, you can later resume the consistency check by using the Resume Consistency Check option. A paused consistency check will also automatically resume when the system is rebooted.

To Pause the Consistency Check

- Right-click the virtual disk on which the consistency check is currently in progress, and click **Pause Consistency Check**.

Resume Consistency Check

This section applies to PERC 2, 2/Si, 3/Si, 3/Di, and CERC SATA1.5/6ch controllers

After having suspended a consistency check by using the Pause Consistency Check option, you can resume the consistency check with the **Resume Consistency Check** option.

To Resume the Consistency Check

- Right-click the virtual disk for which you have previously paused the consistency check, and click **Resume Consistency Check**.

Cancel Consistency Check

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

For PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, you can cancel a consistency check that you initiate.

On PERC 2, PERC 2/Si, PERC 3/Si, PERC 3/Di and CERC SATA1.5/6ch controllers, a consistency check will be auto-initiated when you create a redundant virtual disk. If you cancel that initial consistency check, then the virtual disk is not considered "redundant". To assure the virtual disk is redundant, you must let a consistency check run to completion.

The CERC SATA1.5/2s controllers do not have a pausing feature. You must allow the consistency check to complete or cancel the consistency check.

To Cancel the Consistency Check

- Right-click the virtual disk, and click **Cancel Consistency Check**.

Initialize Virtual Disks

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers

Use this procedure on any virtual disk on a PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controller system. The virtual disk will be formatted by this operation. Note that you need to initialize before using any virtual disk.

Caution Initializing a virtual disk destroys any data that is on the disk.

Note If you initialize a virtual disk and then notice that volume information has not been erased, reboot the server.

To Initialize a Virtual Disk

1. Right-click the virtual disk that you want to initialize.
2. Click **Initialize**. In the right pane, the status of the disk will be **Initializing** until the operation is finished.

Cancel Initialize Virtual Disks

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers

You can cancel the virtual disk initialization while it is in progress.

To Cancel the Virtual disk Initialization

1. Right-click the virtual disk for which you wish to cancel the initialization.
2. Click **Cancel Initialize**.

Format

This section applies to PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

Use this procedure to format a virtual disk on a PERC 2, PERC 2/Si, PERC 3/Si, PERC 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controller. This operation erases the files and removes the file systems on virtual disks but keeps the virtual disk configuration intact.

Caution On a NetWare system, Array Manager does not have the ability to detect whether the SYS volume or the boot partition reside on the virtual disk. Formatting a virtual disk that contains the SYS volume or boot partition will damage the system.

To Format

- Right-click a virtual disk and click **Format.**

Formatting a virtual disk on the PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers removes the Windows disk object and any associated drive mapping. After the format is complete, the disk object is restored, but it may not be mapped to a drive. To restore the drive mapping, either do a global rescan or stop and restart the Disk Management Service.

Note If you format a virtual disk that has been upgraded to dynamic, the disk functions may become inactive or the status of the disk may change to “foreign” or “offline.” You can view a disk’s status by selecting the disk’s **General** tab. When using a PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, or CERC SATA1.5/2s controller, you can resolve these problems by doing a global rescan.

Restore Dead Disk Segments

This section applies to PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

This command recovers the disk space occupied by an orphan or dead disk segment. An orphan disk segment is part of a multipartition virtual disk that cannot access another disk segment within the same virtual disk. A dead disk segment is either no longer needed by the virtual disk or is physically damaged. This operation can be performed only on failed virtual disks of redundant RAID types.

To Restore Dead Disk Segments

- Right-click a failed virtual disk and click **Restore Dead Disk Segments**.

Note You can restore only segments of a redundant RAID type.

Note This operation is not supported in NetWare.

Split Mirror

This section applies to PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

Use the **Split Mirror** command to separate mirrored data originally configured as a RAID 1 or RAID 10 virtual disk. Splitting a RAID 1 mirror creates two concatenated non-redundant virtual disks. Splitting a RAID 10 mirror creates two RAID 0 (striped) non-redundant virtual disks. Data is not lost during this operation.

Note This operation is not supported in NetWare. It is also not supported when using a dynamic virtual disk with a PERC 2, PERC 2/Si, PERC 3/Si, PERC 3/Di or CERC SATA1.5/6ch, or CERC SATA1.5/2s controller.

To Split a Mirror

- Right-click a mirrored virtual disk and click **Split Mirror**.

Caution Your virtual disk will no longer be redundant after performing a Split Mirror operation.

Unmirror

This section applies to PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

Use the **Unmirror** command to separate mirrored data and restore one half of the mirror to free space. Unmirroring a RAID 1 virtual disk results in a single, non-redundant concatenated virtual disk. Unmirroring a RAID 10 virtual disk results in a single, non-redundant RAID 0 (striped) virtual disk. Data is not lost during this operation.

Note This operation is not supported in NetWare. It is also not supported when using a dynamic virtual disk with a PERC 2, PERC 2/Si, PERC 3/Si, PERC 3/Di, CERC SATA1.5/6ch, or, CERC SATA1.5/2s controller.

To Unmirror

- Right-click the mirrored virtual disk and click **Unmirror**.

Caution Your virtual disk will no longer be redundant.

Blink Virtual Disk

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, PERC 2, 2/Si, 3/Si, and 3/Di controllers

The **Blink Virtual Disk** command allows you to locate the array disks included in a virtual disk by blinking the light-emitting diodes (LEDs) on the array disks. The LEDs could blink from 30 seconds to indefinitely.

This procedure is only supported for hotswap array disks (disks that reside in a carrier).

Note This procedure is not supported on the CERC ATA100/4ch or PERC 4/IM controllers.

Unblink Virtual Disk

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, PERC 2, 2/Si, 3/Si, and 3/Di controllers

The **Unblink Virtual Disk** command allows you to cancel the **Blink Virtual Disk** command.

This procedure is only supported for hotswap array disks (disks that reside in a carrier).

Note This procedure is not supported on the CERC ATA100/4ch or PERC 4/IM controllers.

Rebuild (CERC SATA1.5/2s Controller)

A rebuild menu option is available for RAID 1 virtual disks on the CERC SATA1.5/2s controller. The Rebuild option on this controller is only enabled when one array disk in the RAID 1 contains data that is not consistent with the data on the other array disk. This may occur if one of the array disks is removed or the cable is disconnected and the other array disk continues to update its data. See Rebuilding a Virtual Disk on the CERC SATA1.5/2s Controller for more information.

To Rebuild the Virtual Disk

- Right-click the virtual disk and click **Rebuild.**

Virtual Disk Properties

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4/IM, 4e/Si, 4e/Di, CERC ATA 1004/CH, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

This command enables you to display the virtual disk properties.

To display the virtual disk properties:

1. Right-click the virtual disk object in the Array Manager tree view.
2. Select **Properties** from the pop-up menu. This action displays the **Virtual Disk Properties** dialog.

The Virtual Disk Properties dialog has tabs for displaying the properties of the virtual disk. These properties include the read, write, and cache policies of the disk, its status, and its capacity.

Disks 14

This chapter describes conceptual and procedural information about how Array Manager implements basic and dynamic disks. It also includes general disk functions.

Note There are particular considerations regarding dynamic disks and volumes on NetWare, Windows Server 2003, and Linux. See "Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux" for more information.

Windows Operating System Disks: Basic and Dynamic

On a Windows system, you can expand the Disks object in the Array Manager tree view to display the operating system disks recognized by Windows Disk Management. These operating system disks can be either **basic** or **dynamic**. In conceptual terms, a basic disk is similar to a hard disk in that it does not support RAID functions. A dynamic disk is similar to an array disk in that it can support RAID functions. You can convert a basic disk to dynamic in order to apply RAID functions to the disk.

When you install the Windows operating system, your disks are automatically initialized as basic. You can upgrade them to dynamic after the Array Manager installation is complete. See "Upgrading a Basic Disk to Dynamic" for more information.

There are particular considerations regarding dynamic disks and volumes on NetWare, Windows Server 2003, and Linux. See "Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux" for more information.

Basic Disk Definition

A basic disk adheres to the partition-oriented scheme of a Windows operating system. A basic disk can have up to either four primary partitions or three primary partitions plus an extended partition. This extended partition can be subdivided into a number of logical drives.

You can create only primary and extended partitions on basic disks. You should create partitions on basic disks if you want computers running Microsoft Windows 98 or earlier, or MS-DOS to access these partitions.

Dynamic Disk Definition

Dynamic disks are created by upgrading basic disks using Array Manager. A dynamic disk is a physical disk that can contain dynamic volumes created with Array Manager. A dynamic volume organizes space on one or more physical disks by using a specific type of volume layout. The five types of dynamic volumes are simple, spanned, mirrored, striped, and RAID-5.

On a dynamic disk, space is organized through volumes rather than partitions. Because a dynamic disk does not have the partitioning scheme used by Windows 95/98 and MS-DOS, you cannot access dynamic disks through those operating systems.

For more information on the dynamic volume types, see the "Dynamic Volume Types" in the "Volumes" chapter.

There are particular considerations regarding dynamic disks and volumes on NetWare, Windows Server 2003, and Linux. See "Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux" for more information.

Working with Basic Disks

This section's topics are:

- Basic Disk Functions
- Basic Disk Limitations
- Create Primary and Extended Partitions
- Create Logical Drives in an Extended Partition
- Dealing with Removable Media
- Upgrading a Basic Disk to Dynamic

Basic Disk Functions

Array Manager allows you to do the following on basic disks:

- Create and delete primary and extended partitions.
- Create and delete logical drives within an extended partition.
- Format a primary partition and mark it as active.
- Check disk properties such as capacity, unallocated space, and current status.
- Make drive-letter assignments for basic volumes, partitions, CD-ROMs, DVDs, and other removable devices.

- Upgrade a basic disk to dynamic if the basic disk is empty.

Basic Disk Limitations

You cannot perform the following tasks on a basic disk:

- View or change volumes created in NT Disk Administrator.
- Extend basic volumes.

Create Primary and Extended Partitions

Use the Create Volume wizard to create primary partitions and extended partitions on a basic disk.

1. From the Task menu, select **Create Volume**. You can also select a basic disk, right-click to display the context menu, and select **Create Volume**. Both methods display the Create Volume wizard window.

2. The Create Volume wizard's first page displays information about volumes. Read the information and then click **Next** to continue.

3. You are asked to select whether to create a partition or a volume. If you have started with a basic disk, the **Partition** radio button is already selected. The group that the partition will be created in is displayed. Click **Next** to continue.

Select Partition Type

4. The Create Volume wizard asks you to select a partition type (primary or extended) and the size of the partition. Click **Next** to continue.

Note If the disk already has either four primary partitions or three primary partitions and one extended partition, Array Manager displays a message that you have exceeded the number of partitions and will not let you create any more. Click **OK** and return to the Array Manager console.

Select Disks and Region

5. Select a disk and free space on which to create the partition. The dialog box shows a list of available disks and free spaces. Check the box for the disk and free space you want to use, and click to **Next** continue.

6. If there is no free space on the selected disk, the free space selection is shaded. A message appears on the wizard that indicates a lack of available disk space. Click **Next** to continue.

Assign a Drive Letter or Drive Path

7. Assign a drive letter to this volume if you want to format the partition. (Or you can assign a drive letter afterward.) If you are using a Windows 2000 machine, you can also mount the partition at an empty NTFS folder. Click **Next** to continue.

Format

8. Select the **Format this volume** check box if you want to format the partition. Array Manager allows you to format using either FAT or NTFS. In addition, for Windows 2000 only, you can use FAT32.
   - **FAT (File Allocation Table)** — A file system originally designed for MS-DOS, Windows 3.x, and Windows 95/98. Windows 2000 and Windows Server 2003 can also use the FAT file system. The operating system maintains a table to keep track of the status of various segments of disk space used for file storage.
   - **NTFS (NT File System)** — An advanced file system designed for use specifically within the Windows NT and Windows 2000 operating systems. Use this format if you want to use file and folder compression.
   - **FAT32** — (Windows 2000 and Windows Server 2003) A derivative of the file allocation table (FAT) file system. FAT32 supports smaller cluster sizes than FAT, thus providing more efficient space allocation on FAT32 drives. FAT32 is designed for larger disks than FAT.

9. To name the partition, type a name in **Volume Label**. If you do not enter a name, your partition will be named "New Volume" by default.

10. Select an allocation unit size, if you choose to use a size other than the default. Default settings are strongly recommended for general use.

11. If desired, select a formatting option:
    - Quick format — Use this method to skip scanning for bad sectors in the partition during formatting.
    - Enable file and folder compression — Use this method to compress the folders and files that are added to the volume. This option is available only if the partition is being formatted with NTFS.

12. Click **Next** to continue, then **Finish**.

Create Logical Drives in an Extended Partition

An extended partition can be subdivided into as many as 32 logical drives. You can use all or part of the free space in an extended partition when creating logical drives. An extended partition or a logical drive cannot span multiple disks.

1. Select an extended partition from the left panel tree view.

2. Click **Create Logical Drive** to display the Create Volume wizard. Click **Next** to move through the wizard's Select Partition or Volume window. (**Partition** is selected by default.)

Select Partition Type

3. At the Select Partition Type window, the **Logical Drive** button is automatically selected, and displays the largest logical drive size that can be created on that drive.

4. Type the size of the logical drive that you want to create, and then click **OK**.

Select Disks and Region

5. Select the disk and free space for the partition.

Assign Drive Letter

6. Choose a drive letter if you want to name and format the logical drive. If you do not choose a drive letter, the logical drive cannot be formatted. (A drive letter can be assigned later.)

7. Click **Next** to continue, then **Finish**.

   The Status column in list view displays **Formatting** and displays the percentage of completion of the format process. The logical drive status changes to **Healthy** when formatting is complete, and the logical drive becomes available for use.

Dealing with Removable Media

In the graphical view, removable-media devices are labeled **Removable**. They are considered basic disks.

You can create only one primary partition on removable-media devices. You cannot create extended partitions, logical drives, or dynamic volumes on removable-media devices. You cannot upgrade removable-media devices to dynamic volumes. Removable-media devices can contain only one primary partition and can be formatted with FAT or NTFS. You cannot mark a primary partition on a removable-media device as active.

Upgrading a Basic Disk to Dynamic

On Windows Server 2003 you need to use the native Disk Management utility (not Array Manager) to upgrade a basic disk to dynamic. For information on accessing the native Disk Management on Windows Server 2003, see "Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux."

On Windows 2000, Array Manager can upgrade a basic disk that contains partitions, logical drives, or Disk Administrator volumes. Array Manager converts the partitions and logical drives to simple dynamic volumes and the Disk Administrator volumes to dynamic volumes of the appropriate corresponding RAID type. This capability is called encapsulation and is described in more detail in "Upgrade Disks with Legacy Volumes to Dynamic on Windows 2000" in the next chapter, "Volumes."

When Array Manager creates a partition on a basic disk, Array Manager checks whether the partition is the last one on the disk, and then reserves 1MB at the end of the disk for future upgrade needs. When you upgrade from basic to dynamic, the partition needs to align with the cylinder boundary, which may result in a cylinder size (usually 7.8MB) that is left on the disk. This cylinder size is displayed as unallocated disk space.

Warning! After you upgrade a basic disk to a dynamic disk, you will not be able to revert back to a basic disk until you first delete all the dynamic volumes on the disk. Then you can use the Revert to Basic Disk command to revert the disk back to a basic disk. Make sure to back up your data in this scenario before deleting any volumes, or your data will be lost!

Note There are particular considerations regarding dynamic disks and volumes on NetWare, Windows Server 2003, and Linux. See "Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux" for more information.

To Upgrade a Basic Disk to a Dynamic Disk

1. Right-click the disk you wish to upgrade and select **Upgrade Dynamic Disk** from the context menu that appears.

2. The Upgrade Disk wizard explains information about upgrading. Click **Next** to continue.

3. You are asked to select disks to upgrade. Add the basic disks you want to upgrade to the list of dynamic disks and click **Next** to continue, then **Finish**.

Working with Dynamic Disks

This section's topics are:

- Dynamic Disk Advantages
- Dynamic Disk Limitations
- Creating Volumes on a Dynamic Disk
- Extending a Simple or Spanned Volume on a Dynamic Disk
- Reactivating Dynamic Disks
- Merging a Foreign Disk
- Moving a Dynamic Disk to Another Computer
- "Reverting a Dynamic Disk to Basic"

Dynamic Disk Advantages

When you upgrade a dynamic disk, Array Manager writes a small database onto the upgraded disk to store volume management information about the disk and all other dynamic disks present on your system. This database is what gives a dynamic disk many advantages:

- Users can create dynamic volumes without restarting the system.
- Volumes can be accessed continuously, even across disk reconfiguration events.
- Simple volumes and spanned volumes can be readily extended without the need to find contiguous space on the available disks. These volumes can also be extended in size without rebooting the server. This is known as Online Volume Expansion (OLVE).

- Disk configuration information is contained on the disk, not in the registry. The disk configuration information is also replicated across all dynamic disks so that the failure of one disk will not obstruct access to data on other disks.
- Dynamic disks can be mirrored, spanned, or striped across disks.
- Dynamic disks provide volume migration, which is the ability to move a disk from one Windows 2000 system to another system without loss of data.

There are particular considerations regarding dynamic disks and volumes on NetWare, Windows Server 2003, and Linux. See "Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux" for more information.

Dynamic Disk Limitations

Although Windows 2000 can access dynamic disks created with Array Manager, the dynamic volumes on these dynamic disks cannot be changed without Array Manager.

Virtual disks of Fibre Channel storage in a Storage Area Network (SAN) can be converted to dynamic disks, but this should be done with the following considerations:

- Dynamic disks of SAN storage must not be moved by way of Dell OpenManage Storage Consolidation from server to server. Once a virtual disk of SAN storage is made dynamic, it must remain assigned to the server.
- It is better to upgrade virtual disks of SAN storage to dynamic only on an as-needed basis. This will make it easier to choose dynamic disks when creating dynamic volumes.

There are particular considerations regarding dynamic disks and volumes on NetWare, Windows Server 2003, and Linux. See "Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux" for more information.

Creating Volumes on a Dynamic Disk

You can create dynamic volumes on dynamic disks only. To create and delete simple, spanned, striped, mirrored, and RAID-5 volumes, see "Creating a Dynamic Volume" in the "Volumes" chapter.

There are particular considerations regarding dynamic disks and volumes on NetWare, Windows Server 2003, and Linux. See "Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux" for more information.

Extending a Simple or Spanned Volume on a Dynamic Disk

You can extend simple and spanned volumes without having to locate unallocated contiguous space on the disk. See "Extending a Dynamic Simple or Spanned Volume" in the "Volumes" chapter.

There are particular considerations regarding dynamic disks and volumes on NetWare, Windows Server 2003, and Linux. See "Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux" for more information.

Reactivating Dynamic Disks

A dynamic disk may appear as a missing disk when it is corrupted, powered down, or disconnected. You can reactivate a dynamic disk to bring it back online.

There are particular considerations regarding dynamic disks and volumes on NetWare, Windows Server 2003, and Linux. See "Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux" for more information.

To Reactivate a Missing or Offline Disk

1. Right-click the disk marked **Missing** or **Offline dynamic disk.**
2. Select **Reactivate Disk** from the context menu. The disk should be marked **Online** after the disk is reactivated.

Merging a Foreign Disk

Dynamic disks with a **Foreign** status are disks that have been moved from another computer. You cannot reactivate a foreign disk. To change a foreign disk's status and enable it to be seen as a part of the current computer system, you will need to use the command **Merge Foreign Disk**. For details on how this command is done, see the next topic "Moving a Dynamic Disk to Another Computer."

There are particular considerations regarding dynamic disks and volumes on NetWare, Windows Server 2003, and Linux. See "Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux" for more information.

Moving a Dynamic Disk to Another Computer

Before you disconnect the disks, make sure the status of the volumes on the disks is Healthy. If the status is not Healthy, you should repair the volumes before you move the disks. You need to prepare disks for removal in some cases. See "Prepare to Remove." It is also recommended that you back up all data on your volumes before you move disks to another computer.

Note The ability to move dynamic disks to another computer is not supported with SAN storage. Only basic disks can be moved between servers on a SAN. Refer to the Dell OpenManage™ Storage Consolidation software documentation for more information. In addition, dynamic disks and dynamic volumes are only supported on the Windows 2000 operating system.

To Move a Dynamic Disk to Another Computer

1. Turn the computer off, remove the physical disks, and then install the physical disks on the other computer. Restart the computer.
2. Open Array Manager.
3. In most cases, the disks are merged automatically. However, if the disks appear in the disk list and graphical views as **Foreign**, right-click a new disk, click **Merge Foreign Disks**, and then follow the instructions on your screen.

Notes

- After you merge a dynamic disk from another computer, you can see and use any existing volumes on that disk.
- Every time you remove disks from or merge disks to a computer, you must click **Rescan Disks** and then verify that the disk information is correct.
- You can move multiple disks from different machines to another computer by installing the disks, opening Array Manager, right-clicking any of the new disks, and then clicking **Merge Foreign Disks**. The disks are grouped according to the computer name they were moved from. You can choose which disks from the group you want to add by clicking **Select Disk**. However, if you do not have any dynamic disks installed, then all of the disks are added regardless of the disks you selected.
- After you have merged the disk or disks, another dialog box appears that lists the conditions of the volumes. Review this information carefully.

Reverting a Dynamic Disk to Basic

You can revert a dynamic disk to basic if there are no volumes on the disk. If volumes exist, you must delete them before you can revert the disk. Disks shown in Array Manager to be unreadable or to be **Foreign** disks can be reverted, but doing so causes you to lose all existing data on the reverted disk.

There are particular considerations regarding dynamic disks and volumes on NetWare, Windows Server 2003, and Linux. See "Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux" for more information.

General Disk Functions

This section on general disk functions has the following topics:

- Add a New Disk to Your System
- Replacing a Disk Receiving SMART Alerts
- Update Disk Information Using Rescan
- Refresh the Array Manager Console Display
- Write a Disk Signature
- Monitoring Disk Reliability

Note For disk troubleshooting information, see "Disk Troubleshooting Procedures."

Add a New Disk to Your System

Use this procedure to add a new disk to your system and make the disk visible to Array Manager.

Note If the disk you are adding already contains data, you must format the disk using the controller BIOS before the disk can be used by Array Manager. Until you format the disk using the controller BIOS, the disk will have an Offline or Foreign status and may generate other errors.

1. Install or attach the new physical disk (or disks).
2. Click **View**, and then click **Rescan**, or select **Rescan** from the toolbar.

   Array Manager should detect the new disk after you click **Rescan**. However, if Array Manager does not detect the new disk, you may need to restart your computer.

New disks are added to the computer as basic disks. On Windows 2000, you can update basic disks that contain data to dynamic. See the section "Upgrade Disks with Legacy Volumes to Dynamic on Windows 2000" in the "Volumes" chapter.

Related Information

- See "Replacing a Failed Disk" if you are replacing a disk that is part of a virtual disk.
- See "Virtual Disk Considerations for PERC Subsystem 1 Controllers" or "Virtual Disk Considerations for PERC Subsystem 2 Controllers" if you are intending to include the new disk in a virtual disk.

Replacing a Disk Receiving SMART Alerts

SMART alerts are event messages predicting that a disk may fail in the near future. If an array disk is receiving SMART alerts, you should replace the disk. Use the following procedures to replace a disk receiving SMART alerts.

If the disk is part of a redundant virtual disk:

1. Right-click the disk that is receiving SMART alerts and select **Offline** from the pop-up menu in the Array Manager tree view. See "Offline Array Disk" for more information on the Offline command.
2. Manually remove the disk.
3. Insert a new disk. Make sure that the new disk is the same size or larger as the disk you are replacing. (On some controllers, you may not be able to use the additional disk space if you insert a larger disk. See "Virtual Disk Considerations for PERC Subsystem 1 Controllers" for more information.) After you complete this procedure, a rebuild is automatically initiated because the virtual disk is redundant.

If the disk is not part of a redundant virtual disk:

1. Back up data from the virtual disk.
2. Delete the virtual disk.
3. Replace the disk that is receiving SMART alerts.
4. Create a new virtual disk. Make sure that the new virtual disk is the same size or larger than the original virtual disk. For controller-specific information on creating virtual disks, see "Virtual Disk Considerations for PERC Subsystem 1 Controllers" and "Virtual Disk Considerations for PERC Subsystem 2 Controllers."
5. Restore the backed up data from the original virtual disk onto the newly created virtual disk.

Update Disk Information Using Rescan

When Array Manager rescans disks, it scans all attached disks for disk configuration changes. It also updates information on removable media, CD-ROM drives, basic volumes, file systems, and drive letters.

Every time you remove or add disks to a computer, you must click **Rescan** and then verify that the disk information is correct.

To Rescan Disks

From the **View** pull-down menu, select **Rescan**, or select the **Rescan** icon from the toolbar.

A rescan can take several minutes, depending on the number of devices installed.

Refresh the Array Manager Console Display

Use this option to refresh drive letter, file system, volume, and removable-media information. **Refresh** also checks to see whether unreadable volumes are now readable. However, it does not scan hardware. To update hardware information, use **Rescan** instead.

1. From the **View** pull-down menu, select **Refresh**, or select **Refresh** from the toolbar. The General tab view displays a percentage status line under the **Progress** column in the right pane.

2. After the operation is finished, the Array Manager console displays current information.

Write a Disk Signature

Windows requires that a disk have a signature on it before it can be used. Most controllers require that the write signature be added manually.

When you create a virtual disk and then do a rescan, the newly created virtual disk appears under the Disks storage object. If the disk needs a signature to be written manually, its icon will display an error symbol. Also, its type will be listed as "No Disk Signature" in the disks displayed in the right pane of the console. If left unsigned, the disk cannot be used.

Use the following steps to write a disk signature. This procedure assumes that you have already created the virtual disk.

1. Expand the **Disks** object in the Array Manager tree view.

2. Right-click a virtual disk that needs a write signature to display the **Write Disk Signature** dialog. This dialog displays all virtual disks that require a disk signature.

3. Select the virtual disks for which you wish to do a write signature. You can do a write signature on more than one virtual disk at the same time.

4. Click **Ok**.

Once a signature is written on a disk, the disk is displayed as a basic disk. You can create partitions on the basic disk, or you can upgrade the disk to dynamic to create volumes on it.

Monitoring Disk Reliability

Array Manager supports Self Monitoring Analysis and Reporting Technology (S.M.A.R.T.) on array disks that are S.M.A.R.T. enabled.

S.M.A.R.T. performs predictive failure analysis on each disk and sends alerts if a disk failure is predicted. The RAID controllers check array disks for failure predictions and, if found, pass this information to the Array Manager console. Array Manager immediately displays an alert icon on the disk. Array Manager also raises an alert under the Events tab and in the Windows Event Log.

Note When a controller's I/O is paused, you will not receive S.M.A.R.T. events.

Array and Hard Disk Menu Commands

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, 4/IM, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, CERC SATA1.5/2s, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers.

This section describes the following commands for working with array disks:

- Rebuild
- Assign Global Hot Spare
- Unassign Global Hot Spare
- Online Array Disk
- Offline Array Disk
- Prepare to Remove
- Blink
- Unblink
- Array and Hard Disk Properties

Initialize Array Disks

This section applies to PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers.

The **Initialize** command prepares an array disk for use as a member of a virtual disk. For details, see the "Initialize Array Disks."

Caution This is a data-destructive command. Any virtual disks residing on this array disk will be removed.

Rebuild

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers.

The **Rebuild** function is enabled only when a disk has failed. You can rebuild failed disks only in redundant arrays.

To Rebuild a Failed Disk

1. Right-click the failed disk that you want to rebuild.
2. Click **Rebuild**. In the right pane, the status of the disk is **Rebuilding** and a progress bar shows the percentage of completion.

Note For information on performing a rebuild on the CERC SATA1.5/2s controller, see Rebuild (CERC SATA1.5/2s Controller).

Assign Global Hot Spare

This section applies to 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, and CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

A global hot spare is an unused backup disk that is part of the array group. Hot spares remain in standby mode. When an array disk that is used in a virtual disk fails, the assigned hot spare is activated to replace the failed array disk without interrupting the system or requiring your intervention. When a hot spare is activated, it rebuilds the data for all redundant virtual disks that were using the failed array disk.

You can change the hot spare disk assignment by unassigning a disk and choosing another disk to assign, as needed.

For more information on hot spares, see "Understanding Hot Spares."

Note Global hot spares must be assigned and unassigned manually. They are not assigned to specific virtual disks. On PERC Subsystem 1 controllers, there are size requirements for the hot spare. See "Considerations for Hot Spares (PERC Subsystem 1)" for more information. If you want to assign a hot spare to a virtual disk (it will replace any array disk that fails in the virtual disk) then use the "Configure Dedicated Hot Spare" command. For more information on hot spares, see "Understanding Hot Spares."

To Assign a Global Hot Spare

1. Right-click the array disk that you wish to use as a hot spare.
2. Click **Assign Global Hot Spare**.
3. To confirm the successful completion of the operation, check the properties displayed in the right pane. The status of the array disk must be **Ready** and type must be **Spare Array Disk**.

Unassign Global Hot Spare

This section applies to 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, and CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers

Use this operation to unassign the global hot spare disk.

To Unassign a Global Hot Spare

1. Right-click the disk that is assigned as a hot spare.
2. Click **Unassign Global Hot Spare**.
3. To confirm the successful completion of the operation, check the properties displayed in the right pane. The status of the array disk must be **Ready** and type must be **Array Disk**.

Online Array Disk

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers.

You can set to **Online** only an array disk that is part of a redundant virtual disk that has been set to **Offline**.

To Online an Array Disk

1. Right-click the disk that you want to set to **Online**.
2. Click **Online**. In the right pane, the status of the disk is changed from **Offline** to **Online**.

Offline Array Disk

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers.

You can set to **Offline** only an array disk that is part of a redundant virtual disk that has been set to **Online**.

To Offline an Array Disk

1. Right-click the disk that you want to set to **Offline**.
2. Click **Offline**. A confirmation dialog box appears.
3. Click **OK** to continue. In the right pane, the status of the disk will be changed from **Online** to **Failed**.

Prepare to Remove

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers.

Before you remove any physical disk from an enclosure, it is recommended that you perform this operation to prevent data loss.

Note This procedure is not supported on the CERC ATA100/4ch, CERC SATA1.5/6ch, or CERC SATA1.5/2s controllers. In addition, this procedure is only supported for hotswap array disks (disks that reside in a carrier).

To Prepare to Remove an Array Disk

1. Right-click the disk that you want to remove.
2. Click **Prepare to Remove,** then click **Yes** to continue.
3. You can remove the disk under either of the following conditions:
   - Wait for about 30 seconds to allow the disk to spin down.
   - Wait until you notice the initial blink pattern has changed into a different pattern.

Note On PERC 2, 2/Si, 3/Si, and 3/Di controllers, this procedure is only supported for hotswap array disks (disks that reside in a carrier). In addition, you may need to perform a rescan for the new disk to be displayed in the Array Manager tree view.

Blink

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers.

The **Blink** command allows you to find an array disk within an enclosure by blinking one of the light-emitting diodes (LEDs) on the disk. The LEDs could blink from 30 seconds to indefinitely. You may want to use this command to locate an array disk that has failed.

To blink an array or hard disk:

1. Right-click the disk object in the tree view.
2. Select **Blink** from the pop-up menu.

Note This procedure is not supported on the CERC ATA100/4ch, CERC SATA1.5/6ch, or CERC SATA1.5/2s controllers. In addition, this procedure is only supported for hotswap array disks (disks that reside in a carrier).

Note When using an Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, or LSI 1030 controller, the **Blink** and Unblink commands apply to hard disks contained in carriers that can be inserted into a server or an enclosure. If the hard disk is not contained in a carrier but is instead designed to be connected with a SCSI cable (typically a ribbon cable), then the **Blink** and **Unblink** commands are disabled.

Unblink

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers.

The **Unblink** command allows you to cancel the **Blink** command.

To unblink an array or hard disk:

1. Right-click the disk object in the tree view.
2. Select **Unblink** from the pop-up menu.

Note This procedure is not supported on the CERC ATA100/4ch, CERC SATA1.5/6ch, or CERC SATA1.5/2s controllers. In addition, this procedure is only supported for hotswap array disks (disks that reside in a carrier).

Note When using an Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, or LSI 1030 controller, the Blink and **Unblink** commands apply to hard disks contained in carriers that can be inserted into a server or an enclosure. If the hard disk is not contained in a carrier but is instead designed to be connected with a SCSI cable (typically a ribbon cable), then the **Blink** and **Unblink** commands are disabled.

Remove Orphan/Dead Disk Segments

This section applies to PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers.

This command removes the disk space occupied by an orphan or dead disk segment.

A segment may be designated as dead if it is logically bad, is physically damaged, or resides on a disk that has been disconnected from and then reconnected to the controller.

Array and Hard Disk Properties

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, 4/IM, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, CERC SATA1.5/2s, Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers.

Use this command to display the array or hard disk properties. Array disks are attached to RAID controllers (PERC Subsystem 1 and 2) and hard disks are attached to non-RAID SCSI controllers (SCSI Subsystem). The PERC 3/Di and PERC 4/Di controllers, however, can run in RAID/SCSI mode; these controllers handle one channel in RAID mode and the other channel in non-RAID SCSI mode.

To display the disk properties:

1. Right-click the disk object in the Array Manager tree view.
2. Select **Properties** from the pop-up menu. This action displays the **Array Disk Properties** dialog which displays the disk properties.

Disk properties may include:

- **Disk**— This field displays the name of the disk object in the Array Manager tree view.
- **Type** — This field indicates the type of storage component. For disk properties, the value in this field is “Array Disk.”
- **Status** — This field displays the current status of the disk. Possible values are:
  - **Ready** — The disk is functioning normally.
  - **Online** — The array disk is part of a virtual disk and is functioning normally.
  - **Degraded** — The disk has suffered a failure of a component and is operating in a degraded state.
  - **Failed** — The disk has suffered a failure of one or more components and is no longer functioning. In certain circumstances, a failed array disk will not display a Failed state. See Array Disk State, Event Messages and Hot Spares on PERC Subsystem 1 Controllers for more information.
  - **Rebuilding** — Data from a redundant virtual disk is currently being rebuilt onto the array disk.
  - **SMART: Failure predicted, Ready** — A SMART alert has predicted that this disk may fail in the near future. It is strongly recommended that you replace the disk immediately.

- **Removed** — Indicates that the controller cannot communicate with the disk. This state may occur when the disk was part of a virtual disk but has lost communication with the controller. A user may have removed the disk or it may have lost power. A disk that is in Removed state will no longer be displayed by Array Manager after a system reboot or stopping and restarting the Array Manager service.

- **Capacity**— This field displays the total capacity of the disk.
- **Unallocated Space** — This field displays the available space on the disk.

Note If you have PERC Subsystem 1 and PERC Subsystem 2 RAID controllers, the **Unallocated Space** field displays the available space. If you have SCSI (non-RAID) controllers, the field displays "N/A" (Not Applicable). If a controller is running in RAID/SCSI mode, the disks attached to the SCSI channel are displayed as "Hard Disk 0:1" instead of "Array Disk 0:1". Because the disks attached to the SCSI channel are not used to build virtual disks, they will also display "N/A" in the **Unallocated Space** field, just like the disks on non-RAID controllers.

- **Device Type** — This field displays the SCSI devices or system ports to which the disk is attached. See "Display of IDE Devices" for more information on how Array Manager displays IDE devices.
- **Target ID** — This field displays the SCSI ID of the disk.
- **Hardware Vendor** — This field displays the disk's hardware vendor.
- **Product ID** — This field displays the disk's product ID.
- **Revision** — This field displays the disk's revision number.
- **Serial Number** — This field displays the disk's serial number.
- **Negotiated Speed**— This field displays the speed of data transfer that the drive negotiated while spinning up and upon initial communication with the controller. This speed is dependent on the speed of the drive, the capable speed of the controller, the current speed of the controller on that channel, and the speed of the EMM (Enclosure Management Module) on the enclosure.
- **Capable Speed**— This field displays the highest possible speed that the device can transfer data.

Note Negotiated and capable speeds are not displayed for PERC 2/SC, PERC 2/DC, CERC ATA100/4ch, CERC SATA1.5/6ch, CERC SATA1.5/2s, Ultra SCSI, or Ultra2 SCSI. New SCSI drives manufactured since the introduction of the Ultra 320 SCSI drives will display this negotiated and capable speed information from Array Manager if your controller supports reporting them. If the drive was manufactured before that time, the negotiated and capable speed fields might display "not available".

Note Array disks for the PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers display Ready status if they are not part of a virtual disk and Online status if they are part of a virtual disk.

Note Array disks for the PERC 2, PERC 2/Si, PERC 3/Si, PERC 3/Di, CERC SATA1.5/6ch**,** and CERC SATA1.5/2s controllers display Ready status regardless of whether they are a part of a virtual disk or not. Ready status for these controllers means the disk is operational.

Volumes 15

This chapter describes conceptual and procedural information about how Array Manager implements basic and dynamic volumes.

Note There are particular considerations regarding dynamic disks and volumes on NetWare, Windows Server 2003, and Linux. See "Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux" for more information.

Volume Overview

A volume is a logical entity that is made up of portions of one or more physical disks. A volume can be formatted with a file system and can be accessed by a drive letter. The maximum size of a volume depends on the quantity of free disk space and the type of volume selected.

Volumes can be:

- Simple Volumes
- Primary and Extended Partitions
- Logical Drives
- Basic Volumes
- Dynamic Volumes

Basic Volumes

Basic volumes refer to all the volumes that are on basic disks. Basic volumes can be primary or extended partitions or simple logical drives that reside on extended partitions.

Dynamic Volumes

A dynamic volume is a logical volume that can be created from one or more dynamic disks using Array Manager software. Dynamic volume types include simple, spanned, striped, mirrored, and RAID-5. By using spanned dynamic volumes, you can dynamically increase the size of your dynamic volumes as the need arises. For more information, see "Working with Dynamic Volumes" in this section.

There are particular considerations regarding dynamic disks and volumes on NetWare, Windows Server 2003, and Linux. See "Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux" for more information.

Simple Volumes

In Array Manager, a simple volume is a volume that occupies contiguous space on a single disk. You can extend a simple volume within the same disk or onto additional disks. If you extend a simple volume across multiple disks or across noncontiguous areas on the same disk, it becomes a spanned volume. You can create simple volumes only on dynamic disks. Simple volumes are not fault tolerant but can be mirrored.

When a basic disk with a partition is upgraded, the partition becomes a simple volume. An extended partition on a basic disk also becomes a simple volume when the disk is upgraded to dynamic.

Primary and Extended Partitions

In Array Manager, primary and extended partitions and logical drives exist only on basic disks. Basic disks use the traditional disk partitioning mechanism used by MS-DOS or a Windows operating system. A basic disk can have up to either four primary partitions or three primary partitions plus an extended partition. This extended partition can be subdivided into as many as 32 logical drives.

Logical Drives

A logical drive is a simple volume that resides on an extended partition of a basic disk. You can use all or part of the free space in an extended partition when creating logical drives.

General Volume Functions

Just as basic and dynamic disks share some common disk procedures, basic and dynamic volumes also share common volume procedures. Those procedures are documented in this section. The section's topics include:

- Add, Change, or Remove a Drive Letter or Path
- Mount a Volume at an Empty Folder
- Unmount a Volume
- View All Drive Paths
- Check Partition or Volume Properties
- Format a Partition or Volume
- Refresh Drive Letter, File System, and Partition or Volume Information
- Delete a Partition or Volume
- Expanding a Volume
- Array and Volume Migration

Add, Change, or Remove a Drive Letter or Path

A drive path refers to the path to an NTFS volume that is mounted at an empty folder. Mounting a volume at an empty folder is a feature of the Windows operating system. On Windows 2000, the Volume context menu contains the command **Change Drive Letter and Path**.

The **Change Drive Letter and Path** or **Change Drive Letter** function is not available when formatting or upgrading to dynamic any disks contained in the volume. The format or upgrade to dynamic operation must complete before you can change the drive letter or path.

Note that you cannot reassign the drive letter of your active volume, the volume that contains your boot disk. You also cannot mount a volume on your boot disk.

Caution Changing the drive letter or drive path on a partition or volume may cause programs installed on that partition or volume to no longer run.

Note Array Manager does not support volume management on Windows Server 2003. To implement volume management on Windows Server 2003, you need to use the native Disk Management Utility. For more information on using the native Disk Management utility for Windows Server 2003, see "Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux."

To Add, Change, or Remove a Drive Letter or Path on a Windows 2000 Partition or Volume

1. Right-click the designated volume and select **Change Drive Letter and Path** from the context menu that comes up.

   The Drive Letter and Paths screen appears.

2. Click the **Add**, **Remove**, or **Modify** button, depending on what you want to do.

   If a partition or volume does not have a drive letter, click **Add**. If you want to change an existing drive letter, click **Modify**. If you want to remove a drive letter or drive path, click **Remove**.

3. Follow these steps if you clicked the **Add** button:

   The Add Drive Letter and Path screen appears.

   - Select a drive letter from the pull-down list and click **OK**.

     If a drive letter is already assigned, you will not be able to add another drive letter. The option will be inactive (dimmed).

   Or

   - Click the button for mounting the volume at an empty folder, then browse to select the empty folder, and click **OK**.

     You can add multiple drive paths to a partition or volume.

   For details on mounting a volume at an empty folder, see "Mount a Volume at an Empty Folder" in the next section of this chapter.

4. If you clicked the **Remove** button, click **Yes** in the confirmation screen that appears.

5. Follow these steps if you clicked the **Modify** button:

   The Modify Drive and Path screen appears.

   - Select a drive letter from the pull-down list and click **OK**.

     If you try to assign a drive letter to a volume that is in use, a dialog box will ask you whether you want to forcibly reassign the drive letter. You can click **Yes** to do so.

   Or

   - Click the button for mounting the volume at an empty folder, then browse to select the empty folder, and click **OK**.

If you want to modify an existing drive path, you will not be able to do so in this screen. You have to remove the drive path and recreate it. Click **Cancel** to return to the first screen and click **Remove**. Then you have to invoke the command again to recreate the drive path with the **Add** button.

You can add multiple drive paths to a partition or volume.

For details on mounting a volume at an empty folder, see "Mount a Volume at an Empty Folder" in the next section of this chapter.

Mount a Volume at an Empty Folder

Array Manager allows you to mount a volume at any empty folder on a local NTFS volume if you are using Windows 2000.

Array Manager assigns a drive path to the volume rather than a drive letter. Drive paths are useful because they eliminate the 24-drive-letter limit on hard-disk volume names. (Drive letters A and B are reserved for floppy drives.) The volume can be a partition, a logical drive that was created in Disk Administrator, or a dynamic volume.

The only restrictions are that the folder you mount the volume to must be empty and it must be on an NTFS volume on a local computer. You cannot mount volumes to network-accessible drives. You can, however, assign multiple drive paths to a single volume.

Mounted drive folders are distinguished from ordinary folders by having a drive icon instead of the normal folder icon.

You can assign the drive path when you create the volume with the Create Volume wizard at the Assign Drive Letter screen. You can also use the command **Change Drive Letter and Path** after the volume is created. A volume that already has a drive letter assigned can also have one or more drive paths. It does not have to be an unlabeled volume.

To Mount a Volume at an Empty Folder

1. Create an empty folder on a drive with free space.

   If you have not already created the volume, use the Create Volume wizard to create the volume. Do the following:

   - At the Assign Drive Letter screen, click **Mount at an empty NTFS folder**, type the path to the empty folder, or click **Browse** to locate it.
   - Complete the remaining steps for creating the volume.

2. If the volume was already created, use the command **Change Drive Letter and Path**. Its screen is similar to the Assign Drive Letter screen in the Create Volume wizard:

- Highlight the volume you want to change to a mounted volume.
- Right-click to bring up a context menu and select the command Change Drive Letter and Path.
- Click the **Add** button.
- Click **Mount this volume at an empty folder,** which supports drive paths, type the path to the empty folder, or browse to locate it.

  If you have not already created the folder, you can click the **New Folder** button to have Array Manager create the folder for you.
- Click **OK** to complete the command.

Unmount a Volume

To unmount a volume with a drive path:

1. Right-click the volume in the tree view and select **Change Drive Letter and Paths** from the pop-up menu.

2. Verify that the correct volume is highlighted in the **Drive Letter and Paths** dialog box and click **Remove**.

Note To modify a drive path, you must first remove it, and then create a new drive path using the new location. You cannot modify the drive path directly.

The ability to use Array Manager to mount and unmount a volume at an empty folder is only supported in Windows 2000. In Windows Server 2003, you can perform this function using the native Disk Management utility. For information on accessing the native Windows Server 2003 Disk Management utility, see "Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux."

View All Drive Paths

This command lets you view all the mounted volumes on a selected server. These volumes have been assigned a drive path rather than a drive letter. The capability of mounting a volume at a folder is only supported in Windows 2000 when the native disk management has been replaced by Array Manager.

To View All Mounted Volumes

1. Right-click the computer icon and select **View All Drive Paths**.

   The View All Drive Paths window appears.

2. If you want to unmount any volume listed in this window, select the drive path and click **Remove**.

   You can also use the **Unmount** command or the **Change Drive Letter and Path** command to remove a drive path.

3. Click **Close**.

See also "Mount a Volume at an Empty Folder" in this section.

Check Partition or Volume Properties

To Check the Properties of a Partition or Volume

1. Right-click the partition or volume to be checked.

2. Choose **Properties** from the context menu that appears.

   The Properties window comes up.

Format a Partition or Volume

Caution All data in the volume or partition will be lost when you format it.

To Format a Partition or Volume

1. Right-click the volume or partition you want to format and then click **Format**.

2. Choose either FAT, NTFS, or FAT32 as the file system to use.

   Array Manager allows you to format using either FAT or NTFS. In addition, for Windows 2000 only, you can use FAT32.

- **FAT (File Allocation Table)** — A file system used by MS-DOS® and all versions of Windows. The operating system maintains a table to keep track of the status of various segments of disk space used for file storage.

- **NTFS (NT File System)** — An advanced file system designed for use specifically within the Microsoft Windows NT and Windows 2000 operating systems. Windows Server 2003 can also use the NTFS file system. Use this format if you want to use file and folder compression.
- **FAT32** — (Windows 2000 and Windows Server 2003) A derivative of the file allocation table (FAT) file system. FAT32 supports smaller cluster sizes than FAT, thus providing more efficient space allocation on FAT32 drives. FAT32 is designed for larger disks than FAT.

3. Enter a label for the volume. This label appears on the Array Manager console. If a name has been selected, this name appears in the text box, but you can change it here.

4. Leave the allocation unit size as Default.

Note It is strongly recommended that you leave the allocation unit size as **Default**. Selecting any other value requires an advanced knowledge of Windows file systems. You should thoroughly review the Microsoft Windows documentation before selecting any value other than **Default**.

5. Select file system type and formatting options:
   - Quick format — Formats the volume or partition without scanning for bad sectors in the volume or partition. Check this box to use this format method. Default is unchecked.
   - Enable file and folder compression —- Can be used only if you choose NTFS format. Check this box to use this format method. Default is unchecked.

6. Select **OK** to begin formatting. A progress bar will display in the list view.

Note You can cancel formatting at any time during the process by selecting **Cancel format** from the menu. However, the **Cancel** option does not attempt to restore any of the volume's content.

Refresh Drive Letter, File System, and Partition or Volume Information

To refresh, click **View** and then **Refresh**. Progress bars display the percentage of progress while this command runs.

When you click the **Refresh** command, Array Manager updates displayed data with the current information. **Refresh** also checks to see whether unreadable volumes are now readable. However, it does not scan hardware. To update hardware information, use **Rescan** instead.

Delete a Partition or Volume

Caution All data in the volume or partition will be lost when you delete it. Make sure that the information has been backed up onto another storage medium and verified, or that it is no longer needed.

To Delete a Partition or Volume

1. Right-click the designated volume and select **Delete Volume**.
2. You are prompted for verification. Click **Yes** to delete, or **No** to cancel. The volume is removed immediately if you click **Yes**.

Note You cannot delete the volume with the system files (the boot volume).

Expanding a Volume

There are different methods for expanding a volume and the virtual disk on which the volume resides. Each method has particular advantages and disadvantages associated with it. These methods are described in the following sections:

- Expanding a Volume Using Any RAID Level by Replacing Disks
- Expanding a Volume Using Any RAID Level by Adding New Disks
- Expanding a Redundant RAID Volume

Expanding a Volume Using Any RAID Level by Replacing Disks

You can expand a volume by removing the disks on which the volume resides and replacing them with larger disks. This method can be used on any RAID level. Using this method causes the volume to be unavailable until you have completed the procedure.

This method has the following advantages and disadvantages.

Advantages:

- This method can be used on any RAID level whether the volume (or virtual disk) is redundant or not.

- You can expand the volume by replacing drives currently being used. This means that you do not need empty drive slots on your system for increased storage.

Disadvantages:

- The volume will be inaccessible for a long period of time while you replace disks and restore from backup.

To expand a volume regardless of RAID level by replacing disks:

1. Back up all data (including all of the volumes) that reside on the virtual disk that contains the volume you wish to expand. This step is critical.

2. Remove the disks containing the virtual disks on which the volume resides.

3. Insert new disks with greater storage capacity. After inserting the new disks, you may need to perform a rescan in order for the disks to be displayed in the Array Manager tree view. (On some controllers, there are considerations for maximizing disk space usage. See "Virtual Disk Considerations for PERC Subsystem 1 Controllers" for more information.)

4. Create a new virtual disk using the disks that you have just inserted. This virtual disk should be larger than the original one.

5. Create a volume on the virtual disk you have just created.

6. Restore the data from your backup onto the virtual disk.

Expanding a Volume Using Any RAID Level by Adding New Disks

If your system has slots available for inserting new disks, then you can expand a volume by adding new disks and extending the volume onto those disks.

This method has the following advantages and disadvantages.

Advantages:

- This method can be used on any RAID level whether the volume (or virtual disk) is redundant or not.
- This is the quickest method for expanding a volume.

Disadvantages:

- This method requires that your system have empty slots available for the insertion of additional disks.

To expand a volume regardless of RAID level by adding new disks:

1. Back up all data (including all of the volumes) that reside on the virtual disk that contains the volume you wish to expand. This step is highly recommended.

2. Insert additional disks in the free slots available on your system. After inserting the new disks, you may need to perform a rescan in order for the disks to be displayed in the Array Manager tree view.

3. Create a new virtual disk using the disks that you have just inserted. (On some controllers, there are considerations for maximizing disk space usage. See "Virtual Disk Considerations for PERC Subsystem 1 Controllers" for more information.)

4. Expand the **Disks** object in the Array Manager tree view and locate the disks corresponding to both the original virtual disk on which the volume resides and the new virtual disk that you just created. Right click each of these disks and select **Upgrade Dynamic Disk** from the pop-up menu. Use the Upgrade Disk wizard to complete the upgrade. See "Upgrading a Basic Disk to Dynamic" for more information.

5. Expand the **Volumes** object in the Array Manager tree view.

6. Right-click the volume you intend to expand and select **Extend Volume** from the pop-up menu. This action displays the Extend Volume Wizard.

7. Click **Next** to advance past the wizard's Welcome screen.

8. Select **Custom Mode** on the Select Extend Volume Method screen and click **Next**.

9. Select the new disks that you have added to the system. They should be displayed in the **Select disks to extend this volume** portion of the Specify Size and Disks to Extend Volume screen. After you have selected these disks, specify their combined size in the **Extend Volume by** text box.

10. Click **Next** and continue through the Extend Volume Wizard.

Expanding a Redundant RAID Volume

You can use this method if you have a redundant RAID level and you wish to extend the volume by replacing disks while the volume remains up and running.

This method has the following advantages and disadvantages.

Advantages:

- The volume remains accessible while being extended.
- You can expand the volume by replacing drives currently being used. This means that you do not need empty drive slots on your system for increased storage.

Disadvantages:

- This method can only be used with a redundant RAID level.
- It may take a long time to rebuild each redundant disk that you replace.
- You will be creating a new virtual disk when using this method. Some controllers have considerations for maximizing disk space usage when creating virtual disks. See "Virtual Disk Considerations for PERC Subsystem 1 Controllers" for more information.

To expand a redundant RAID volume:

1. Back up all data (including all of the volumes) that reside on the virtual disk that contains the volume you wish to expand. This step is highly recommended.
2. Remove one of the disks containing the virtual disk on which the volume resides and replace it with a larger disk. Because the volume is redundant, the system automatically starts a rebuild on the new disk. Allow the rebuild to complete.
3. Repeat Step 2 for each disk included in the virtual disk.
4. Create a new virtual disk using the additional space created by inserting the larger disks.
5. Expand the **Disks** object in the Array Manager tree view and locate the disks corresponding to both the original virtual disk on which the volume resides and the new virtual disk that you just created. Right click each of these disks and select **Upgrade Dynamic Disk** from the pop-up menu. Use the Upgrade Disk wizard to complete the upgrade. See "Upgrading a Basic Disk to Dynamic" for more information.
6. Expand the **Volumes** object in the Array Manager tree view.
7. Right-click the volume you intend to expand and select **Extend Volume** from the pop-up menu. This action displays the Extend Volume Wizard.
8. Click **Next** to advance past the wizard's Welcome screen.
9. Select **Custom Mode** on the Select Extend Volume Method screen and click **Next**.
10. Select the new disks that you have added to the system. They should be displayed in the **Select disks to extend this volume** portion of the Specify Size and Disks to Extend Volume screen. After you have selected these disks, specify their combined size in the **Extend Volume by** text box.
11. Click **Next** and continue through the Extend Volume Wizard.

Array and Volume Migration

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers

This section on array and volume migration issues covers the following topics:

- Required Conditions
- Migrating Virtual Disks and Volumes to Another System
- Saving the New Controller Configuration

Required Conditions

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers

After some planning and preparation, it is possible to migrate array disks, virtual disks and array volumes between one controller and another, between one enclosure and another, or between one server and another. You can perform array and volume migration on PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers as long as the following conditions are met:

1. PERC 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, virtual disks cannot be migrated to a PERC 2/SC or PERC 2/DC controller. PERC 2/SC and PERC 2/DC virtual disks can be migrated to a PERC 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, and 4e/Di controller.

2. When moving the array disks from one enclosure to another, the SCSI ID for each disk must remain the same.

3. When moving the arrays disks from one enclosure to another or when moving an external enclosure from one server to another, the enclosure must be connected to the same channel number on the controller as in the original condition.

   In the case where a virtual disk consists of array disks on multiple channels, each array disk must be migrated to an enclosure that is connected to the same channel number that the array disk or enclosure was originally connected to. This also prevents migration of disks on channel 1 of a PERC 2/DC controller to a PERC 2/SC controller because the PERC 2/SC has only the single channel 0.

4. A virtual disk cannot be migrated between the PERC 2, PERC 2/Si, PERC 3/Si, PERC 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controller family and the PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controller family.

5. A virtual disk cannot be migrated from one controller to another unless the receiving controller's configuration has been cleared.

6. Virtual disks from one controller cannot be merged with virtual disks from another controller.

Migrating Virtual Disks and Volumes to Another System

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers

To migrate virtual disks and volumes from one system to another, do the following:

1. Shut down the system that the array disks are being moved from.

2. If the receiving controller has a preexisting virtual disk configuration on attached array disks, use the following procedure for clearing the configuration:
   - Shut down the receiving server.
   - Remove all the array disks from the controller.
   - Start up the receiving server and clear the configuration from the controller BIOS. After making the changes, power down the server when it indicates to reboot the system.

3. If the receiving controller does not have a preexisting virtual disk configuration, then shut down its server.

4. Place the array disks into the new enclosure.

5. Start up the system connected to the receiving controller. When the system connected to the new enclosure comes up, use the BIOS access (Ctrl-M) to update the controller configuration with the information saved on the array disks. Exiting the BIOS access will require one more reboot of the system.

The migration is complete. The virtual disk is now manageable through Array Manager.

Saving the New Controller Configuration

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers

Use the following procedure to enter the controller BIOS configuration screens to save the new configuration:

1. Reboot the PowerEdge server.
2. Enter the PowerEdge RAID Controller BIOS by pressing **Ctrl+M** when prompted during the power on self test (POST) of the controller during the bootup sequence.
3. Select the **Configure** menu option and press Enter.
4. Select **View/Add Configuration** and press Enter.

   This will read the data on the array disks to restore the RAID controller's proprietary configuration information. Press Esc and save the new configuration when prompted.
5. Press the **Esc** key and save the new configuration when prompted.
6. Reboot the system.

Working with Basic Volumes

The functions that apply only to basic volumes or partitions include:

- Create Primary and Extended Partitions
- Create Logical Drives in an Extended Partition
- Mark a Primary Partition as Active
- Upgrade Disks with Legacy Volumes to Dynamic on Windows 2000

The first two functions are covered in the "Disks" chapter. The remainder are covered in this section.

Mark a Primary Partition as Active

You can change the partition your system uses to start up by marking a selected primary partition as active.

The names commonly used for partitions containing the startup and operating system files are the **system** and **boot partitions**, respectively. The system partition refers to the partition containing the startup files. The boot partition contains the operating system files.

The system partition must be a primary partition that has been marked as active for startup purposes and must be located on the disk that the computer accesses when starting up the system. There can be only one active system partition at a time, which displays as **Active** in the status field. If you want to use another operating system, you must first mark its system partition as active before restarting the computer.

Caution It is strongly recommended that you do not change the drive letter for the partition containing \WINNT on a Windows system. This is because the Windows registry contains multiple references to the drive letter of the partition that contains \WINNT. Changing this drive letter can cause severe problems.

To Mark a Partition as Active

1. Select the primary partition that contains the startup files for the operating system that you want to activate.
2. Right-click to display the **Partition** menu and click **Mark Partition Active**.
3. A message appears, advising you that the partition will be marked **Active** and that the operating system on that partition will be started when you restart your computer. Click **Yes** to proceed.
4. The status of the partition displays **System**.

The Dell Utility Partition

The utility partition is a factory-installed, bootable partition on the hard-disk drive that provides utilities and diagnostics for your Dell PowerEdge™ system. When activated, the partition boots and provides an executable environment for the partition's utilities. When the utility partition is not booted, it is designated as a non-MS-DOS partition.

This partition is placed on your hard-disk drive at the factory; it is recommended that you leave it as is.

Upgrade Disks with Legacy Volumes to Dynamic on Windows 2000

Under Windows 2000, you can upgrade partitions to dynamic volumes by upgrading the disk(s) they are on to dynamic. This feature is called encapsulation.

With Windows 2000, you can upgrade basic disks containing these volumes to dynamic disks, and Array Manager will convert the partitions into volumes of the appropriate type, preserving all data contained in them. This upgrade can be done on any basic disk, including those containing the system and/or boot partitions.

When you upgrade disks that contain partitions, these partitions become simple volumes on the dynamic disks, and the other volumes become dynamic volumes of the corresponding type:

- A basic spanned volume (volume set) becomes a simple or spanned dynamic volume.
- A basic mirrored volume (mirror set) becomes a mirrored volume.
- A basic striped volume (stripe set) becomes a striped dynamic volume.
- A basic RAID-5 volume (stripe set with parity) becomes a dynamic RAID-5 volume.

For detailed instructions for converting a basic disk to a dynamic disk, see the section "Upgrading a Basic Disk to Dynamic" in the previous chapter, "Disks."

There are particular considerations regarding dynamic disks and volumes on NetWare, Windows Server 2003, and Linux. See "Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux" for more information.

Working with Dynamic Volumes

Dynamic volumes are volumes created on dynamic disks with Array Manager. On a dynamic disk, storage is divided into volumes instead of partitions. A dynamic volume consists of a portion or portions of one or more physical disks of any possible dynamic volume type. The five dynamic volume types are: simple, spanned, mirrored, striped, and RAID-5.

Note You cannot create a dynamic mirrored or RAID-5 volume if you are using a Windows 2000 Professional machine. Also, dynamic volumes cannot be created on Microsoft Cluster Server (MSCS) shared resources (that is, cluster-enabled array disks that are attached to a PERC 3/DC controller).

You can create any number of dynamic volumes in the unallocated space on a disk or create volumes that span two or more disks. Each volume on a disk can have a different file system, such as the file allocation table (FAT or FAT32) file system or the Microsoft NTFS file system.

Remember that any changes you make to your disk are immediately available. You do not need to quit Array Manager to save them or reboot your computer to implement them.

Topics covered in this section include:

- Dynamic Volume Types
- Creating a Dynamic Volume
- Extending a Dynamic Simple or Spanned Volume
- More on Mirrored Volumes
- More on RAID-5 Volumes

There are particular considerations regarding dynamic disks and volumes on NetWare, Windows Server 2003, and Linux. See "Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux" for more information.

Dynamic Volume Types

This section covers the five dynamic volume types:

- Simple
- Spanned
- Striped
- Mirrored
- RAID-5

There are particular considerations regarding dynamic disks and volumes on NetWare, Windows Server 2003, and Linux. See "Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux" for more information.

Simple

A dynamic simple volume consists of a single contiguous region on a disk. Dynamic simple volumes can be extended or mirrored. You can extend a simple volume within the same disk or onto additional disks. If you extend a simple volume across noncontiguous areas on the same disk or across multiple disks, it becomes a spanned volume. Dynamic simple volumes by themselves are not fault tolerant, but you can add data redundancy to a simple volume by mirroring it to another disk.

Spanned

A dynamic spanned volume consists of two or more subdisks (contiguous disk regions) on one or more disks. With a spanned volume, you can combine sections of unallocated space from multiple dynamic disks into one large volume. The areas of unallocated space used to create spanned volumes can be different sizes. Spanned volumes are organized sequentially—that is, Array Manager sequentially allocates space on each disk until that disk is full and then continues with the next disk until the volume size is reached.

Spanned volumes can be extended. After a volume is extended, no portion of it can be deleted without deleting the entire spanned volume. Spanned volumes by themselves are not fault tolerant, but you can add data redundancy to a spanned volume by mirroring it to another disk.

When you create a dynamic volume, simple and spanned dynamic volumes are combined together in an option called **concatenated**.

Note Spanning volumes across multiple disk controllers is not recommended.

Striped

Striped volumes are created by combining areas of free space on two or more disks into one logical volume. Data is divided into blocks and spread in a fixed order among all the disks in the volume. With a striped volume, data is written to multiple disks, similar to spanned volumes. However, striping writes files across all disks so that data is added to all disks at the same rate.

Striped volumes offer the best performance of all the disk management strategies. However, striped volumes do not provide fault tolerance. If a disk in a striped volume fails, the data in the entire volume is lost. You cannot mirror or extend a striped volume.

Mirrored

A mirrored volume is a volume that duplicates your data on two physical disks. A mirror provides redundancy by simultaneously writing the same data onto two separate volumes that reside on different disks. If one of the disks fails, data continues to be written to and read from the unaffected disk. You can mirror only a simple or spanned dynamic volume. Those two volume types combined together are known as concatenated.

RAID-5

A RAID-5 volume is a fault-tolerant volume with data and parity striped alternately across three or more physical disks. If one of the physical disks fails, data that was on the failed disk is recreated from the remaining data and parity on the other disks. RAID-5 volumes are a good solution for data redundancy in a computer environment in which most activity consists of reading data.

You need a minimum of three disks for a RAID-5 volume. Array Manager can accommodate up to 32 disks in a RAID-5 volume.

For more information about RAID strategies, see "What Is RAID?" and "Choosing RAID Levels and Concatenation" in the "Understanding Storage Concepts" chapter.

Creating a Dynamic Volume

This section on creating a dynamic volume has some useful background information, as well as the detailed steps for successfully carrying out the procedure.

There are particular considerations regarding dynamic disks and volumes on NetWare, Windows Server 2003, and Linux. See "Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux" for more information.

Background Information

Dynamic Volume Types

Refer to the previous section, "Dynamic Volume Types," for a detailed description of the five dynamic volume types—simple, spanned, striped, mirrored, and RAID-5.

Volume Size in the Create Volume Wizard

It is important to understand that the total volume size you indicate for the volume in the Create Volume wizard is the usable size for storing the data. Depending on the layout type, the volume can take more space on the disk. The size shown in the wizard for

simple, striped, and spanned volumes is the actual volume size on the disk. A RAID-5 volume requires additional space for parity information, and a mirrored volume is a multiple of the space taken by the original plex of the volume.

Steps for Creating a Dynamic Volume

Use the following procedure to create a dynamic volume.

There are particular considerations regarding dynamic disks and volumes on NetWare, Windows Server 2003, and Linux. See "Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux" for more information.

1. From the toolbar, click the **Create Volume** icon. The Create Volume wizard appears. You can read information about partitions and volumes and then click **Next** to continue.

2. The **Select Partition or Volume** screen enables you to choose whether to create a partition or volume. Under **Volume Types**, make sure that the **Dynamic Volume** button is selected.

   Once the **Dynamic Volume** button is selected, the dynamic group to which the volume will belong is automatically displayed. Click **Next** to continue.

3. The **Select Dyanmic Disks** screen enables you to choose the dynamic disks on which the volume will be created. Array Manager selects the dynamic disk for you when the **Auto select disks** button is highlighted.

   To choose the dynamic disk yourself, do the following:

   a. Select the **Manually select disks** button.

   b. Highlight the disk of your choice in the **Available disks** list.

   c. Click the **Add** button to move the disk from the **Available disks** list to the **Selected disks** list.

4. The **Select Volume Type** screen enables you to select the RAID types and the size of the volume.

a. Select the appropriate button under **Layout** to indicate whether the volume should be concatenated, striped, a RAID-5, or mirrored. The following table provides information on these selections. For additional information on concatenation and the RAID types, see "What Is RAID?" and "Choosing RAID Levels and Concatenation."

Selection	Description
Concatenated	Default — Simple or spanned volume.
Striped	You must have unallocated space on two dynamic disks to create a striped volume. The number of columns is automatically selected as 2.
RAID-5	You must have at least three dynamic disks with sufficient unallocated space to create a RAID-5 volume. Note that you cannot create a RAID-5 volume if you are using a Windows 2000 Professional machine.
Mirror	You must have two dynamic disks with sufficient unallocated space to create a mirror. You can mirror only a simple or spanned volume. Any other selection will cause the mirror option on the menu to be inactive (dimmed). You cannot mirror a volume if you are using a Windows 2000 Professional machine.

b. Type the size of the volume you are creating in the **Total volume size** text box. You can specify the size in either megabytes or gigabytes. To identify the maximum possible size of the volume, click the **Query Max Size** button.

c. **(Windows 2000 only.)** Identify the maximum possible size of the volume. If you are creating the volume on a Windows 2000 system, you will see a **Query Max Size** button. The **Query Max Size** button works differently, depending on whether you had a disk selected when you chose the Create Volume command.

- If no disk was selected, clicking the **Query Max Size** button will display the maximum for all available disks in the Total volume size entry box.
- If a disk was selected, clicking the **Query Max Size** button will display the maximum size for the selected disk. However, if you click the button a second time, the size will be the maximum for all available disks.

5. The **Verify Disks** screen enables you to verify the disk on which the volume will be created. If you wish to choose a different disk, click **Modify** and make the appropriate change. When ready, click **Next** to continue.

6. The **Assign Drive Letter** screen enables you to assign a drive letter to the volume. Do one of the following:
   - Choose a drive letter. Click **Next** to continue. Your next step will be to format the volume.
   - If you do not want to assign a drive letter, select **do not assign drive letter**. Click **Next** to continue, then **Finish.** The volume appears as an icon with no name in the Array Manager console. You can assign a drive letter later, if you like.
   - If you are using a Windows 2000 machine, you can also mount the volume at an empty NTFS folder. See the section "Mount a Volume at an Empty Folder" for more information.

7. The **Format Volume** screen enables you to select the file system, label, allocation size, and compression method.

   a. Format the volume using either NTFS or FAT. FAT32 is also available for Windows 2000. File and folder compression is supported only on NTFS volumes.

   b. Leave the allocation unit size as Default.

Note It is strongly recommended that you leave the allocation unit size as **Default**. Selecting any other value requires an advanced knowledge of Windows file systems. You should thoroughly review the Microsoft Windows documentation before selecting any value other than **Default**.

   c. Choose file system type and, if desired, a formatting option:
      - Quick format — Formats the disk without checking for bad sectors.
      - Enable file and folder compression — Can be used only if you choose NTFS format.

8. Click **Next** to continue, then **Finish**.

Extending a Dynamic Simple or Spanned Volume

The capability of extending a simple or spanned volume, which is also known as Online Volume Expansion (OLVE), lets you increase the size of existing simple or spanned volumes. This is independent of the type of storage used, but is most beneficial when the virtual disks are first increased in size.

Note It is recommended that you use only hardware redundant (such as RAID-1 and RAID-5) virtual disks for spanned volumes.

If you are planning to use OLVE with SAN storage, see the section that follows, "Considerations for Using OLVE with SAN Storage," before completing the instructions in this section.

You can extend a volume only if:

- The volume is formatted as NTFS
  and
- there is unallocated space on a dynamic disk onto which volume can be extended.

You cannot extend a volume if any of the following are true:

- The volume is formatted as FAT or FAT32.
- The volume is a striped, mirrored, or a RAID-5 volume.
- There is no unallocated space available on a dynamic disk.

You can extend simple and spanned volumes on dynamic disks onto a maximum of 32 dynamic disks. Once a volume is extended, it cannot be mirrored or striped. No portion of a spanned volume can be deleted without deleting the entire spanned volume.

To Extend a Simple or Spanned Volume

1. Right-click the simple or spanned volume you want to extend and then click **Extend Volume.** The selected volume appears in the dialog box along with its current size.

Note The **Extend Volume** option is disabled if the volume cannot be extended.

2. Enter the amount to extend the volume and click **OK**.

   Ignore any system messages that you may see while the volume extends. You do not have to run any special utilities. Examples of these messages are as follows:

   System Process - Corrupt File

   (X) The file or directory is corrupted and unreadable.

   Please run the Chkdsk utility.

3. Click **OK** to finish. The volume will now show the extended size.

Considerations for Using OLVE with SAN Storage

When using OLVE with SAN storage, the following need to be considered:

- Dynamic disks cannot be moved between servers in a SAN, so any dynamic volumes will be dedicated to a single server.

- It is strongly recommended that you have only a single spanned dynamic volume per server connected to a SAN.
- If you need large volumes, it is better to create larger virtual disks than to create dynamic volumes that span across multiple virtual disks. It is best to have a dynamic volume span across a minimum of virtual disks.
- Remember that online expansion is for dynamic volumes only—basic volumes cannot be extended. On Windows 2000 servers, basic volumes can be converted to dynamic volumes. See the section "Upgrade Disks with Legacy Volumes to Dynamic on Windows 2000" in this chapter.
- Extending dynamic volumes does not maintain any redundancy data on the dynamic disks. If any dynamic disk of a spanned dynamic volume fails, the entire dynamic volume fails, making its data unavailable and possibly corrupt. A spanned dynamic volume is only as reliable as its least-reliable dynamic disk. Therefore, use dynamic disks that are protected by hardware redundant RAID.
- Do not span a dynamic volume across SAN and non-SAN storage or across multiple storage systems in a SAN.
- Do not span more than one dynamic volume to a single dynamic disk. When spanning a dynamic volume, use the entire space of the dynamic disk. In other words, a dynamic disk should contain only a portion of one spanned dynamic volume.
- To ease the procedure of spanning dynamic volumes in a SAN, create dynamic disks only as you need them. This will make it easy to choose which dynamic disk to use in the new spanned dynamic volume.
- Once a dynamic volume is extended, the expansion cannot be undone. Therefore, once a volume is spanned across two disks, it will always be a two-disk volume.
- A dynamic simple or spanned volume can be increased in size but not decreased.
- The Dell PowerVault™ 530F system's snapshot feature is not supported on dynamic volumes.
- Dell PowerEdge FE100 Clusters do not support dynamic volumes.

It is recommended that you use the following procedure for extending a simple or spanned dynamic volume across a new virtual disk in a SAN:

1. Create a new virtual disk (or LUN) on your disk storage system. For the PowerVault 660F, this can be done with Array Manager.

2. Assign the new virtual disk to the SAN server, using the Dell OpenManage™ Storage Consolidation software.

Note Never reassign a SAN LUN that is part of a spanned dynamic volume to another server or Free List. Data loss will occur. Also, it is strongly recommended that you use redundant RAID levels such as RAID-1 or RAID-5 for a virtual disk that will be used as part of a dynamic volume.

3. Find the new disk under the Disks folder using Array Manager and upgrade it to dynamic.

4. Continue with the step-by-step instructions in the main section for "Extending a Dynamic Simple or Spanned Volume" to complete the process of extending the volume. When specifying the amount to extend the volume, be sure to use all the space on the new dynamic disk.

More on Mirrored Volumes

This section gives more information on mirrored volumes and has the following topics:

- Two Methods for Creating a Mirror
- Adding Mirrors
- Removing and Breaking a Mirror
- Reactivating a Mirrored Volume
- Repairing a Mirrored Volume

Two Methods for Creating a Mirror

There are two methods for creating a mirror:

- To create a new mirrored volume, use the Create Volume wizard. See "Creating a Dynamic Volume" in this chapter.
- To create a mirror on an existing volume, use the Add Mirror wizard. See the next topic, "Adding Mirrors."

Note You cannot extend or mirror a dynamic mirrored volume.

Adding Mirrors

The **Add Mirror** option is disabled if the volume cannot be mirrored or if there is no unallocated space on another dynamic disk large enough to mirror the volume.

To Add a Mirror to a Volume

Right-click a volume and then click **Add Mirror**.

Removing and Breaking a Mirror

You can *remove* or *break* a mirror. You should understand the differences between these two procedures.

Removing a mirror from a volume "removes" or destroys the data from the selected mirror and leaves the other mirror intact. After you remove a mirror, the space on the disk used by the removed mirrored volume becomes unallocated space. The remaining (no longer mirrored) volume becomes a simple volume on the disk.

Breaking a mirror creates two simple volumes with individual drive letters. Each volume contains the data on the mirror at the time the mirror was broken. The data is no longer redundant, but it remains intact.

To Remove a Mirror from a Mirrored Volume

Right-click the volume from which you want to remove a mirror and click **Remove Mirror**. You can then follow the instructions on your screen.

Caution Once you remove a mirror from a mirrored volume, the removed part of the mirror becomes unallocated space and the volume is no longer fault tolerant. All of the data on the removed mirror is deleted.

To Break a Mirrored Volume

1. Right-click the volume from which you want to break a mirror and click **Break Mirror.**

2. Choose the mirror to break off from the volume. This volume becomes a simple volume with a different drive letter from the mirrored volume.

Note that the remaining piece of the mirrored volume also becomes a simple volume. It retains the same drive letter of the original mirrored volume.

Reactivating a Mirrored Volume

If there is a problem with one of the disks on which a mirrored volume resides, make sure that the disk is properly connected, and then try to bring that disk back online by first using **Rescan**. If **Rescan** does not work, then use the **Reactivate Disk** command. In most cases, **Reactivate Disk** will bring the disk online and all of the volumes will be healthy.

However, if any of the mirrors or plexes of the mirrored volume are still not healthy, you may need to resynchronize the unhealthy volume with the **Reactivate Volume** command. Resynchronizing makes the data consistent on the plexes of a mirrored volume.

Right-click the volume to bring up the context menu, and then click **Reactivate Volume**.

Caution When reactivating a volume, be aware that the volume's data is restored, even if it is stale, corrupt, or out-of-date. See "Reactivate a Dynamic Volume" for more information on the consequences of reactivating a volume.

The mirrored volume should be brought back to a healthy state unless there is a serious problem with the volume or one of its disks has actually failed. In that case, see the next section, "Repairing a Mirrored Volume."

See also "Rescan to Update Information," "Reactivate a Disk," and "Reactivate a Dynamic Volume" in the "Troubleshooting" chapter.

Repairing a Mirrored Volume

When a disk on which a mirrored volume resides fails, the volume displays a Failed Redundancy status. The disk's name is changed to Missing, and an icon (X) appears on the Missing icon. The status of the disk will be Offline.

To repair a dynamic mirrored volume, you must first use **Remove Mirror** to remove the part of the volume that is on the missing disk and then use **Add Mirror** to add a new mirror, using space from another available dynamic disk.

The difference between reactivating a mirrored volume and repairing it is that reactivating is resynchronizing the volume to bring it to a normal state in its existing location, while repairing involves removing the damaged section of the volume and recreating it in another location on a healthy dynamic disk.

More on RAID-5 Volumes

This section's topics include:

- Reactivating a RAID-5 Volume
- Repairing a RAID-5 Volume

Reactivating a RAID-5 Volume

If there is a problem with one of the disks on which a RAID-5 volume resides, make sure that the disk is properly connected, and then try to bring that disk back online by first using **Rescan**. If **Rescan** does not work, then use the **Reactivate Disk** command. In most cases, **Reactivate Disk** will bring the disk online and all of the volumes will be healthy.

However, if the volume is still not healthy, try to resynchronize the unhealthy volume with the **Reactivate Volume** command. Right-click the volume to bring up the context menu, and then click **Reactivate Volume**.

Caution When reactivating a volume, be aware that the volume's data is restored, even if it is stale, corrupt, or out-of-date. See "Reactivate a Dynamic Volume" for more information on the consequences of reactivating a volume.

The RAID-5 volume should be brought back to a healthy state unless there is a serious problem with the volume or one of its disks has actually failed. In that case, see the next section, "Repairing a RAID-5 Volume."

See also "Rescan to Update Information," "Reactivate a Disk," and "Reactivate a Dynamic Volume" in the "Troubleshooting" chapter.

Repairing a RAID-5 Volume

You can repair a RAID-5 volume if it is in a state of Failed Redundancy, and if there is unallocated space on another dynamic disk available. To avoid data loss, you should attempt to repair the volume as soon as possible.

To Repair a RAID-5 Volume

1. Right-click the volume and select **Repair RAID-5 Volume**.
2. A message appears that indicates that the repair will be attempted if there is another dynamic disk with adequate unallocated space. Click **Yes** to confirm the repair.
3. The volume should be brought back to a healthy state.

Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux

If you have a NetWare or Windows Server 2003 system or are planning on installing Linux, review the following information:

- On NetWare, dynamic disks and dynamic volumes are not supported. Refer to the NetWare documentation about mirroring and NSS volumes for more information.
- In order to implement disk and volume management on Windows Server 2003, you need to use the native Windows Server 2003 Disk Management utility. To locate the Disk Management utility on Windows Server 2003, open the Computer Management (Local) and Storage utilities.

- Linux does not provide support for dynamic volumes. Array Manager also does not support Linux. Therefore, if you use Array Manager to create dynamic disks and volumes and then subsequently install Linux, you will not be able to manage the dynamic disks and volumes you created since Array Manager will not be available on Linux nor does Linux provide dynamic volume support.

Tape 16

This section applies to Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers

The SCSI Subsystem controllers enable tape and CDROM management.

Tape Menu Commands

This section applies to Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers

This section describes the following tape commands:

- Tape Properties

Tape Properties

This section applies to Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers

You can display the properties for tape devices attached to the SCSI subsystem.

To display the tape properties:

1. Right-click the tape object in the Array Manager tree view.
2. Select **Properties** from the pop-up menu. This action displays the **Tape Properties** dialog which displays the tape properties.

Tape properties may include:

- **Name** — This field displays the name of the tape object in the Array Manager tree view. The last digit of the name indicates the tape's SCSI ID.
- **Vendor** — This field displays the tape vendor's name.

- **Status** — This field displays the current status of the tape device. Possible values are:
  - **Ready** — The tape device is functioning normally.
  - **Degraded** — The tape device has suffered a failure of a component and is operating in a degraded state.
  - **Failed** — The tape device has suffered a failure of one or more components and is no longer functioning.
  - **No Media** — The tape device does not have available storage tape.
  - **Device not monitored** — A tape device is attached to the controller, but Array Manager does not monitor the status of this device.
- **Type** — This field indicates the type of storage component. For tape properties, the value in this field is "Tape."
- **Hardware Vendor** — This field displays the tape hardware vendor's name.
- **Product ID** — This field displays the product ID of the tape device.
- **Revision** — This field displays the revision number of the tape device.
- **Serial Number** — This field displays the serial number of the tape device.

CDROM 17

This section applies to Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers

The SCSI Subsystem controllers enable tape and CDROM management.

CDROM Menu Commands

This section applies to Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers

This section describes the following CDROM commands:

- Load (CDROM)
- Eject (CDROM)
- CDROM Properties

Load (CDROM)

This section applies to Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers

You can load media inserted in the CDROM drive.

To load media on a CDROM:

1. Insert the CDROM in the CDROM drive.
2. Right-click the CDROM object in the Array Manager tree view.
3. Select **Load** from the pop-up menu.

Eject (CDROM)

This section applies to Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers

You can eject media inserted in the CDROM drive.

To eject a CDROM from the CDROM drive:

1. Right-click the CDROM object in the Array Manager tree view.
2. Select **Eject** from the pop-up menu.

CDROM Properties

This section applies to Ultra SCSI, Ultra2 SCSI, Ultra160 SCSI, LSI 1020, and LSI 1030 controllers

You can display the properties for CDROM devices attached to the SCSI subsystem.

To display the CDROM properties:

1. Right-click the CDROM object in the Array Manager tree view.
2. Select **Properties** from the pop-up menu. This action displays the **CDROM Properties** dialog which displays the CDROM properties.

CDROM properties may include:

- **Name** — This field displays the name of the CDROM object in the Array Manager tree view. The last digit of the name indicates the CDROM's SCSI ID.
- **Vendor** — This field displays the CDROM vendor's name.
- **Status** — This field displays the current status of the CDROM device. Possible values are:
  - **Ready** — The CDROM device is functioning normally.
  - **Degraded** — The CDROM device has suffered a failure of a component and is operating in a degraded state.
  - **Failed** — The CDROM device has suffered a failure of one or more components and is no longer functioning.
  - **No Media** — The CDROM device does not have a CDROM available.

- **Type** — This field indicates the type of storage component. For CDROM properties, the value in this field is “CDROM.”
- **Hardware Vendor** — This field displays the tape hardware vendor’s name.
- **Product ID** — This field displays the product ID of the tape device.
- **Revision** — This field displays the revision number of the tape device.
- **Capacity** — This field indicates the available capacity on the CDROM.

The Dell PowerVault 660F Storage System 18

This section applies to the PV660F storage system

This chapter describes how to configure and manage arrays with the Dell PowerVault™ 660F Storage System.

Note A Dell PowerVault 660F Storage System must be managed by only one Array Manager managed system. Having more than one Array Manager managed system enabled to manage the same Dell PowerVault 660F Storage System will result in improper function in Array Manager. See "Adding and Removing Fibre Channel RAID and Enclosure Management" in the "Array Manager Utilities" chapter for details.

PowerVault 660F Storage System Description

The Dell PowerVault 660F Storage System is a high-performance Fibre Channel host to Fibre Channel disk array controller. It provides one external Fibre host channel and two internal fibre drive channels conforming to Fibre Channel standards. The logical protocol used for both host and disk communications is 100 MB/s Fibre Channel Protocol (FCP) SCSI over fibre. The storage system provides Fibre Channel host performance and fault-tolerant RAID disk operations for Fibre Channel disk environments. The Fibre Channel RAID controllers used in the storage system are intelligent, caching controllers that support concatenation and RAID levels 0, 1, 3, 5, and 0+1. The storage system can be used in either a dedicated direct-attach configuration, or it can allow multiple hosts in a Storage Area NetWork (SAN) to access the array of physical disk drives (array disks), which can be configured as one or more virtual storage devices (virtual disks).

The Dell PowerVault 660F storage system also provides continuous access to data in the event of a single controller failure. This capability comes with a redundant controller system, using two controllers that share access to the same disk array. In the event of a controller failure, controller operations are assumed by the surviving controller through a failover process. The failed controller can then be removed and replaced while the system is online. The new controller resumes processing array operations in a failback process. During failover and failback, write cache coherency is maintained with the array disks.

The controller supports a battery backup unit (BBU) for maintaining memory content in case of an AC power failure. The principal purpose of the BBU is to provide cache fault tolerance during a power glitch; however, the BBU is capable of sustaining memory content for hours at a time, depending on the type and size of the memory. The fault management features are based on the SCSI-3 Enclosure Services (SES) device interface.

The following features are available on the Dell PowerVault 660F Storage System:

- **Drive Channel Failover** — Provides support for using both ports of a dual-ported fibre disk drive in an active-inactive topology. Each controller can access any of the drives through either of the ports by using the dual fibre channels (loops) in the Dell PowerVault 660F enclosure or PowerVault 224F expansion enclosures. When the channel between the controller and drive fails, the controller automatically fails over to the other channel.
- **Cache Coherency** — Allows server access to a virtual disk (or LUN) through either controller when operating in an active-active mode. It achieves this by operating each controller such that the virtual disks configured with the write-back caching policy are treated as though they were configured for write-through operation and the cache is flushed.
- **Memory DIMM Support** — Additional support for memory DIMM sizes 128 MB, 256 MB, and 512 MB.
- **Enclosure Management** — Provides notification of enclosure status changes for fans, power supplies, and temperature through Array Manager.
- **Conservative Cache Mode** — Provides a higher level of data protection after a critical system component has failed. It automatically switches virtual disks with a write-back cache policy to write-through cache policy and flushes their cache.
- **Enclosure Alarms** — Enabling of enclosure alarms based on drive fault conditions.
- **32 Virtual Disk (LUN) Support** — Allows a maximum of 32 virtual disks per storage system.
- **Real-Time Online Virtual Disk Expansion** — Allows for the expansion of a current virtual disk or adding one or more array disks to the virtual disk's disk group while the storage system is online.
- **Background Initialization (Instant RAID Availability)** — This is only available in version 7.7 or later of the controller firmware. This feature allows a virtual disk to be used while it is being initialized in the background. This feature allows host I/O operations to use a virtual disk immediately after the disk is created without waiting for the virtual disk initialization to complete.
- **112 Array Disk Support** — Supports 112 Fibre Channel array disks assembled in a maximum of 8 enclosures containing a maximum of 14 array disks each.
- **16 Disks per Disk Group** — Allows a maximum of 16 array disks per disk group.

- **Cluster Support** — Supports clusters with the Microsoft Cluster Server.

Supported RAID Levels (PV660F)

This section applies to the PV660F storage system

The Dell PowerVault 660F Storage System supports the following RAID levels:

- Concatenation
- RAID-0
- RAID-1
- RAID-3 and RAID-5
- RAID-0+1

Note With RAID-3 the parity is usually dedicated to a single array disk. The PowerVault 660F storage system, however, implements RAID-3 in the same manner as RAID-5, with distributed parity across multiple disks.

Supported RAID Levels

RAID Level	Description	# of Drives Min.	Max.	Fault Tolerant
Concatenation	Storing data on a single array disk	1	1	No
0	Block striping is provided, which yields higher performance than is possible with individual disk drives. No redundancy is provided.	2	16	No
1	Disk drives are paired and mirrored. All data is 100% duplicated on an equivalent disk drive.	2	2	Yes
3 and 5	Data is striped across several physical disk drives. Parity protection is used for data redundancy.	3	16	Yes
0+1	Combination of RAID levels 0 and 1. Data is striped across disks with mirrors of the stripe chunks on adjacent disks. This can be configured to have an odd or even number of disks.	3	16	Yes

Maximum Virtual Disks per PV660F Storage System

This section applies to the PV660F storage system

You can create a maximum of 32 virtual disks on a Dell PowerVault 660F Storage System.

Conservative Cache Mode

This section applies to the PV660F storage system

Conservative Cache mode is a feature of the Dell PowerVault 660F Storage System. It is an operating mode in which virtual disks configured with the write-back caching policy are treated as though they were configured for write-through operation and the cache is flushed. Conservative Cache mode is entered to provide a higher level of data protection after a critical system component has failed. When the condition forcing Conservative Cache mode is resolved, the cache policies on the system drives are changed back to their original settings.

Conditions that cause Conservative Cache execution are:

- The **Conservative Cache Mode** controller option is enabled for a redundant controller configuration, and one of the controllers has either failed or is disabled.
- A power supply has failed (not power supply removal or power supply not present).
- An enclosure supporting redundant power (multiple power supply support) is operating on a single power supply.
- An enclosure temperature exceeds the warning threshold.
- SES monitoring has failed.

For information on how to change the conservative cache mode setting, see "Advanced Fibre Channel RAID Controller Options." By default, conservative cache mode is enabled.

Understanding Virtual Disk Initialization (PV660F)

This section applies to the PV660F storage system

Initializing a virtual disk makes the disk space on the member array disks available to the operating system for future read and write operations. Array Manager provides two methods for initializing a virtual disk:

- **Foreground initialization.** Use the "Initialize"command to perform a foreground initialization.

- **Background initialization.** Use the "Background Initialize"command to perform a background initialization.

Foreground and Background Initialization

When using foreground initialization, the system cannot perform read and write operations to the virtual disk being initialized. Foreground initialization also writes zeroes (0) to the disks being initialized, effectively destroying any data or redundancy information previously residing on those disks.

When using background initialization, the initialization is performed as a background process which enables the system to perform read and write operations to the virtual disk as it is being initialized.

In contrast to foreground initialization, background initialization does not zero out the existing data on the disks. Instead, background initialization only sets the redundancy or parity data on the virtual disk without changing the user data. This means that for a newly created virtual disk that was background initialized, the data region of the virtual disk will contain any data that existed on the array disks before the virtual disk was created. In order to clear the user data regions in an initialization, you must use the foreground initialize function instead.

Background initialization can only be used with virtual disks configured as redundant RAID levels.

Background initialization is one of the "Advanced Fibre Channel RAID Controller Options." By default, background initialization is enabled.

How the Controller Implements Background Initialization

As long as the **Autostart Background Init** option in the "Advanced Fibre Channel RAID Controller Options" is set to the default setting of enabled, then the system automatically performs a background initialization on any uninitialized virtual disks when a write operation is performed on the virtual disk.

Operations that Suspend Background Initialization

The following operations cause the background initialization process to be suspended:

- Foreground initialization
- Array disk rebuild
- Check consistency
- Expand Capacity
- Add Virtual Disk

System Performance Impacts of Background Initialization

Background initialization does impact system performance. You may notice that read and write operations to virtual disks contained in the same disk group as the virtual disk being initialized are slower when performing a background initialization.

Once the background initialization has started, the system dedicates a certain percentage of the processor bandwidth to the background initialization, regardless of what other operations are being performed on the controller. This percentage of the processor bandwidth will continue to be dedicated to the background initialization until the initialization either completes or is paused.

The impact that background initialization has to system performance is based on the rebuild rate. You can change the rebuild rate by using the "Expert Fibre Channel RAID Controller Options."

Background Initialization and a Disabled Partner Controller

If the partner controller is disabled while a background initialization is in process, the background initialization will experience a short interruption and then continue processing. The disabled controller (or its replacement) are held in reset until the background initialization completes.

You must wait for the background initialization to complete before initiating the "Enable Partner" command. When the "Enable Partner" command is initiated, the partner controller fails back.

Events Generated by Background Initialization

The following events are generated by a background initialization and displayed in the Array Manager event log:

- 905 — Logical drive background initialization started.
- 906 — Logical drive background initialization stopped.
- 908 — Logical drive background initialization restarted.
- 910 — Logical drive background initialization completed.

Redundant Fibre Channel RAID Controller Concepts

This section describes concepts associated with redundant Fibre Channel RAID controllers. This section's topics are:

- Redundant Fibre Channel RAID Controller Configurations
- Redundant Fibre Channel RAID Controller Configuration Requirements

- Fibre Channel Concepts

Redundant Fibre Channel RAID Controller Configurations

Redundant Fibre Channel RAID controller configurations interconnect two identical controllers that share a common set of array disks. This configuration allows a surviving controller to take over resources of a failed controller. This failover process is transparent to the applications running on the host.

Redundant Fibre Channel RAID controller support provides the system with the mechanisms for the following:

- Initializing two controllers to process host requests
- Establishing a controller-controller nexus (C-C nexus) for controller redundancy
- Detecting failure of a controller that is part of a C-C nexus and initiating a failover
- Detecting the replacement of a failed controller, initiating a failback and re-establishment of a C-C nexus.

Note A controller-controller nexus refers to the state in which both redundant controllers are in communication. In this state, each controller can copy write-back data to its partner controller and can determine whether the other controller is operating.

The two Fibre Channel RAID controllers in the disk enclosure will both be connected to a server in a direct attach configuration or in a SAN configuration. The two controllers communicate with each other to verify that both are functioning properly through a ping/acknowledgment sequence. Failure to acknowledge the ping triggers failover.

The redundant Fibre Channel RAID controller configuration also supports dual or multiple server communication in a SAN, offering the advantages of being able to sustain data access in the event of a host failure. If configured in a cluster or high-availability environment, it is also able to sustain data access in the event of the failure of a server or a host bus adapter (HBA). This configuration requires implementation of alternate path software on the server.

Redundant Fibre Channel RAID Controller Configuration Requirements

Both controllers in a redundant Fibre Channel RAID controller system must be identical. The following requirements must be satisfied in order to provide optimal operation.

- Both controllers must have the same firmware version and type.
- Both controllers must have the same amount of memory.

Failover and Failback

In redundant Fibre Channel RAID controller configurations, maintaining continuous access to data requires that a failed controller be replaced in a manner that is transparent to the host. Drive channel ports on the controllers are configured as inactive ports until those ports are needed to respond to requests of a failed controller.

In the event of a controller failure in a redundant Fibre Channel RAID controller system, the failed controller's operations are assumed by the surviving controller. The failed controller can then be removed and replaced while the system is online. When the surviving controller detects the presence of the new controller, the new controller resumes processing array operations. During failover and failback, write cache coherency is maintained with the disk drives.

When a controller is participating in a controller-controller nexus (C-C nexus) and detects a communication error with its partner controller, it initiates the failover process. The following steps outline the failover process executed by the surviving controller:

1. On detection of the controller failure, the surviving controller holds the failed controller in disable partner mode.
2. The surviving controller activates the failover port.
3. Cached data is flushed to the disk drives.
4. Conservative Cache is enabled (if the **Conservative Cache Mode** controller option has been enabled).
5. The surviving controller begins handling I/O requests for the failed controller.

When a failed controller is replaced, the system either automatically detects the replacement (if the **Auto Restore** controller option has been enabled) or is informed of the replacement by the **Enable Partner** command. The following steps outline the failback process executed by the surviving controller:

1. A replacement controller is detected.
2. The surviving controller enables the replacement controller.
3. Once the replacement controller completes initialization and is ready to resume I/O requests, the surviving controller quiesces both ports by responding with BUSY status to new I/O requests.
4. The surviving controller disables the failover port.
5. The surviving controller clears the BUSY condition.

6. The replacement controller enables its primary ports.

7. Both controllers disable Conservative Cache mode (if enabled) for write-back system drives and resume normal redundant Fibre Channel RAID controller operation.

Fibre Channel Concepts

The following sections discuss basic fibre channel concepts. For more information on Fibre Channel technology, refer to the following websites:

www.fibrechannel.com

www.T11.org

www.ncits.org

www.scsita.org

Fibre Channel

A data transfer interface technology that allows for high-speed I/O and networking functionality in a single connectivity technology. The Fibre Channel Standard supports several topologies, including Fibre Channel Point-to-Point, Fibre Channel Fabric (generic switching topology), and Fibre Channel Arbitrated Loop (FC_AL).

Fibre Channel Point-to-Point

The simplest Fibre Channel connection. This provides a direct connection between the transmit output of one system and the receive input of a second system. A second connection is provided between the opposite connectors in order to complete the signals. The physical connection between the two systems is called a link. No switches, loops, or fabric elements are needed. This is also known as a direct attach configuration.

Fibre Channel Fabric

A topology that uses one or more switches to route frames between systems in a Fibre Channel network. The routing of frames is transparent to the systems or devices. This is also known as a Storage Area Network (SAN) configuration.

Fibre Channel Arbitrated Loop

A fibre channel arbitrated loop is a method for connecting multiple devices on a single contiguous loop. All systems in the loop share the same bandwidth. A fibre channel loop provides certain advantages such as fault tolerance and hotswapping among others.

Dell PowerVault 660F Storage System Configuration Overview

This section summarizes the steps for configuring the storage attached to a Dell PowerVault 660F Storage System.

1. Determine requirements for disk groups, virtual disks, and hot spares, including the storage capacity and RAID type for each virtual disk. Make sure the requirements for operating system partitions, fault tolerance, and total capacity are considered during this determination. Note that hot spares are used both for recovery from disk failure and for future capacity expansion.

2. Identify requirements for future capacity expansion. Note that if a virtual disk is to be expanded, it must be the only virtual disk in its disk group. Keep this in mind when determining whether to create one or many virtual disks within a disk group.

3. Create required disk groups and virtual disks. Note the following considerations:

   - Although virtual disks can be added to a disk group at a later time, if the size of a virtual disk is to be expanded later, then the disk group can contain only a single virtual disk.
   - A virtual disk's RAID type must be defined at the time of creation.
   - It is strongly recommended that you fully initialize virtual disks before the virtual disks are accessed by a host computer operating system. Parity information may not be correct until the virtual disk is initialized. To minimize initialization time, you can begin initialization of each virtual disk before creating the next disk group. You can also use background initialization if the virtual disk is needed before initialization is fully complete. See "Initialize" and "Background Initialize" for more information.
   - The RAID types available depend on the number of array disks in the disk group.
   - The stripe size for all virtual disks managed by a PowerVault 660F storage system is set to the same value. It is determined by the stripe size setting on the first virtual disk created. For the remaining virtual disks, the same stripe size is used, so the strip size field is inactive (dimmed).

4. Assign one or more hot spares to enable auto-rebuild in the event of drive failure within a fault-tolerant virtual disk.

5. On a Windows system, if the PowerVault 660F storage system is part of a SAN, you must first assign the virtual disk to a server by using the Dell OpenManage™ Storage Consolidation software before a server can use the virtual disk. Once assigned to a server (or if configured in a "direct-attach" hardware configuration) you must perform a rescan in order to access the virtual disk.

6. Create and format operating system partitions to complete the configuration. See the "Disks" and "Volumes" chapters for details on how to perform this step with Windows 2000.

Dell PowerVault 660F Storage System Commands and Procedures Summary

The remainder of this chapter provides detailed information on the Dell PowerVault 660F Storage System commands and procedures and has the following main topics:

- Array Group Commands (PV660F)
- Disk Group Commands (PV660F)
- Virtual Disk Commands (PV660F)
- Array Disk Commands (PV660F)
- Fibre Channel RAID Controller Commands
- Subsystem Commands (PV660F)
- Enclosure Management for the PV660F Storage System
- How the Battery Backup Unit Works

Array Group Commands (PV660F)

This section applies to the PV660F storage system

An **array group** contains disk groups of array disks that are controlled by a particular storage system. Array disks are numbered according to the enclosure number and the enclosure slot number. For example, Array Disk 1:7 indicates that the disk belongs to the first enclosure and resides in slot 7 of the enclosure. The enclosure slot numbers are accessible on the exterior of the enclosure allowing for easy identification of a particular array disk.

Because the Dell PowerVault 660F Storage System can support 112 array disks in a single array group, array disks are further divided into disk groups with up to 16 array disks each. You create one or more virtual disks from a disk group.

You can perform the following operations from the Array Group context menu:

- Create groups of array disks — Use the Create Disk Group function to select specific array disks to be configured as a disk group.
- Add global hot spare disks — Use the Add Spares function to add global hot spare disks to the array group.
- Delete all information on the controller — Use the Reset Configuration function to clear the controller in order to create a fresh configuration.

- Customize RAID policies — Use the Edit RAID Policy Wizard to create, edit, or delete RAID policies for the Fibre Channel RAID controller.
- View information on a selected array group — Use the Array Group Properties command.

To bring up an array group's context menu, right-click the array group object in the tree view.

Create Disk Group

This section applies to the PV660F storage system

Use this function to create a disk group of available physical array disks. The disk group of array disks will be used to create virtual disks. You can have up to 16 array disks per disk group. Before you create a disk group, consider the following items:

- The number of disks needed for the RAID levels you want to implement.
- The number of virtual disks to include in the disk group. Note that if a virtual disk is to be expanded later, it must be the only virtual disk in the disk group and use all the disk space of the member array disks.
- A disk group can be left with free space available on all its array disks. This available space can be later used to create one or more additional virtual disks.

To Create a Disk Group

1. Right-click an array group.
2. Click **Create Disk Group**. The Create Disk Group dialog box appears.
3. Select one or more array disks from the Available Array Disks list. Click **Add Disk** to add the disks to the Add Array Disks list. Continue adding disks as required by your configuration.
4. Disks can be removed from the Add Array Disks list by selecting the specific disk and clicking on **Remove Disk**, or remove all disks from the Add Array Disks list by clicking on **Remove All**.
5. Verify that the Add Array Disks list contains the appropriate disks. Click **OK** to invoke the changes. (Click **Cancel** to discontinue the operation.)

Add Spares

This section applies to the PV660F storage system

Use this function to create a global hot spare disk. A global hot spare disk is an array disk that will be used to rebuild disk groups containing redundant virtual disks. A hot spare is also required before you can use the "Expand Capacity" or "Add Virtual Disk" commands.

Caution Only disk groups that don't contain non-redundant virtual disks are eligible to have a global hot spare rebuild a failed array disk. Disk groups containing concatenated virtual disks or RAID 0 virtual disks will not have a failed array disk rebuilt with a global hot spare disk.

1. Right-click an array group.
2. Click **Add Spares**. The Add Spares dialog box appears.
3. Select one or more array disks from the Available Array Disks list. Click **Add Disk** to add the disks to the Add Array Disks list. Continue adding spare disks as required by your configuration.
4. Disks can be removed from the Add Array Disks list by selecting the specific disk and clicking on **Remove Disk**, or remove all disks from the Add Array Disks list by clicking on **Remove All**.
5. Verify that the Add Array Disks list contains the appropriate disks. Click **OK** to make the changes. (Click **Cancel** to discontinue the operation.)

A specific array disk with Ready status can also be made a global hot spare by right clicking on it and selecting **Make global hotspare**.

Note If multiple array disks that are different sizes are configured as global hot spares, a rebuild will choose the array disk that best fits the array disk that failed.

Reset Configuration

This section applies to the PV660F storage system

Use this function to delete all configuration information on the Fibre Channel RAID controller, so you can create a fresh configuration.

Caution Resetting a configuration permanently destroys all data on the Fibre Channel RAID controller and array disks. This operation is not recoverable.

Note It is strongly recommended that you first delete all the volumes (and unassign the virtual disks using Dell OpenManage Storage Consolidation if the array is configured in a Storage Area Network) before performing a Reset Configuration.

1. Right-click an array group.
2. Click **Reset Configuration**. A warning/confirmation dialog box appears warning that all virtual disks and data will be lost. The user is asked to confirm this operation.
3. Click **Yes** to continue. (Click **No** to exit the function.)

Edit RAID Policy

This section applies to the PV660F storage system

Use this function to create, edit, or delete RAID policies.

To Create a New RAID Policy

1. Right-click a disk group.
2. Click **Edit RAID Policy**. The RAID Policy Editor screen appears.
3. Click **New Policy**. An empty detailed RAID Policy Editor screen appears.
4. Click **Add Disk**. A disk with empty attribute fields appears on the screen.
5. Edit the fields by typing the appropriate information, or double-click in the field for a context menu of possible settings.

Note When creating multiple disks in a RAID Policy, keep in mind the restrictions of RAID levels and the number of array disks needed. For example, having a RAID policy containing a RAID 1 and RAID 5 isn't possible since RAID 1 is only valid with 2 array disks, and RAID 5 requires at least 3 array disks.

6. Add more disks as required for your configuration.

The last disk in the list can be removed from the configuration by clicking **Remove Disk**.

7. Enter a name in the Policy Name field.

8. When all settings are correct, click **OK** to return to the main RAID Policy Editor screen. (Click **Cancel** to discontinue the operation.)

9. Click **OK** to save the configuration. (Click **Cancel** to discontinue the operation.)

To Edit an Existing RAID Policy

1. Right-click a disk group.

2. Click **Edit RAID Policy**. The RAID Policy Editor screen appears.

3. Select the RAID policy you wish to modify.

4. Click **Modify Policy**. A detailed RAID Policy Editor screen appears showing the attributes of the selected RAID policy.

5. Edit the fields by typing the appropriate information, or double-click in the field for a context menu of possible settings.

6. Add or delete disks as required by your configuration.

7. When all settings are correct, click **OK** to return to the main RAID Policy Editor screen. (Click **Cancel** to discontinue the operation.)

8. Click **OK** to save the configuration. (Click **Cancel** to discontinue the operation.)

To Delete a RAID Policy

1. Right-click a disk group.

2. Click **Edit RAID Policy**. The RAID Policy Editor screen appears.

3. Select the RAID policy you wish to delete.

4. Click **Delete Policy**.

5. Click **OK** to save the configuration. (Click **Cancel** to discontinue the operation.)

Array Group Properties

This section applies to the PV660F storage system

The **Properties** command for a selected array group displays the following information:

- **Name** — This field displays the array group number, which is equivalent to the Fibre Channel RAID controller number.
- **Controller** — This field displays the number of the Fibre Channel RAID controller with which this array group is associated. There is one number for the pair of controllers which corresponds to the Array Manager controller object number.
- **Status** — This field displays the current status of the array group. Possible values are:
  - **Ready** — The array group is operational.
  - **Online** — The array group has been configured and is functioning normally.
  - **Offline** — The array group has not been configured, or the configuration has just been cleared.
  - **Failed** — The array group has failed to operate properly.
  - **Rebuild** — The array group has been configured and is in the process of being rebuilt.

Note If the status value is anything other than Ready, a red X appears to the left of the status to indicate a situation that requires attention.

Disk Group Commands (PV660F)

This section applies to the PV660F storage system

You can perform the following operations from the Disk Group context menu:

- Create virtual disks — Use the Create Virtual Disk (PV660F) Wizard to create virtual disks.
- Apply RAID policies — Use the Apply RAID Policy Wizard to apply RAID policies to virtual disks and volumes.
- Add capacity to a virtual disk — Use the Expand Capacity function to increase the size of the virtual disk.
- Add virtual disks to a disk group — Use the Add Virtual Disk function to add virtual disks to a already full array disk group.

- Delete a disk group — Use the Delete Disk Group function to remove a disk group from the array group.
- View information on a selected disk group — Use the Disk Group Properties command.

To bring up a disk group's context menu, right-click the disk group object in the tree view.

Create Virtual Disk (PV660F)

This section applies to the PV660F storage system

Creating virtual disks allows you to choose from a variety of RAID configurations to make the best use of your array disk storage. You must create a disk group before creating a virtual disk, because the array disks for a virtual disk are chosen by a disk group. See "Create Disk Group" in this chapter for details.

The Dell PowerVault 660F Storage System supports the following RAID configurations: concatenation, RAID-0, RAID-1, RAID-3, RAID-5, and RAID-0+1. For more information about RAID configurations, see "What Is RAID?" and "Choosing RAID Levels and Concatenation" in the "Understanding Storage Concepts" chapter.

Note When you create virtual disks, you may want to consider how virtual disks must be deleted. Because of hardware constraints, virtual disks must be deleted in reverse order of their creation. In other words, you must delete the most recently created virtual disk first and continue to delete in reverse order.

1. Right-click a disk group. For information on creating a disk group, see "Create Disk Group."
2. Select **Create Virtual Disk** from the **Task** pull-down menu. The Welcome to the Create Virtual Disk Wizard screen appears.
3. Read the welcome screen and click **Next** to continue. (Click **Cancel** to discontinue the operation.) The Select Virtual Disk Type screen appears.
4. Select the Type, Size, Stripe Size (only settable for the storage system's first virtual disk), and Write Policy for the virtual disk.

   Supported virtual disk types or RAID levels are available in the drop-down menu of the Type field. The choices shown in the drop-down menu depends on the number of array disks in the disk group.

Note When creating virtual disks in a disk group, use non-redundant RAID levels (such as concatenated or RAID 0) only when the disk group will not need hot-spare disk replacement. If a disk group contains at least one non-redundant virtual disk, any failed array disks in the disk group cannot be rebuilt with a hot-spare array disk.

The actual size of the virtual disk may be larger than what you specify. This is because the virtual disk size is rounded up according to the RAID level and the number of array disks in the disk group. For more information on how the actual disk size is determined, see "Virtual Disk Size."

The stripe size is selectable only on the first virtual disk created. Supported stripe size values are available in the drop-down menu of the Stripe Size field. When you create additional virtual disks, the stripe size setting is inactive (dimmed) because there can be only one stripe size for all virtual disks on the same Dell PowerVault 660F storage system.

The Write Policy toggles between Write Cache Enabled and Write Cache Disabled. When write cache is enabled, the Fibre Channel RAID controller returns a completion status to the host as soon as the cache memory receives the data. The array disk receives the data at a more appropriate time. When write cache is disabled, the controller writes data to the array disk before returning completion status to the host.

5. When the virtual disk attributes are set appropriately, click **Next** to continue. (Click **Back** to return to the previous screen. Click **Cancel** to discontinue the operation.) The Completing the Create Virtual Disk Wizard screen appears.

Note Initializing the new virtual disk is strongly recommended and may take approximately one hour to complete for foreground initialization, depending on the size of the virtual disk. Multiple simultaneous disk foreground initializations are supported, but having more than eight virtual disk initializations running at the same time can slow down the response time of commands from Array Manager to that PowerVault 660F fibre channel array. Array Manager does not allow initialization to immediately occur after virtual disk creation if eight virtual disk initializations are already ongoing. You will have to later manually start an initialization when one or more of the current initializations are finished, or you can plan to use background initialization to have the virtual disks initialized.

6. Check the **Initialize Disk** option if you want to perform a foreground initialization on the new virtual disk immediately, and click **Finish** to create the virtual disk as determined in the previous steps. (Click **Back** to return to the previous screen. Click **Cancel** to discontinue the operation.)

Note If you want, initialization can be performed later by choosing **Initialize** or **Background Initialize** (if available) from the Virtual Disk context menu. Also, if the **Autostart Background Init** controller option is set, the virtual disk is automatically

queued up for background initialization when it is created. The Background Initialize controller option is only available when using PowerVault 660F controller firmware 7.7 or later.

Note Starting a foreground initialization (by checking the Initialize Disk option) will suspend any active background initializations on another virtual disk. The background initialization will continue once all foreground initializations are either completed or canceled.

If this is the first virtual disk created, a dialog box notifying the user that the controller needs to reset appears. The reset operation may take more than 30 seconds to complete.

7. After foreground initialization completes (or background initialization has started to run), do one of the following:

- If the PowerVault 660F array is part of a SAN, assign the virtual disk to a server by using the Dell OpenManage Storage Consolidation software. After assigning the virtual disk, you should rescan the console on Windows 2000 and Windows Server 2003.
- If the PowerVault 660F Fibre Channel RAID controller is connected through direct attach configuration, you will need to perform an Array Manager console rescan for the server to be able access the virtual disk.

8. Add a disk signature to the virtual disk.

   For instructions on writing a disk signature, see the section "Write a Disk Signature" in the "Disks" chapter.

Your virtual disk is now ready to be used by the server.

Virtual Disk Size

Because the PowerVault 660F storage system rounds up the size of the virtual disk, the actual size of the virtual disk may be larger than what you specified when creating the virtual disk. The method for rounding up the virtual disk size varies depending on the RAID level and the number of array disks in the disk group. The following table indicates how the virtual disk size is rounded up.

Virtual Disk Size on a PowerVault 660F Storage System

RAID Level	Method for Determining Virtual Disk Size
RAID 0 and Concatenated	2 MB x (number of array disks in the disk group)
RAID 1 and RAID 0 + 1	1 MB x (number of array disks in the disk group)
RAID 3 and RAID 5	2 MB x (number of array disks in the disk group - 1)

Apply RAID Policy

This section applies to the PV660F storage system

Use this function to apply RAID policies to create virtual disks.

Note If you use this method, all your virtual disks will be in one disk group. You will not be able to use the **Expand Capacity** command to add array disks in the future if the policy contains more than one virtual disk.

1. Right-click a disk group.
2. Click **Apply RAID Policy**. The Welcome to the RAID Policy Wizard screen appears.
3. Read the welcome screen and click **Next** to continue. (Click **Cancel** to discontinue the operation.) The Select RAID Policy screen appears.
4. Use the RAID Policy Browser to select a RAID policy to apply to the array disk group. Check the disk(s) for which this RAID policy is to be applied. Changes can be made to this policy's settings by double-clicking on any of the fields. These changes will be applied "on the fly" without editing or creating a new RAID policy. Click **Revert to default** to revert to the existing policy's default settings.

Note When changing the RAID level for disks in a RAID Policy, keep in mind the restrictions that apply to the number of array disks required for a particular RAID level. For more information on these restrictions, see "Supported RAID Levels (PV660F)." If you select a RAID level that is inappropriate for the number of available array disks, an error message displays.

5. When the RAID policies are set appropriately, click **Next** to continue. (Click **Back** to return to the previous screen. Click **Cancel** to discontinue the operation.) The Completing the RAID Policy Wizard screen appears.

Note Initializing the new virtual disk is strongly recommended and may take one hour or longer to complete for foreground initialization, depending on the size of the virtual disk. Multiple simultaneous disk foreground initializations are supported, but having more than eight virtual disk initializations running at the same time can slow down the response time of commands from Array Manager to the PowerVault 660F fibre channel array. Array Manager will not allow initialization to immediately occur after virtual disk creation if eight virtual disk initializations are already ongoing. You will have to later manually start a initialization when one or more of the current initializations are finished.

6. Check the Initialize Disk option if you want to immediately perform a foreground initialization on the new virtual disk, and click **Finish** to create the virtual disk as determined by the RAID policy chosen. (Click **Back** to return to the previous screen. Click **Cancel** to discontinue the operation.)

Note If you wish, initialization can be performed later by choosing **Initialize** or **Background Initialize** from the Virtual Disk context menu. Also, if the **Autostart Background Init** controller option is set, the virtual disk is automatically queued up for background initialization when it is created.

Note Starting a foreground initialization (by checking the Initialize Disk option) will suspend any active background initialization occurring on another virtual disk. The background initialization will continue once all foreground initializations are either completed or canceled.

If this is the first virtual disk created, a dialog box notifying the user that the controller needs to reset appears. The reset operation may take more than 30 seconds to complete.

7. After the foreground initialization completes (or background initialization starts), do the following:

- If the PowerVault 660F Fibre Channel RAID controller is part of a SAN, assign the virtual disk(s) to a server by using the Dell OpenManage Storage Consolidation software. After assigning the virtual disk, you should rescan the console on Windows 2000 or Windows Server 2003.
- If the PowerVault 660F Fibre Channel RAID controller is connected through direct attach, you must rescan for the server to be able access the virtual disk(s).

8. Add a disk signature to the virtual disk.

   For instructions on writing a disk signature, see the section "Write a Disk Signature" in the "Disks" chapter.

After applying the RAID policy to a virtual disk, you may note that the size of the virtual disk is larger than the size you specified when creating the virtual disk. This is because the virtual disk size is rounded up according to the RAID level and the number of array disks in the disk group. For more information on how the actual disk size is determined, see "Virtual Disk Size."

Expand Capacity

This section applies to the PV660F storage system

Use this function to increase the capacity of an existing disk group while the storage system is online. The **Expand Capacity** command appends the added array disk's capacity to the existing virtual disk in the disk group. In contrast, the **Add Virtual Disk** command expands a disk group by adding a new virtual disk, using the existing and added array disks.

The **Expand Capacity** command is not available (menu option is dimmed) unless the following conditions are met:

- The PowerVault 660F storage system is operating with the partner controller held in reset. Use the **Disable Partner** command to place the partner controller in reset before invoking **Expand Capacity**.
- This command can be performed only on disk groups with one virtual disk, which must not be configured as a concatenated RAID level.
- The array disks to be added to the disk group are configured as hot spares. See "Add Spares" or "Make Global Hot Spare" for information on configuring a disk as a hot spare.

Note that the following considerations apply when using the **Expand Capacity** command:

- Starting an Expand Capacity suspends any background initialization occurring on another virtual disk. The background initialization will continue once the Expand Capacity is completed or canceled.

- When you expand the capacity of a virtual disk with RAID level 0, the virtual disk's RAID level is automatically converted to RAID 0+1.

Use the following procedure to increase the capacity of an existing virtual disk in a disk group:

1. Right-click a disk group containing the virtual disk to be expanded.
2. Click **Expand Capacity**. The Expand Capacity dialog box appears.
3. Select one or more array disks from the Available Array Disks list. Click **Add Disk** to add the disks to the Add Array Disks list. Continue adding disks until you want to stop or you reach the 16 array disk limit.

   Array disks can be removed from the Add Array Disks list by selecting the specific disk and clicking on **Remove Disk**, or remove all disks from the Add Array Disks list by clicking on **Remove All**.
4. Verify that the Add Array Disks list contains the appropriate disks. Click **OK** to make the changes. (Click **Cancel** to discontinue the operation.)

   A temporary disk group appears that contains the selected array disks. This disk group remains until the **Expand Capacity** procedure is completed. When the procedure is finished, the temporary disk group will be removed and the selected array disks are displayed under the disk group that contains the virtual disk that has been expanded.

   The **Expand Capacity** process may take hours to complete, depending on the number of disks involved and their capacity.
5. After the **Expand Capacity** process has completed, you MUST enable the partner controller to re-establish controller redundancy.

   See the Enable Partner commands.

Caution No rebuilds of redundant virtual disks will take place until after the **Expand Capacity** operation is completed and the partner controller is enabled.

6. If the virtual disk is on a Windows 2000 or Windows Server 2003 server, you must perform a console rescan for the server to recognize the additional space.
7. You are now ready to either expand existing file systems on the free space, or put a format and a file system on the free space.

   Refer to "Working with Dynamic Volumes" in the "Volumes" chapter.

Add Virtual Disk

This section applies to the PV660F storage system

This function allows additional array disks to be added to a disk group, using the additional space of the added array disks as a new virtual disk. Existing virtual disks in the disk group are restriped across the new set of array disks.

The **Add Virtual Disk** command is not available (menu option is dimmed) unless the following conditions are met:

- The PowerVault 660F storage system is operating with the partner controller disabled. Use the **Disable Partner** command to hold the partner controller in reset before invoking **Add Virtual Disk**.
- The array disks to be added to the disk group for the new virtual disk are configured as hot spares. See "Add Spares" or "Make Global Hot Spare" for information on configuring a disk as a hot spare.
- The number of virtual disks for the storage system is less than 32. When creating virtual disks on the PowerVault 660F storage system, you cannot exceed the 32 virtual disk limit.

Note Starting an Add Virtual Disk suspends any background initialization occurring on another virtual disk. The background initialization will continue once the Add Virtual Disk is completed or canceled.

The Add Virtual Disk adds a new virtual disk, using the existing array disks.

1. Right-click a disk group.
2. Click **Add Virtual Disk**. The Welcome to the Add Virtual Disk Wizard screen appears.
3. Read the welcome screen and click **Next** to continue. (Click **Cancel** to discontinue the operation.) The Select Virtual Disk Type screen appears.
4. Select the Type and Write Policy for the virtual disk.

   Supported virtual disk types or RAID levels are available in the drop-down menu of the Type field.

   Supported RAID levels are described in the Supported RAID Levels table near the beginning of this chapter.

Note When creating virtual disks in a disk group, use non-redundant RAID levels (such as concatenated or RAID 0) only when the disk group will not need hot-spare disk replacement. If a disk group contains at least one non-redundant virtual disk, any failed array disks in the disk group cannot be rebuilt with a hot-spare array disk.

The Write Policy toggles between Write Cache Enabled and Write Cache Disabled. When write cache is enabled, the controller returns a completion status to the host as soon as the cache receives the data. The target device receives the data at a more appropriate time. When write cache is disabled, the controller writes data to the target device before returning completion status to the host.

5. When the virtual disk attributes are set appropriately, click **Next** to continue. (Click **Back** to return to the previous screen. Click **Cancel** to discontinue the operation.) The Select Disk screen appears.

6. Select one or more spare disks for this virtual disk by checking the appropriate disk(s).

7. When the spare disk or disks are selected, click **Next** to continue. (Click **Back** to return to the previous screen. Click **Cancel** to discontinue the operation.) The Completing the Add Virtual Disk Wizard screen appears.

8. Click **OK** to continue.

9. After the **Add Virtual Disk** process has completed, you MUST enable the partner controller to re-establish controller redundancy. See the Enable Partner command for more information.

Caution No rebuilds of redundant virtual disks will take place until after the **Add Virtual Disk** operation is completed and the partner is enabled.

10. If you are using a SAN, see the *SAN Administrator's Guide* for instructions on how to assign the new virtual disk to a server, using the Dell OpenManage Storage Consolidation software.

After using the Add Virtual Disk command to increase the size of a virtual disk, you may note that the size of the virtual disk is larger than expected. This is because the virtual disk size is rounded up according to the RAID level and the number of array disks in the disk group. For more information on how the actual disk size is determined, see "Virtual Disk Size."

Delete Disk Group

This section applies to the PV660F storage system

Use this function to remove the selected disk group from the array group.

Note Before you can delete a disk group, you must have deleted the virtual disks in the disk group.

1. Right-click a disk group.
2. Click **Delete Disk Group**. A warning/confirmation dialog box appears notifying the user that the selected disk group will be deleted. The user is asked to confirm this operation.
3. Click **Yes** to continue. (Click **No** to exit the function.)

Note After deleting the last disk group, you must reset the configuration. See "Reset Configuration" in this chapter.

Disk Group Properties

This section applies to the PV660F storage system

The **Properties** command for a disk group gives the following information on the selected disk group:

- **Group Size** — This field displays the total size of the disk group. This is the sum of all selected disks' physical sizes configured in this disk group.
- **Used Space** — This field displays the amount of used or configured space.
- **Free Space** — This field displays the amount of free or unused space.
- **Minimum Stripe Size** — This field displays the minimum stripe size capable by the PowerVault 660F storage system.
- **Maximum Stripe Size** — This field displays the maximum stripe size capable by the PowerVault 660F storage system.
- **Default Stripe Size** — This field displays the current stripe size being used. Supported stripe sizes are 4 KB, 8 KB, 16 KB, 32 KB, and 64 KB.
- **Status** — This field displays the current status of the disk group. Possible values are:
  - **Ready** — The disk group is operational.

- **Online** — The disk group has been configured and is functioning normally.
- **Offline** — The disk group has not been configured or the configuration has just been cleared.
- **Failed** — The disk group has failed to operate properly.
- **Rebuild** — The disk group has been configured and is in the process of being rebuilt.

Note If the status value is anything other than Ready, a red X appears to the left of the status to indicate a situation that requires attention.

Virtual Disk Commands (PV660F)

This section applies to the PV660F storage system

You can perform the following operations from a virtual disk's context menu:

- Change write caching policy — Use the Change Policy function to change the write caching policies of your virtual disk.
- Delete a virtual disk — Use the Delete Virtual Disk function to delete the virtual disk. This is only selectable on the virtual disk that was created last on that storage system.
- Initialize a virtual disk in the foreground— Use the Initialize function to manually start a foreground initialization on the virtual disk.
- Cancel foreground initializations — Use the Cancel Initialize function to cancel all foreground initializations that are in progress.
- Initialize a virtual disk in the background --- Use the Background Initialize function to perform a background initialization on the virtual disk. This allows the virtual disk to be simultaneously used while still being initialized by the controller.
- Cancel a background initialization --- The Cancel Background Initialize function stops a background initialization on a virtual disk.
- Check redundant data information — Use the Check Consistency (PV660F) function to verify the integrity of data on the virtual disk.
- Cancel a check consistency — Use the Cancel Check Consistency function to cancel a consistency check that is in progress.
- Force a virtual disk online — Use the Force Online function to force a failed virtual disk (due to offline array disks) to an online status.
- View information on a selected virtual disk — Use the Virtual Disk Properties (PV660F) command.

To bring up a virtual disk's context menu, right-click the virtual disk object in the tree view. For more information why you would want to use virtual disks, see the section "Create Virtual Disk (PV660F)" in this chapter.

Change Policy

This section applies to the PV660F storage system

Use this function to change the write cache policy of a virtual disk. This setting will be used when the PowerVault 660F storage system is not in conservative cache mode.

1. Right-click a virtual disk.
2. Click **Change Policy**. The Write Policy dialog box appears.

   Select either Write Cache Enabled or Write Cache Disabled. The Write Policy toggles between Write Cache Enabled and Write Cache Disabled. When write cache is enabled, the controller returns a completion status to the host as soon as the cache receives the data (improving write performance). The array disk receives the data at a more appropriate time. When write cache is disabled, the controller writes data to the array disk before returning completion status to the host.
3. Click **OK** to make the change. (Click **Cancel** to discontinue the operation.)

Note If the controller is in conservative cache mode, the virtual disk will operate with write cache disabled regardless of the write cache policy setting.

Delete Virtual Disk

This section applies to the PV660F storage system

The **Delete** command for virtual disks is not available (menu option is dimmed) unless the following conditions are met:

- The virtual disk is the last virtual disk that was created on the PowerVault 660F storage system.
- None of the virtual disks in the array are being rebuilt.
- None of the virtual disks in the array are being initialized.
- A consistency check is not being performed on any virtual disk in the array.

- None of the disk groups are being expanded with the Expand Capacity or Add Virtual Disk commands.

Caution Deleting a virtual disk permanently deletes all information contained on that disk. It is strongly recommended that you first delete the volumes from the virtual disk (and offline the virtual disk from the server using Dell OpenManage Storage Consolidation if the PowerVault 660F is part of a Storage Area Network) before deleting the virtual disk.

Note This function is enabled only for the virtual disk that was created last. Because of hardware constraints, virtual disks must be deleted in the reverse order of their creation. In other words, you must delete the most recently created virtual disk first and continue to delete in reverse order.

Use this function to delete the virtual disk that was created last.

1. Right-click the virtual disk that was created last.
2. Click **Delete**. A warning/confirmation dialog box appears warning that all data on the selected virtual disk will be lost. The user is asked to confirm the operation.
3. Click **Yes** to continue. (Click **No** to exit the function.) The virtual disk disappears from the left pane and the right pane in the Array Manager console if the array group is selected.

Initialize

This section applies to the PV660F storage system

Use this function to manually start a foreground initialization of a virtual disk. It is strongly recommended that you fully initialize a virtual disk before the virtual disk is accessed by a host computer operating system. Parity information may not be correct until the virtual disk is initialized.

The system will not be able to perform read and write operations on the virtual disk while it is being foreground initialized. You can also use background initialization, which enables the system to perform read and write operations during the initialization process. For more information on foreground and background initialization, see "Understanding Virtual Disk Initialization (PV660F)."

Note It is possible to have foreground initializations ongoing for multiple virtual disks, but performing more than eight simultaneous virtual disk initializations at the same time can slow down the response time of commands between Array Manager and

the PowerVault 660F array where the storage is being initialized. Array Manager will not allow more than eight simultaneous foreground initializations on a PowerVault 660F array.

Note Starting a foreground initialization will suspend any background initialization occurring on another virtual disk. The background initialization will continue once all foreground initializations are either completed or canceled.

Caution Performing a foreground initialization on a virtual disk will destroy all the data on that virtual disk.

1. Right-click a virtual disk.
2. Click **Initialize**. A warning/confirmation dialog box appears warning that all current data on the virtual disk will be lost and recommends that the user not perform any disk or volume operations on the virtual disk while the initialization is in progress.
3. Click **Yes** to continue. (Click **No** to exit the function.) The initialization progress is indicated in the right pane. The virtual disk status indicates “Initializing” and the progress indicates the percent completed.

Cancel Initialize

This section applies to the PV660F storage system

Use this function to manually stop ALL foreground virtual disk initializations that are in progress for a single PowerVault 660F storage system.

1. Right-click on a virtual disk that has an active foreground initialization.
2. Click **Cancel Initialize**. A warning/ dialog box appears asking you to confirm the operation.
3. Click **Yes** to continue. (Click **No** to exit the function.)

Caution A redundant RAID configured virtual disk must have a complete initialization in order to set consistency on the disk. This can be done with either foreground or background initialization.

Background Initialize

This section applies to the PV660F storage system

Note The Background Initialize controller option is only available when using PowerVault 660F controller firmware 7.7 or later.

Use this command to initialize a virtual disk without interrupting the normal functioning of the virtual disk and to preserve the data on it. The **Background Initialize** command enables a server to use a virtual disk while the background initialization process is active. Background initialization can only be used with virtual disks configured as redundant RAID levels.

It is strongly recommended that you fully initialize a virtual disk before the virtual disk is accessed by a host computer operating system. Parity information may not be correct until the virtual disk is initialized. You can use the **Background Initialize** command to fully initialize the virtual disk.

For more information on virtual disk initialization, see "Understanding Virtual Disk Initialization (PV660F)." This information includes a discussion of background initialization's impact on system performance and a list of operations that interrupt a background initialization.

To manually start a background initialization:

1. Right-click on the virtual disk for which you want to start background initialization.
2. Click **Background Initialize**. A confirmation dialog box is displayed.
3. Click **Yes** to continue. (Click **No** to exit the function.) The background initialization process for that virtual disk will begin when the controller receives its next write request.

Note The **Background Initialize** menu will not be enabled if another virtual disk is already being initialized (in foreground or background). To queue up virtual disks for background initialization, you can set the **Autostart Background Init** option in the Advanced Controller Options. See "Advanced Fibre Channel RAID Controller Options" for more information on setting the **Autostart Background Init** option.

For information on canceling the background initialization, see "Cancel Background Initialize."

Cancel Background Initialize

This section applies to the PV660F storage system

Note The Cancel Background Initialize controller option is only available when using PowerVault 660F controller firmware 7.7 or later.

Use this function to manually stop a background virtual disk initialization that is in progress.

1. Right-click on the virtual disk that is being background initialized.
2. Click **Cancel Background Initialize**. A warning/confirmation dialog box appears asking you to confirm the operation.
3. Click **Yes** to continue. (Click **No** to exit the function.)

Caution A redundant RAID configured virtual disk must have a complete initialization in order to set consistency on the disk.

Check Consistency (PV660F)

This section applies to the PV660F storage system

Use this function to begin a consistency check. A consistency check verifies that mirror or parity information matches the stored data on the redundant RAID configured virtual disk (RAID-1, RAID-3, RAID-5, or RAID-0+1). If the parity block information is inconsistent with the data blocks, events will be generated in the event log showing this.

Note If a consistency check fails for a virtual disk, it will be necessary to initialize the virtual disk and restore the data from backup to correct the virtual disk.

Note Starting a check consistency will suspend any background initialization occurring on another virtual disk. The background initialization will continue once the check consistency is either completed or canceled.

1. Right-click a virtual disk.
2. Click **Check Consistency**. A warning/confirmation dialog box appears warning that other disk or volume operations on the corresponding disks not be performed while the consistency check is in progress. The user is asked to confirm this operation.

3. Click **Yes** to continue. (Click **No** to exit the function.)

Cancel Check Consistency

This section applies to the PV660F storage system

Use this function to manually stop a consistency check that is in progress.

1. Right-click a virtual disk.
2. Click **Cancel Check Consistency**. A confirmation dialog box appears asking whether you want to cancel the consistency check.
3. Click **Yes** to continue. (Click **No** to exit the function.)

Force Online

This section applies to the PV660F storage system

Use this function to force a failed virtual disk to online status. This command is needed to recover a failed virtual disk after one or more failed array disks are replaced.

Note If a failed redundant RAID configured virtual disk is forced online, you must initialize the virtual disk to establish consistency. Data previously on the virtual disk will be lost (since the virtual disk did suffer a failure).

1. Right-click the offline virtual disk.
2. Click **Force Online**. The status change is indicated in the right pane. All virtual disks in the disk group will have a status change from offline to online.

Virtual Disk Properties (PV660F)

This section applies to the PV660F storage system

The **Properties** command for a virtual disk gives the following information on the selected virtual disk:

- **RAID Leve**l — This field displays the RAID level of the virtual disk. Refer to Supported RAID Levels (PV660F) for more information.

- **Target ID** — This field displays the target ID for this virtual disk.
- **Cache Policy** — This field displays the write cache policy assigned to this virtual disk. Refer to Change Policy for more information. This setting is used when the controller is not in Conservative Cache mode.
- **LUN** — This field displays the Logical Unit Number (LUN) assigned to this virtual disk.
- **Physical Size** — This field displays the configured size of the virtual disk.
- **Used Array Capacity** — This field is not applicable at this time and always displays 0.00 MB.
- **Cache Line** — This field is not applicable at this time and is always equal to the Stripe Size.
- **Stripe Size** — This field displays the stripe size of the virtual disk. Refer to Create Virtual Disk (PV660F) for more information.
- **Status** — This field displays the current status of the virtual disk. Possible values are:
  - **Ready** — The virtual disk is operational.
  - **Online** — The virtual disk has been configured and is functioning normally.
  - **Offline** — The virtual disk has not been configured or the configuration has just been cleared.
  - **Failed** — The virtual disk has failed to operate properly.
  - **Rebuild** — The virtual disk has been configured and is in the process of being rebuilt.

Note If the status value is anything other than Ready, a red X appears to the left of the status to indicate a situation that requires attention.

Array Disk Commands (PV660F)

This section applies to the PV660F storage system

Array disks are the physical hard drives connected to the controller through the enclosure. The array disks are configured into disk groups to be configured as one or more virtual disks. You can perform the following operations from the Array Disk context menu:

- Locate a specific array disk — Use the Blink Disk (PV660F) function to blink the LED on the enclosure for the selected array disk.
- Rebuild a failed array disk — Use the Rebuild function to return a degraded virtual disk to online status by rebuilding the redundant data on a hot spare or replacement array disk.
- Cancel a rebuild operation — Use the Cancel Rebuild function to cancel a rebuild that is in progress.
- Prepare an array disk for removal — Use the Prepare to Remove function to spin down the array disk in preparation for removal from the enclosure.
- Define an array disk as a global hot spare disk — Use the Make Global Hot Spare function to assign a selected array disk to act as a backup hot spare.
- Unassign a hot spare disk — Use the Unassign Global Hot Spare function to remove the hot spare status from the selected array disk.
- Force an array disk to Offline state — Use the Make Offline function to force the selected array disk to offline state.
- View information on a selected disk group — Use the Array Disk Properties (PV660F) command.

To bring up an array disk's context menu, right-click the array disk object in the tree view.

Blink Disk (PV660F)

This section applies to the PV660F storage system

The **Blink Disk** command allows you to find an array disk within an enclosure by blinking one of the light-emitting diodes (LEDs) on the disk. You may want to use this command to locate an array disk that has failed.

Use the following steps to blink an array disk:

1. Right-click an array disk.
2. Click **Blink Disk**. The Locate Disk dialog box appears.

3. Click **Blink** to activate the LED for the selected array disk. The LED on the hard-disk carrier may take 30 seconds to start blinking.

4. Click **Done** when finished. The LED may continue blinking for another 30 seconds.

Rebuild

This section applies to the PV660F storage system

The **Rebuild** function starts a rebuild on a replace array disk for a degraded virtual disk. You can do only one rebuild at a time.

1. Right-click an array disk.

2. Click **Rebuild**.

   A warning/confirmation dialog box appears recommending that you not perform any disk or volume operations on the corresponding disks while the rebuild is in progress, and asking whether you want to rebuild this array disk.

3. Click **Yes** to continue. (Click **No** to exit the function.)

The status and progress are indicated in the right pane. The array disk status indicates "Rebuilding."

Note Starting a rebuild will suspend any background initialization occurring on another virtual disk. The background initialization will continue once the rebuild is either completed or canceled.

Cancel Rebuild

This section applies to the PV660F storage system

Use this function to manually stop a rebuild that is in progress.

1. Right-click an array disk.

2. Click **Cancel Rebuild**. A confirmation dialog box appears asking whether you want to cancel the array disk rebuild.

Click **Yes** to continue. (Click **No** to exit the function.)

Prepare to Remove

This section applies to the PV660F storage system

This command spins down the array disk to prepare it for removal from the enclosure. Only array disks that have a hardware failure or that have been brought offline can be prepared for removal.

Note Unused array disks that are not part of a virtual disk or set as hot spares are already spun down and can be removed at any time.

1. Right-click an array disk.
2. Click **Prepare to Remove**.

   A confirmation dialog box appears asking whether you want to prepare this array disk for removal.
3. Click **Yes** to continue. (Click **No** to exit the function.)

   The array disk will spin down so that it can be removed.

Note After doing a Prepare to Remove operation on an array disk in a PV660F array, the array disk cannot be used for a rebuild. If you need to use this disk for a rebuild, then remove and reinsert the disk to automatically start the rebuild.

Make Global Hot Spare

This section applies to the PV660F storage system

Use this function to designate the selected array disk as a global hot spare disk. This function is also used to make a failed / offline array disk usable as a global hot spare.

1. Right-click an array disk.
2. Click **Make Global Hot Spare**. The array disk is labeled as a hot spare in the left pane. The information in the right pane is updated to indicate that the array disk status is Ready and the type is Spare Array Disk.

Unassign Global Hot Spare

This section applies to the PV660F storage system

Use this function to unassign the selected array disk as a global hot spare disk.

1. Right-click an array disk.
2. Click **Unassign Global Hot Spare**. The type change is indicated in the right pane. The array disk status indicates Ready and the type is Array Disk.

Make Offline

This section applies to the PV660F storage system

Use this function to force the selected disk to Offline status.

1. Right-click an array disk.
2. Click **Make Offline**. A warning / confirmation dialog box appears warning that loss of data may occur. The user is asked to confirm this operation.
3. Click **Yes** to continue. (Click **No** to exit the function.) The status change is indicated in the right pane. The array disk status indicates a change from Online to Offline

Note Forcing an array disk that is part of a disk group offline may start a rebuild or may fail one or more virtual disks in the disk group.

Note An array disk that remains in offline state can be brought back online. To do this, you must make it a global hot spare. If desired, it can later be unassigned as a hot spare and used as part of a disk group.

Array Disk Properties (PV660F)

This section applies to the PV660F storage system

The **Properties** command for an array disk gives the following information on the selected array disk:

Disk Properties

- **Vendor** — This field displays the array disk manufacturer or vendor name.
- **Product** — This field displays the array disk product number or model number.
- **Revision** — This field displays the current version of firmware operating on the array disk.

Disk

- **Serial #** — This field displays the array disk serial number
- **Bus Type** — This field displays the bus type supported by this array disk. Currently only fibre array disk drives are supported.
- **Port** — This field displays the drive port to which this array disk is attached.
- **Target ID** — This field displays the target ID for the selected array disk.
- **Enclosure ID** — This field displays the ID of the enclosure where this array disk is located. Be aware that the enclosure ID number displayed by Array Manager does not match the Enclosure ID set on the switch module ID indicator on the back of the PowerVault 660F or 224F enclosures. When the switch module ID indicators on all the enclosures are configured properly, the enclosure ID numbers start at 0 and continue sequentially through 7. The enclosure ID numbers displayed by Array Manager, however, start at 1 and continue sequentially through 8. For this reason, the enclosure ID number displayed by Array Manager will be one greater than the number set on the switch module ID indicator on the back of the PowerVault 660F or 224F enclosures. For example, if the switch module ID indicator identifies the enclosure as 0, the Array Manager enclosure ID will be 1.
- **LUN** — This field displays the Logical Unit Number (LUN) assigned to this array disk.

Size

- **Free Space** — This field indicates the amount of unused space available on the array disk.
- **Used Space** — This field indicates the amount of disk space used or configured.

- **Physical Size** — This field indicates the actual size of the array disk.
- **Largest Contiguous Free Space** — This field indicates the largest amount of free space and is equal to the amount of free space listed above.
- **LUN** — This field displays the LUN ID for the selected disk drive.

Status

- The Status field indicates the array disk's current status. The possible values are:
  - **Ready** — The array disk is operational and is either configured as a hot spare or is a member of a disk group.
  - **Online** — The array disk is powered on, part of a configuration, and functioning normally.
  - **Offline** — The array disk has been inserted, powered on into an unconfigured slot, or the configuration has just been cleared. The array disk is not part of a configuration.
  - **Failed** — The array disk has failed to operate properly.
  - **Rebuild** — The array disk is powered on, part of a configuration, and in the process of being rebuilt.

Note If the status value is anything other than Ready, a red X appears to the left of the status to indicate a situation that requires attention.

Fibre Channel RAID Controller Commands

This section applies to the PV660F storage system

You can perform the following controller operations, which gives you the ability to configure and manage your system:

- Modify controller parameters — Use the Fibre Channel RAID Controller Options function to customize the Dell PowerVault 660F Storage System parameters and features to meet the specific needs of your operating environment.
- Add drive enclosures — Use the Add Enclosure option to add additional PowerVault 224F drive enclosures while the subsystem is still active.
- Change firmware versions — Use the Download function to upgrade the Dell PowerVault 660F Storage System firmware or the enclosure's SES Processor firmware when new versions are available.

- Force the partner controller to be held in reset (i.e., offline partner controller) — Use the Disable Partner function to hold the partner controller in reset.
- Revive the partner controller held in reset — Use the Enable Partner function to return the partner controller to full functionality.
- Perform a reboot of the controllers — Use the Reset function to reboot both controllers of the Dell PowerVault 660F Storage System.
- Prepare the controller for shutdown — Use the Disable BBU function to disable the write cache and prepare the Dell PowerVault 660F Storage System for power down.
- Return the controller to normal operation — Use the Enable BBU function to return the controller to normal cache operation.
- Monitor battery functionality on both controllers — Use the C0 Battery and C1 Battery menu to manage battery functionality.
- To view the properties of the battery — Use the Battery Properties command.
- To view the properties of the controller — Use the Fibre Channel RAID Controller Properties command.

To bring up a controller's context menu, right-click the controller object in the tree view.

Fibre Channel RAID Controller Options

This section applies to the PV660F storage system

Each Dell PowerVault 660F Storage System is shipped from the factory with initial settings that have been found to work well in a majority of applications and environments. These settings are listed as the controller options and vary depending on product and user requirements. You can customize controller options and features to meet the specific needs of your operating environment. The Controller Options dialog box has the following tabs:

- Advanced Fibre Channel RAID Controller Options
- Expert Fibre Channel RAID Controller Options
- Fibre Channel RAID Controller Options

You can also revert any settings you make to the default factory options.

Revert to Factory Controller Options

This section applies to the PV660F storage system

Use the revert function to return to the Dell PowerVault 660F storage system factory default values at any time from any controller option page.

To Revert to Factory Default Values

1. Right-click a controller.
2. Click **Controller Options**. The Controller Options multi-tabbed dialog box appears with the **Advanced** tab open.
3. Click **Revert**. The Revert to Factory Defaults dialog box appears.
4. Select either **Current Page Only** or **All Controller Options**.
5. **Current Page Only** returns the controller options on this page only to the factory default values.

   All Controller Options returns all controller options to the factory default values.
6. Click **OK** to make the changes. (Click **Cancel** to discontinue the operation.)
7. A controller reset is required when making certain controller option changes. Controller options requiring a controller reset are indicated with a fly-by dialog when the mouse cursor moves over the controller option. A controller reset is recommended after making any changes to the controller configuration. If a dialog box does NOT appear with the warning that a controller reset will occur, refer to the Reset command for controller reset instructions.

Advanced Fibre Channel RAID Controller Options

This section applies to the PV660F storage system

To make changes to the Advanced controller options settings:

1. Right-click a controller.
2. Click **Controller Options**.
3. Click the **Advanced** tab to continue. The **Advanced** tab opens.

Note Checking a specific option activates that option, either enabling or disabling the option as stated in the dialog box. No check indicates that the option is disabled.

4. Make changes to the following **Advanced** controller options as required by your operating environment:

   Customizations

   - **Conservative Cache Mode—Default=Enabled.** This function is provided to allow a controller an extra degree of data safety when operating in a degraded state. This function switches write-back caching to write-through operation after a partner controller fails or is disabled. When the partner controller is back online, write-back caching operations resume. Conservative Cache mode is entered automatically when a power supply failure, over-temperature condition, Expand Capacity operation, or SES failure exists.
   - **Restrict Reassign to One Block—Default=Disabled.** When enabled, this function limits reassigning I/O failures to the single failed block. For example, if the I/O process encounters a failure due to a faulty disk block, only the data that was intended for the disk block will be reassigned to another block on the disk. This option is limited to recovered errors and medium errors. If the sense data of the error does not indicate a recovered or medium error, then this setting does not apply. When **Restrict Reassign to One Block** is disabled, the entire I/O process is reassigned when encountering a failure. A controller reset is required when making changes to this controller option.
   - **Smart Large Host Transfers—Default=Disabled.** This function allows selection of Coalesce (fewer disconnects on large transfers) or As Available (more disconnects) on host data transfers. The Smart Large Host Transfers function takes effect for transfers larger than the stripe size.
   - **True Verification Of Data—Default=Disabled.** If this option is disabled, the verify command returns a status without data checking. If this option is enabled, the verify command checks data before returning a status. Enabling this option compromises performance. A controller reset is required when making changes to this controller option.
   - **Write Through Verify—Default=Disabled.** This function enables Force Unit Access for reads and writes during error handling. Force Unit Access bypasses the cache and forces all reads and writes directly to or from the array disk. A controller reset is required when making changes to this controller option. For some devices, enabling Force Unit Access reduces sequential write performance by up to 86%.
   - **Autostart Background Init—Default=Enabled.** *This controller option is only available on version 7.7 of the controller firmware or later.* This function enables you to automatically start a background initialization for a virtual disk . A virtual disk

normally must be foreground initialized after you create it and before using it. The normal initialization process is time consuming (perhaps hours). When using background initialization, you can use the virtual disk immediately without waiting for the initialization process to complete. When you enable the **Autostart Background Init** function, background initialization is automatically initiated for uninitialized virtual disks when the controller receives a write request. The background initialization tasks are queued in the numerical order of the virtual disks. For example, virtual disk 0 is initialized first, followed by virtual disk 1 and so forth. Only one virtual disk can have background initialization active. Background initialization can only be used with virtual disks configured as redundant RAID levels. See the "Background Initialize" section for details on the background initialize function.

Internals

- **Enclosure Management—Default=Enabled.** This function allows the controller to take autonomous actions when a failure occurs. Actions that the Enclosure Management function monitors and reports include array disk failures, background activity completion status, and enclosure events. This function should remain enabled during normal controller operation. Do not disable this function unless specifically instructed to do so as part of a troubleshooting diagnostic activity (such as upgrading the array disk firmware). The Enclosure Management function works in conjunction with Automatic Rebuild Management features in the PowerVault 224F disk array enclosures to detect the removal of a failed array disk. The **Enclosure Management** option must be enabled to allow the PowerVault 660F storage system to exercise its full array of management capabilities.
- **Automatic Rebuild Management—Default=Enabled.** This function allows the controller to take autonomous actions when a failed array disk is replaced or a configured global hot spare disk drive is present. The Automatic Rebuild Management function works in conjunction with Enclosure Management features in the PowerVault 224F disk array enclosures to detect the removal of a failed array disk. The Automatic Rebuild Management function also performs an automatic rebuild after a replacement array disk is installed into a redundant (fault tolerant) virtual disk (RAID-1, RAID-3, RAID-5, and RAID-0+1).
- **Coalesce Device Queues—Default=Enabled.** This function provides device queuing coalescing optimization. The function enables data traffic coalescing (combining of address adjacent I/Os) on the traffic of each device. This joins the data from adjacent I/Os into a single I/O to improve performance.
- **Queue Limit—Default=6.** This option specifies the allowed queue depth for tagged commands to all attached array disks. This value is further limited to the array disk's own tag limit. Allowed values are 1 to 230. A setting of 1 is similar to no tags. Tagged command queuing is an optional SCSI-II command feature that

enables multiple “tagged commands” to be sent to the array disk for processing. The commands are executed in a sequence that requires the least amount of rotational latency from the disk, thus improving performance.

5. When the **Advanced** controller options are set appropriately, click **OK** to make the changes. (Click **Revert** to revert controller options to factory default values. Click **Cancel** to discontinue the operation.)

Expert Fibre Channel RAID Controller Options

This section applies to the PV660F storage system

To make changes to the **Expert** Fibre Channel RAID controller options:

1. Right-click a controller.
2. Click **Controller Options**.
3. Click the **Expert** tab to continue. The **Expert** window opens.

Note Checking a specific option either enables or disables the option as stated in the dialog box. No check indicates that the option is disabled.

4. Make changes to the following **Expert** controller options as required by your operating environment:
   - **On Q Full Give Busy—Default=Disabled**. Enabling this option sets the controller to return a Busy status when a Queue Full condition is detected. Disabling this option sets the controller to return a Queue Full status. When a command is received and the controller detects a Queue Full condition, it normally returns Queue Full status. This option is intended to help hosts that are confused by Queue Full. A controller reset is required when making changes to this controller option.
   - **Disable Busy Status During Failback—Default=Disabled**. The **Disable Busy Status During Failback** function allows the controller to disregard new requests without returning a Busy status. If enabled, during failback, the surviving controller ignores all new requests and does not return any status. If disabled, the surviving controller returns a Busy status to new commands received from the host during the cache flush operation. A controller reset is required when making changes to this controller option.

- **Vendor Unique Test Unit Ready Status—Default=Disabled.** This function enables a vendor unique TUR response sent to an offline LUN. If disabled, a Hard Error status is returned. If enabled, a Not Ready status is returned. A controller reset is required when making changes to this controller option.
- **No Pause On Controller Not Ready—Default=Disabled.** This option turns the pause off or on for certain commands when the controller is not ready. Normally, when the controller is starting up, certain commands encounter a pause. This happens only when the controller has not reached Startup Complete. The pause lasts one second. The commands affected are: Prefetch, Read / write, Read / write Extended, Test Unit Ready, Verify, and Write Verify. A controller reset is required when making changes to this controller option.
- **Disable CC For Invalid LUN—Default=Enabled**. When this option is enabled, the controller disables Check Condition for an invalid LUN. This affects the handling of the Inquiry command when the referenced LUN is invalid. If enabled, the Inquiry command returns data with the peripheral qualifier indicating that the peripheral device is not connected. If disabled, the Inquiry command will be failed with a check condition of "Illegal Request, LUN Not Supported." A controller reset is required when making changes to this controller option.
- **Auto Restore (automatic failback)—Default=Enabled.** When enabled in a redundant controller system, **Auto Restore** allows automatic recovery of a partner controller when a replacement is inserted. A controller reset is required when making changes to this controller option.
- **Reset Propagation—Default=Disabled.** This function allows a port to issue an internal reset without causing a reset event to occur on its attached interface. If enabled, a port that issues an internal reset propagates the reset by causing a reset event to occur on its attached interface. If disabled, a port will not cause a reset event on its attached interface as part of issuing an internal reset.
- **Multi-port Reset—Default=Disabled.** This function restricts internal resets to ports that have logical devices reserved through that port. If enabled, an internal reset is executed by a port only if a logical device has been reserved through that port. If disabled, internal resets are not qualified by logical device reservations.
- **Rebuild Rate—Default=50.** Change the default Rebuild Rate to less than or equal to 50. A rate of 50 devotes the maximum allowable resources to a drive rebuild or array expansion, allowing the Rebuild or Expand to proceed at its fastest. Lowering the number devotes more resources to I / Os and consequently slows the Rebuild or Expand process.

5. When the **Expert** controller options are set appropriately, click **OK** to make the changes. (Click **Revert** to revert controller options to factory default values. Click **Cancel** to discontinue the operation.)

Fibre Channel RAID Controller Options

This section applies to the PV660F storage system

To make changes to the Fibre Channel RAID controller options:

1. Right-click a controller.
2. Click **Controller Options**.
3. Click the **Fibre** tab to continue. The **Fibre** tab opens.

Note Checking a specific option either enables or disables the option as stated in the dialog box. No check indicates that the option is not enabled.

4. Make changes to the following **Fibre** controller options as required by your product and configuration:
   - **Node Name Retention—Default=Enabled.** This option disables/enables a failed controller's node or system name to be retained through a controller failure. When disabled, each controller shares its node name with its partner controller through failover; however, when failback occurs, the replacement controller uses its own node name. When enabled, each controller shares its node name with its partner controller and those names are used through all phases of failover and failback. A controller reset is required when making changes to this controller option.
   - **PCI Latency Control—Default=Short.** This function allows adjustment of the Fibre Channel chip's use of the PCI bus. This function controls the amount of data each processor can burst across the primary bus before relinquishing bus ownership to the next device. PCI Latency Control takes effect only when both ports are active and are arbitrating for the bus. PCI Latency Control allows the integrator to tune the controller's operating options for specific applications. Possible settings are Long, Medium, and Short; and these can be selected from the drop-down menu. For maximum throughput, Long is recommended and is equivalent to the time necessary to transfer 1024 bytes; Medium is equivalent to 684 bytes, and Short is equivalent to 512 bytes. A controller reset is required when making changes to this controller option.
   - **Frame Control—Default=2KB.** This option sets the host Fibre Channel data frame size. Possible settings are 512 bytes, 1 KB or 2 KB; and these can be selected from the drop-down menu. A frame size of 2 KB is recommended since it provides the largest packet transfers and the most throughput. A controller reset is required when making changes to this controller option.

- **Set Hard Loop IDs—Default is the following: Controller 0 Port 0=Enabled, Loop ID=2; Controller 1 Port 0=Enabled, Loop ID=4.** This option allows you to enable or disable use of the controller / host port and, if enabled, to define the hard loop ID. Specifying a hard loop ID means that the same ID will always be requested. The valid range for loop IDs is from 0 to 125. A controller reset is required when making changes to this controller option.

5. When the **Fibre** controller options are set appropriately, click **OK** to make the changes. (Click **Revert** to revert controller options to factory default values. Click **Cancel** to discontinue the operation.)

6. A controller reset is required when making some controller option changes. Controller options requiring a controller reset are indicated with a pop-up window when the mouse cursor moves over the controller option. A controller reset is recommended after making any changes to the controller configuration. If a dialog box does NOT appear warning that a controller reset will occur, refer to the Reset command for the controller reset instructions.

Add Enclosure

This section applies to the PV660F storage system

The **Add Enclosure** command allows users to add device enclosures (one or more) to a configured system while the system continues to operate. After the enclosure or enclosures have been added to the system, the **Add Enclosure** command initiates the SES monitoring process for the added enclosure(s). The user can configure the additional disk drive capacity without rebooting the PowerVault 660F storage system through Array Manager.

To Add Enclosures

1. Attach and cable the new enclosure(s) to the system while the new enclosure(s) are powered off. Be sure to check for and resolve any enclosure shelf ID conflicts. Refer to the appropriate documentation for enclosure installation.

2. Power on the new enclosure(s).

3. Wait for all the drives to complete the spin-up process.

4. Right-click the corresponding controller object in the Array Manager console's left pane.

5. Select **Add Enclosure**.

6. Right-click the PV660 SubSystem object and select **Rescan**.

Download

This section applies to the PV660F storage system

The **Download** command allows you to download controller firmware or SES firmware. To use the command, right-click the controller object. Then select the **Download** command from the context menu that comes up. A submenu appears for selecting either of the two download options. For details, see:

- Fibre Channel RAID Controller Firmware, the subsection directly below.
- SES Firmware, the next subsection.

Note Upgrading from PowerVault 660F Storage System firmware version 7.51 to version 7.7 or higher will change the controller's World Wide Name (WWN). When the WWN changes, you will not be able to see the storage system with Array Manager. Resolving this problem can include updating pathing of the FC HBA, as well as FC switch zoning configuration if zoning is being performed by WWN. See the documentation accompanying the new controller firmware, and the latest SAN or PowerVault 660F documentation that came with your system or that is available from the Dell support document library at http://support.dell.com for information on editing the FC switch zone table and resolving problems associated with a new WWN.

Fibre Channel RAID Controller Firmware

This section applies to the PV660F storage system

Note Before downloading firmware, be sure to suspend all I/O (including Virtual Disk initializations, check consistency, rebuilds, Expand Capacity, and Add Virtual Disk) to the PowerVault 660F storage system.

1. Select **Controller Firmware.**

   The **Download Firmware** dialog box appears. The dialog box provides a standard browser for locating the appropriate firmware image file.

2. Click **Browse** to locate the file and enter the path and filename. Information about the selected file displays.

3. Verify that the selected file is correct. If a valid file is not selected, the command is not executed.

4. Click **Apply** to begin the download process. (Click **Cancel** to discontinue the operation.)

Note If downloading to a redundant controller system, both controllers will receive the firmware simultaneously.

If downloading controller firmware version 7.7 or later onto a storage system running version 7.51, you will have to manually power down and back up the storage system before the new firmware will take affect. Additional steps to configure the SAN for the newly downloaded storage system will be required -- see the Note above for more information.

SES Firmware

This section applies to the PV660F storage system

Use this function to upgrade enclosure SES firmware as new versions become available. Array Manager allows you to download new SES firmware to the SES device. You will be able to download to all SES devices associated with the PowerVault 660F array. The SES Download wizard guides you through the process.

Note When downloading SES firmware, make sure that the SES firmware for all enclosures attached to a PV660F storage subsystem are at the same level.

1. Select **SES Firmware.**

2. Read the welcome screen and click **Next** to continue. (Click **Cancel** to discontinue the operation.)

   The **Select Enclosures** screen appears.

3. Select the enclosures for which you want to download SES firmware. Click **Next** to continue. (Click **Back** to return to the previous screen. Click **Cancel** to discontinue the operation.)

4. Click **Browse** to locate the file and enter the path and filename.

5. Verify that the selected file is correct. If the correct file is not selected, the command is not executed.

6. The Completing the SES Download Wizard screen appears. Click **Next** to continue. (Click **Back** to return to the previous screen. Click **Cancel** to discontinue the operation.)

7. Click **Finish** to continue. (Click **Back** to return to the previous screen. Click **Cancel** to discontinue the operation.)

While the SES download is ongoing, the SES Download Wizard will show the progress of the download. If you open another window over the Wizard window, it will temporary erase the existing progress information. Once the Wizard window is updated with new information, all existing progress information will be redisplayed.

Note After successfully downloading the SES firmware, you may receive a warning stating that the firmware is mismatched. Rescan the PV660F subsystem to resolve this warning.

Disable Partner

This section applies to the PV660F storage system

Use this function to hold the partner controller in reset and initiate the failover process. If you are planning to use the **Expand Capacity** or **Add Virtual Disk** commands, those commands are not available unless you use this command first to hold the partner controller in reset and transfer control to a single controller. After the **Expand Capacity** or **Add Virtual Disk** command is completed, you need to use the **Enable Partner** command to return to a redundant controller configuration with both controllers operating.

Note Disabling the partner controller forces all I/O to the subsystem to go through the surviving controller. Any servers connected to the Fibre Channel fabric may experience an HBA failover (depending on the fabric configuration) in order to redirect I/O to the surviving controller.

To use the Disable Partner command:

1. Right-click a controller.

2. Click **Disable Partner**. This command is invoked immediately; no warning or confirmation screen appears. The partner controller is held in reset. The surviving controller operates as a failed-over redundant controller.

Enable Partner

This section applies to the PV660F storage system

Use this function to return the partner controller to full functionality. Enable the partner controller following an **Expand Capacity** or **Add Virtual Disk** operation to return the system to a redundant controller configuration.

Note Enabling the partner controller allows I/O to be processed by both controllers. For a SAN environment, manual HBA failovers may be necessary to have servers use the newly enabled controller.

To use the Enable Partner command:

1. Right-click the active controller.
2. Click **Enable Partner**. This command is invoked immediately; no warning or confirmation screen appears. The failback process begins, and the partner controller resumes full functionality.

Note If a virtual disk is being initialized in the background, the Enable Partner function will fail. You must either wait until background initialization is complete or canceled before successfully performing an Enable Partner.

Note If after performing the enable partner command, the controller object Properties command only shows one controller, you will have to perform a rescan on the PV660F SubSystem object in the Array Manager tree view to have Array Manager recognize both controllers. This most often occurs when Array Manager is started while the storage system had its partner controller disabled.

Reset

This section applies to the PV660F storage system

Use this function to reboot the controllers in the storage system. After creating a new configuration or making changes to controller options, a controller reset is recommended.

1. Right-click a controller.
2. Click **Reset**. A warning/confirmation dialog box appears notifying the user that a controller reset will take 30 seconds or longer. The user is asked to confirm this operation.

3. Click **Yes** to continue. (Click **No** to exit the function.) A warm power cycle occurs immediately on the controllers.

Note A reset can be done only when the following two conditions both exist: 1) There is no I/O to or from the controllers, and 2) no initializations, rebuilds, check consistency, Expand Capacity, Add Virtual Disk are in progress.

Disable BBU

This section applies to the PV660F storage system

Use this feature to prepare the storage system for power down. This procedure forces the controller into conservative cache mode, flushes the cache, and electronically disconnects the battery.

1. Right-click a controller.
2. Click **Disable BBU**. A warning/confirmation dialog box appears notifying you that the controller will be prepared for removal. You are asked to confirm this operation.
3. Click **Yes** to continue. (Click **No** to exit the function.)

Enable BBU

This section applies to the PV660F storage system

Use this function to return the controller to normal cache operation. This enables the BBU to protect the write-cache. If the array is in an optimal state, this operation sets the controller to normal cache mode.

1. Right-click on a controller object.
2. Click **Enable BBU**. A warning/confirmation dialog box appears notifying the user that the controller will return to normal operation. The user is asked to confirm this operation.
3. Click **Yes** to continue. (Click **No** to exit the function.)

C0 Battery and C1 Battery

This section applies to the PV660F storage system

The C0 Battery represents the battery for the C0 controller (right side on the rear view of the storage system), and the C1 Battery represents the battery for the C1 controller (left side on the rear view).

Use these menu items to manage the following battery functions:

- Initiate a battery recondition sequence — Use the **Recondition** function to perform a full battery conditioning cycle.
- Cancel a battery recondition cycle — Use the **Cancel Recondition** function to cancel a recondition that is in progress.
- Change the battery threshold settings — Use the **Change Low Threshold** function to modify the threshold limits.

Recondition

This section applies to the PV660F storage system

Certain conditions make it beneficial for the BBU battery to be reconditioned before it can be used. If you are installing a new controller or replace the battery, then the battery should be reconditioned. This will fully validate the maximum battery charge possibly.

A full battery recondition cycle consists of fully discharging and recharging the battery. The recondition cycle must start with a fully charged battery. When the battery is first connected, it will run a calibration phase which will identify the battery charge.

The recondition cycle must be initiated manually and must complete without an interruption of power. In addition, the operator must not interrupt the cycle while it is in process. If the cycle is interrupted for any reason, the batteries must be recharged and the recondition process must be restarted.

Note During the reconditioning, write back cache mode will be disabled as the battery charge falls below the battery threshold. This may slow performance of the storage system. Therefore, schedule this operation accordingly.

To recondition a battery on a PV660F controller:

1. Right-click a controller.
2. Click **C0 Battery** or **C1 Battery** for the battery associated with the selected controller.

3. Click **Recondition**. The recondition cycle begins immediately and must not be interrupted.

For more information about the BBU, see the following:

- Battery Properties
- How the Battery Backup Unit Works
- How to Detect, Recondition, and Replace a Failed Battery Backup Unit

Cancel Recondition

This section applies to the PV660F storage system

To cancel a battery recondition on a PV660F controller:

1. Right-click a controller.
2. Click **C0 Battery** or **C1 Battery** for the battery associated with the selected controller.
3. Click **Cancel Recondition**. The recondition cycle is cancelled immediately.

Change Low Threshold

This section applies to the PV660F storage system

Use this function to modify the power level used to trigger an alarm when the battery power level drops below the set value. Once the battery's power level falls below this threshold, the controllers will enter conservative cache mode (i.e., write cache disabled).

1. Right-click a controller.
2. Click **C0 Battery** or **C1 Battery** for the battery associated with the selected controller.
3. Click **Change Low Threshold**. The Change BBU Low Threshold dialog box appears. The current value in minutes is shown.
4. Enter the new value. Click **OK** to make the changes. (Click **Cancel** to discontinue the operation.)

Battery Properties

This section applies to the PV660F storage system

This command includes the following battery properties information:

General Characteristics

- **Battery Type** — The battery pack type. Currently, NiMH is the only battery pack type supported.
- **Hardware Revision** — Version number for the controller. Version numbers are from 1 to 255. The version number used for the first release is 1 for the controller production release.

Power

- **Maximum Power** — The highest level that the battery can be charged, expressed as hours or minutes of charge. This value will change as the battery ages and is reconditioned. The older the battery gets, the less maximum charge it can contain.
- **Current Power** — The current power level of the battery, expressed as hours or minutes of charge. This value will change as the battery is discharged or charged.
- **% Charge Level** — The battery's charge level, where 100% means the battery has been charged up to its maximum level.
- **Low Threshold** — This is a preset or modifiable power level used to trigger an event when the battery power level drops below the set value. When the battery falls below this threshold, the controller will switch over to Conservative Cache Mode. Refer to Change Low Threshold for information on setting this value.
- **Threshold Unit** — The threshold level is expressed in hours or minutes.

Attributes

This is an informational field describing the current attribute(s), for example, Fast Charge is not active or Recondition is not active.

States

This field provides a status icon and information about the current state of the battery. A blue check indicates good status, and a red X indicates a problem or poor status.

- **Reconditioning** — This field monitors the reconditioning status and indicates whether the battery has or has not been reconditioned within 30 discharge / charge cycles.

- **Reconditioned** — This field indicates whether or not the battery has been fully discharged before it was charged up. A new battery indicates Never Reconditioned until it has been fully discharged and then charged by the Recondition procedure.
- **Power** — This field indicates the current power value.

Status

This field displays the current status of the battery. Possible values are:

- **Ready** — The battery is functioning normally.
- **Degraded** — The battery has suffered a failure of a component and is operating in a degraded state.
- **Failed** — The battery has suffered a failure and is no longer functioning.

Note If the status value is anything other than Ready, a red X appears to the left of the status to indicate a situation that requires attention.

Fibre Channel RAID Controller Properties

This section applies to the PV660F storage system

This command displays the following categories of properties for one or both controllers in the Dell PowerVault 660F storage system:

General Fibre Channel RAID Controller Characteristics

- **Model** — This field displays the controller number and model that is present in the selected slot.
- **Firmware Version** — This field displays the current version of firmware operating on the controller.
- **WWName** — This field displays the node or system World Wide Name (WWN) for this controller.
- **Cache Size** — This field displays the amount of cache memory available on this controller.
- **Memory Size** — This field displays the amount of memory available on this controller.
- **Host Ports** — This field displays the number of host ports available on this controller.
- **Drive Channels** — This field displays the number of device ports or drive channels available on this controller.

- **Logical Drives** — This field displays the number of logical drives or virtual drives created on this controller.
- **Physical Devices** — This field displays the number of physical devices attached to this controller.

Dual Controller Status

- **Fault Tolerant** — This field indicates whether the PV660F is operating in a fault tolerant or dual-active controller configuration.
  - **Yes** indicates dual-active controller configuration. This means that the PV660F storage subsystem is operating with two controllers and is therefore fault tolerant. In order for the server to have fault-tolerant access to the PV660F storage subsystem, the server must have two fibre channel host bus adapters (HBAs), each having a fibre channel connection to each of the controllers in the subsystem.
  - **No** indicates either single controller configuration or failed-over dual-active controller configuration.
- **Partner Status** — This field indicates the partner controller's status. Possible values are:
  - **OK** — The partner controller is functioning normally.
  - **Failed** — The partner controller is no longer functioning, and the surviving controller is operating in a failed-over state.

Host Connection Information

- **OS Controller** — This field displays the controller's identification number assigned by the operating system. This value is based on the number of peripheral devices recognized by the operating system, including controllers and HBAs.
- **Channel** — This field displays the channel number on which this controller is located.
- **Target** — This field displays the target ID for this controller.

Port

- **Name** — This field displays the port name for this controller.
- **Port ID** — This field displays the port ID for this controller.

Status

- This field displays the current status of the controller. Possible values are:
  - **Ready** — The controller is functioning normally.

- **Degraded** — The controller has suffered a failure of a component and is operating in a degraded state.
- **Failed** — The controller has suffered a failure of more than one component and is no longer functioning.

Note If the status value is anything other than Ready, a red X appears to the left of the status to indicate a situation that requires attention.

Subsystem Commands (PV660F)

This section applies to the PV660F storage system

The PV660F SubSystem folder object consists of physical and logical arrays. The physical arrays include objects for the RAID controllers, enclosures, disks, fans, power supplies, and temperature probes. The logical arrays include objects for the virtual disks on the storage system. The logical arrays are created, monitored, and maintained by Array Manager.

Right-clicking a PV660F SubSystem folder object in the tree view brings up a context menu with subsystem commands. Similarly, right-clicking a physical or logical array object also brings up a context menu.

You can perform the following operations from the Subsystem context menu:

- Scan for configuration changes, or new Dell PowerVault 660F storage systems — Use the Rescan Subsystem command.
- View information on a selected subsystem — Use the Subsystem Properties (PV660F) command.

The physical array object context menu allows you to:

- View general information on the selected physical array — Use the Physical Array Properties (PV660F) command.

The logical array object context menu allows you to:

- View information on the selected logical array — Use the Logical Array Properties (PV660F) command.

Note Commands related to the enclosure, fans, power supply, and enclosure temperature are covered in the next major section, "Enclosure Management for the PV660F Storage System."

Rescan Subsystem

This section applies to the PV660F storage system

You can scan for new PowerVault 660F Storage Systems or hardware configuration changes to existing storage systems by using the **Rescan** function.

1. Right-click a subsystem.
2. Click **Rescan**. This command is invoked immediately; no warning or confirmation screen appears. The added components appear in the left pane.

Subsystem Properties (PV660F)

This section applies to the PV660F storage system

The **Properties** command for the PV660F SubSystem object provides the name, vendor, and status of the subsystem. Possible values for status are:

- **Ready** — The subsystem is functioning normally.
- **Degraded** — The subsystem has suffered a failure of a component and is operating in a degraded state.
- **Failed** — The subsystem has suffered a failure of more than one component and is no longer functioning.

Note If the status value is anything other than Ready, a red X appears to the left of the status to indicate a situation that requires attention.

Physical Array Properties (PV660F)

This section applies to the PV660F storage system

The Physical Array Properties dialog box displays the name, vendor, and status of the physical array. Possible values of the status field are:

- **Ready** — The logical array is functioning normally.
- **Degraded** — The logical array has suffered a failure of a component and is operating in a degraded state.
- **Failed** — The logical array has suffered a failure of more than one component and is no longer functioning.

Note If the status value is anything other than Ready, a red X appears to the left of the status to indicate a situation that requires attention.

Logical Array Properties (PV660F)

This section applies to the PV660F storage system

The Logical Array Properties dialog box displays the name and status of the logical array. Possible values of the status field are:

- **Ready** — The logical array is functioning normally.
- **Degraded** — The logical array has suffered a failure of a component and is operating in a degraded state.
- **Failed** — The logical array has suffered a failure of more than one component and is no longer functioning.

Note If the status value is anything other than Ready, a red X appears to the left of the status to indicate a situation that requires attention.

Enclosure Management for the PV660F Storage System

This section applies to the PV660F storage system

Dell PowerVault 224F enclosures have enclosure management capability. This section describes how enclosure management works and then details the specific enclosure management commands. The specific topics in the section are:

- Enclosure Management Overview
- Enclosure Commands (PV660F)

Enclosure Management Overview

This section applies to the PV660F storage system

Array Manager allows you to view all the external Fibre Channel system probes within the same integrated Array Manager console. You are notified of any enclosure status changes while the console is open, and logged events appear in the Events tab as well as in the Windows Event Log.

The Fibre Channel subsystem of Array Manager provides direct support for the display of instrumentation of Fibre Channel enclosures. No other software is required to be installed.

The storage objects associated with enclosures and enclosure management are:

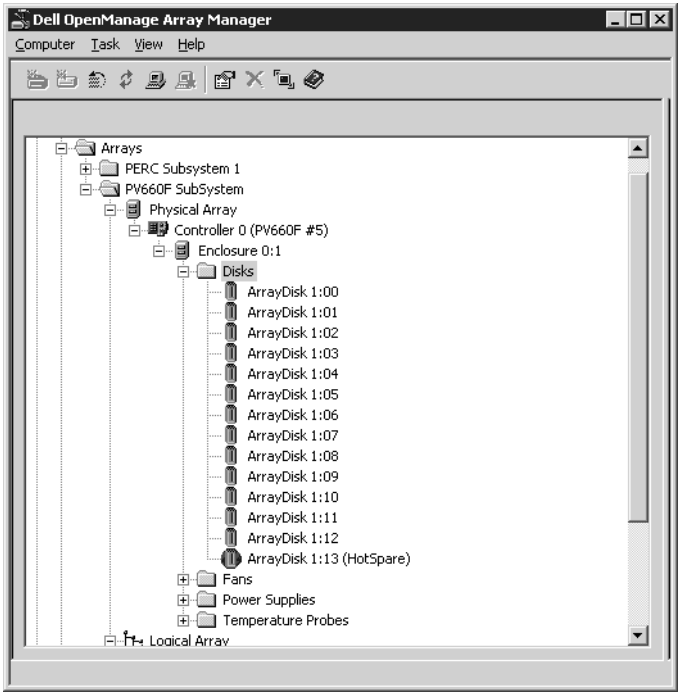
- Enclosures
- Disks
- Fans
- Power Supplies
- Temperature Probes
- LS Modules (SES)

Enclosures

This section applies to the PV660F storage system

When you expand the PV660F SubSystem storage object in the tree view, you will see one or more controllers attached to the subsystem along with one or more enclosures that are attached to each controller. As shown in the sample screen below, an enclosure has Disks, Fans, Power Supplies and Temperature Probes storage objects associated with it.

Enclosure Management Storage Objects



In the above screen, the enclosure is identified as “Enclosure 0:1.” The zero refers to the number of the controller the enclosure is attached to. The “1” means that it is the first enclosure attached to that controller. If there were a second and a third enclosure attached to Controller 0, they would be called “Enclosure 0:2” and “Enclosure 0:3,” respectively. If you right-click an Enclosure storage object, you will get a context menu with enclosure commands. For details, see the section “Enclosure Commands (PV660F)” in this chapter.

The subordinate Enclosure storage objects—Disks, Fans, Power Supplies, and Temperature Probes—are described in the sections that follow.

Disks

This section applies to the PV660F storage system

The enclosure’s array disks are displayed under the Disks folder object. When the Disks folder object is highlighted in the left-pane tree view, the right-pane panel of the console window lists these disks along with their operating status and other properties. If you

right-click on one of the subordinate array disks object of the Disks folder object, the Array Disks context menu will come up. See "Array Disk Commands (PV660F)" for details on these commands.

Fans

This section applies to the PV660F storage system

The enclosure's fans are displayed under the Fans folder object. When the Fans folder object is highlighted in the left-pane tree view, the right-pane panel of the console window lists the enclosure's fans along with their operating status and other properties. If you right-click on one of the Fans object, a context menu comes up with a Fan Properties command.

Power Supplies

This section applies to the PV660F storage system

The power supplies that are installed in the enclosure are displayed under the Power Supplies folder object. When the Power Supplies folder object is highlighted in the left-pane tree view, the right-pane panel of the console window lists the enclosure's power supplies along with their operating status and other properties. If you right-click on any of the Power Supplies object, a context menu comes up with a Power Supply Properties command.

Temperature Probes

This section applies to the PV660F storage system

The temperature probes that are installed in the enclosure are displayed under the Temperature Probes folder object. When the Temperature Probes folder object is highlighted in the left-pane view, all of the enclosure's temperature probes are displayed along with their operating status and other properties. The categories are Status, Current Value, Minimum Warning Limit, Maximum Warning Limit, Minimum Error Limit, Maximum Error Limit, and Unit of Measurement. The unit of measurement is degrees Celsius. If you right-click on any of the Temperature Probes objects, a context menu comes up with two commands, Temperature Properties and Set Thresholds for Temperature. The **Set Thresholds for Temperature** command displays the enclosure temperature and the thresholds for critical and warning temperature events. It allows you to set the temperature threshold for warning temperature levels.

Note Only the warning threshold of the LS Modules (SES) temperature probes can be set, and setting the warning threshold of either LS Module will result in the warning thresholds of both LS Modules to be set to the new value. The warning thresholds of the power supply temperature probes are set to a fixed value.

LS Modules (SES)

This section applies to the PV660F storage system

LS (or SES) modules are cards installed in slots at the front of the enclosure. Each LS module has an SES processor that monitors environmental functions. LS modules also have a Loop Redundancy Circuit (LRC) function, which maintains the viability of the Fibre Channel loop. LS modules are not displayed in the Array manager tree view.

Note When downloading firmware to the LS or SES module, make sure the firmware for all modules and all enclosures attached to a PV660F storage subsystem are at the same level. In addition, it is recommended that the firmware for all enclosures in a SAN be at the same level.

For more information, see "SES Firmware."

Enclosure Commands (PV660F)

This section applies to the PV660F storage system

This section describes the commands associated with the Enclosure storage object and its subordinate storage objects. Right-clicking on the enclosure object in the tree view of an Dell PowerVault 224F array disk enclosure brings up a context menu with enclosure commands. Under each enclosure are objects in the tree view for fans, power supplies, and temperature. Right-clicking each of these subordinate objects brings up a context menu as well.

The Enclosure context menu commands are:

- Alarm Enabled — Enables an audible alarm that sounds whenever the fault LED lights.
- Alarm Disabled — Turns off the audible alarm settings.
- Asset Settings — Allows you to enter or change asset information for the enclosure.
- Enclosure Properties (PV660F) — Displays enclosure properties.

The context menu commands for fans, power supplies, and temperature are:

- Fan Properties — Displays fan properties.
- Power Supply Properties — Displays power supply properties.
- Set Thresholds for Temperature — Displays the enclosure temperature and the thresholds for critical and warning temperature levels. Allows the warning threshold to be set for the LS Modules (SES) temperature probes.
- Temperature Properties — Displays temperature-related properties.

Alarm Enabled

This section applies to the PV660F storage system

Use this function to enable the alarm when the enclosure exceeds a warning threshold for temperature; when a disk drive, power supply, fan, or controller fails; and during the controller boot-up process. Once the alarm is enabled, the audible alarm sounds whenever the fault LED lights. To invoke the command, right-click the enclosure and select **Alarm Enabled** from the context menu that appears.

Alarm Disabled

This section applies to the PV660F storage system

Use this function to disable the alarm. To invoke the command, right-click the enclosure and select **Alarm Disabled** from the context menu that appears.

Asset Settings

This section applies to the PV660F storage system

The Asset Settings command brings up a dialog box that allows you to enter or change the asset tag and asset name for the array disk enclosure. Each enclosure of the PowerVault 660F Storage System can have its own asset settings. To perform the command:

1. Right-click on an enclosure object.
2. Click **Asset Settings**.

   The Asset Information dialog box appears with two entry boxes, Asset Tag and Asset Name.
3. If the values shown need to be changed, enter a new value for the asset tag and a new asset name.

The asset tag for the selected enclosure will be added to the controller name in the object tree -- for example, **Controller 1 (Enclos53)**, where "Enclos53" is the asset tag for the first enclosure attached to Controller 1. The first enclosure attached to a controller is the enclosure in which the pair of redundant controllers are installed. The asset tag is also used to identify the PowerVault 660F array in the Dell OpenManage Storage Consolidation software (if the array is used in a Storage Area Network (SAN).

Note The asset tag is limited to 10 characters, and the asset name is limited to 32 characters.

4. Click **Set** to make the changes. Click **Cancel** to exit the function without accepting the changes.

Note If you have recently added an enclosure to the PV660F storage system, then the changes you make to the asset tag may not be applied even after doing a rescan. If this is the case, first rescan the controller and then change the asset tag.

Enclosure Properties (PV660F)

This section applies to the PV660F storage system

Use this function to view the following information about the selected enclosure:

- **Vendor** — This field displays the enclosure vendor's name.
- **Product** — This field displays the enclosure product number or model number.
- **Revision** — This field displays the current version of the LS Modules (SES) (or SES processor) firmware.
- **Enclosure ID** — This field displays the enclosure ID number, beginning with 1. Be aware that the enclosure ID number displayed by Array Manager does not match the Enclosure ID set on the switch module ID indicator on the back of the PowerVault 660F or 224F enclosures. When the switch module ID indicators on all the enclosures are configured properly, the enclosure ID numbers start at 0 and continue sequentially through 7. The enclosure ID numbers displayed by Array Manager, however, start at 1 and continue sequentially through 8. For this reason, the enclosure ID number displayed by Array Manager will be one greater than the number set on the switch module ID indicator on the back of the PowerVault 660F or 224F enclosures. For example, if the switch module ID indicator identifies the enclosure as 0, the Array Manager enclosure ID will be 1.
- **Asset Tag** — This field displays the enclosure asset number that is determined by the **Asset Settings** command.

- **Asset Name** — This field displays the enclosure asset name that is determined by the **Asset Settings** command.
- **Service Tag** — This field displays the enclosure service tag that is established during manufacturing.
- **Status** — This field displays the current status of the enclosure. Possible values are:
  - **Ready** — The enclosure is functioning normally.
  - **Degraded** — The enclosure has suffered a failure of a component and is operating in a degraded state.
  - **Failed** — The enclosure has suffered a failure off more than one component and is no longer functioning.

Note If the status value is anything other than Ready, a red X appears to the left of the status to indicate a situation that requires attention.

Fan Properties

This section applies to the PV660F storage system

Use this command to view the following fan properties:

- **Vendor** — This field displays the fan vendor or manufacturer name.
- **Product** — This field displays the fan number.
- **Enclosure ID** — This field indicates the ID number of the enclosure where this fan is located. Be aware that the enclosure ID number displayed by Array Manager does not match the Enclosure ID set on the switch module ID indicator on the back of the PowerVault 660F or 224F enclosures. When the switch module ID indicators on all the enclosures are configured properly, the enclosure ID numbers start at 0 and continue sequentially through 7. The enclosure ID numbers displayed by Array Manager, however, start at 1 and continue sequentially through 8. For this reason, the enclosure ID number displayed by Array Manager will be one greater than the number set on the switch module ID indicator on the back of the PowerVault 660F or 224F enclosures. For example, if the switch module ID indicator identifies the enclosure as 0, the Array Manager enclosure ID will be 1.
- **Fan Speed** — This field displays the current fan speed. Possible values are High, Low, and N/A (meaning the fan speed is not available).
- **Controller** — This field indicates the ID number of the controller with which this fan is associated.
- **Status** — This field displays the current status for the fan. Possible values are:

- **Ready** — The fan is operating normally.
- **Degraded** — The enclosure has suffered a failure of a fan component and is operating in a degraded state.
- **Failed** — The fan has failed to operate properly.

Note If the status value is anything other than Ready, a red X appears to the left of the status to indicate a situation that requires attention.

Power Supply Properties

This section applies to the PV660F storage system

Use this command to view the following power supply properties:

- **Vendor** — This field displays the power supply vendor or manufacturer name.
- **Product** — This field displays the power supply number.
- **Part No** — This field displays the part number and part revision number of the power supply.
- **Enclosure ID** — This field indicates the ID number of the enclosure where this power supply is located. Be aware that the enclosure ID number displayed by Array Manager does not match the Enclosure ID set on the switch module ID indicator on the back of the PowerVault 660F or 224F enclosures. When the switch module ID indicators on all the enclosures are configured properly, the enclosure ID numbers start at 0 and continue sequentially through 7. The enclosure ID numbers displayed by Array Manager, however, start at 1 and continue sequentially through 8. For this reason, the enclosure ID number displayed by Array Manager will be one greater than the number set on the switch module ID indicator on the back of the PowerVault 660F or 224F enclosures. For example, if the switch module ID indicator identifies the enclosure as 0, the Array Manager enclosure ID will be 1.
- **Status** — This field displays the current status for the power supply. Possible values are:
  - **Ready** — The power supply is operating normally.
  - **Failed** — The power supply has failed to operate properly.

Note If the status value is anything other than Ready, a red X appears to the left of the status to indicate a situation that requires attention.

Set Thresholds for Temperature

This section applies to the PV660F storage system

Note This command can only be performed on the LS Modules (SES) temperature probes.

This command displays the enclosure temperature and the thresholds for critical and warning temperature levels. The warning thresholds can be changed. All values are expressed in degrees Celsius. To perform the command:

1. Right-click on a LS module temperature probe object in the tree view.
2. Select **Set Thresholds for Temperature** from the context menu that appears.

   A dialog box comes up that displays the following temperature-related information:

   - **Current Temperature** — This is the current reading for the selected temperature probe.
   - **Critical Threshold** — These values are set in the array disk enclosure and cannot be changed.

     High or 55 degrees Celsius: This value is the high critical limit. If this value is exceeded, the system sound the audible alarm (if enabled).

     Low or 1 degree Celsius: This value is the low critical limit. If temperature drops below this value, the system will sound the audible alarm (if enabled).
   - **Warning Thresholds** — You can set these thresholds. The warning threshold is a general setting that will apply to the temperature probes on the LS modules of the enclosure. The temperature probes on the power supplies also have warning thresholds, but these can not be changed.

     High: When this value is exceeded, a warning alarm sounds, if the alarm is turned on. The upper limit for the high warning threshold is 35 degrees Celsius.

     Low: When the temperature drops below this value, a warning alarm sounds, if the alarm is turned on. The lower limit is 20 degrees Celsius.
3. Make any necessary changes to the Warning Threshold values.

Note Thresholds are set based on room temperature expectations. The temperatures reported by the subsystem are actual values. The subsystem automatically adjusts the entered thresholds to account for internal temperature limitations.

4. Click **OK** to make the changes. Click **Cancel** to discontinue the operation.

Temperature Properties

This section applies to the PV660F storage system

Use this command to view the following temperature-related properties:

Note All values are expressed in degrees Celsius.

- **Vendor** — This field displays the temperature probe vendor or manufacturer name.
- **Product** — This field displays the temperature probe number.
- **Enclosure ID** — This field indicates the ID number of the enclosure where this temperature probe is located. Be aware that the enclosure ID number displayed by Array Manager does not match the Enclosure ID set on the switch module ID indicator on the back of the PowerVault 660F or 224F enclosures. When the switch module ID indicators on all the enclosures are configured properly, the enclosure ID numbers start at 0 and continue sequentially through 7. The enclosure ID numbers displayed by Array Manager, however, start at 1 and continue sequentially through 8. For this reason, the enclosure ID number displayed by Array Manager will be one greater than the number set on the switch module ID indicator on the back of the PowerVault 660F or 224F enclosures. For example, if the switch module ID indicator identifies the enclosure as 0, the Array Manager enclosure ID will be 1.
- **Current Temperature** — This is the current reading for the selected temperature probe.
- **Critical Thresholds** — Thresholds set by the program:
  - **High** or 55 degrees Celsius: This value is the high critical limit. If this value is exceeded, the audible alarm will sound (if enabled).
  - **Low** or 1 degree Celsius: This value is the low critical limit. If the temperature value drops below this value, the audible alarm will sound (if enabled).
- **Warning Thresholds** — Thresholds the user can adjust with the Set Thresholds for Temperature command:
  - **High**: When this value is exceeded, a warning alarm sounds, if the alarm is turned on.
  - **Low**: When this value is exceeded, a warning alarm sounds, if the alarm is turned on.
- **Status** — This field displays the current status of the temperature probe. Possible values are:
  - **Ready** — The temperature probe is functioning normally.

- **Degraded** — The temperature probe has suffered a failure of a component and is operating in a degraded state.
- **Failed** — A temperature probe is not present.

Note If the status value is anything other than Ready, a red X appears to the left of the status to indicate a situation that requires attention.

RAID Policies (PV660F)

This section applies to the PV660F storage system

Today's storage management tools used for configuring and managing storage subsystems require the system administrator to have a great deal of knowledge about RAID technology. To simplify storage management, the PV660F storage subsystem enables you to use RAID Policies, which are automated RAID configuration methods that are application-based.

Overview

This section applies to the PV660F storage system

RAID Policies provide the following functions:

- A set of standard RAID Policies based on commonly used server applications, such as mail servers and database servers.
- A means to create, save, and edit a custom RAID Policy.
- A set of generic RAID Policies based on redundancy and performance.
- A means to execute a given RAID Policy on different array groups.

If RAID-based server hardware is used to set up an application server (Exchange, Notes, SQL, or Oracle), then Array Manager's RAID Policies make this task simple.

Using Standard RAID Policies

This section applies to the PV660F storage system

Most common mail server and database server applications require a specialized storage configuration. Some of the application server software components require a highly redundant storage configuration; others require a high-performance storage configuration. Array Manager standard RAID Policies are designed to meet the storage configuration needs of the following software:

- Microsoft® Exchange
- Lotus Notes
- Microsoft SQL Server
- Oracle

To Use a Standard RAID Policy

1. Right-click an array group and select **Apply RAID Policy** to launch the RAID Policy wizard.

2. Read the welcome screen. Click **Next** to display the Select RAID Policy window or cancel the operation by clicking **Cancel**.

3. Select a RAID Policy by highlighting an option under **Mail Server Policies**, **Database Server Policies,** or **Generic Policies**.

4. Review the storage configuration for the selected RAID Policy. You can modify the policy attributes in the right pane of the Select RAID Policy window or use the defaults.

   When ready click **Next** to display the Select Disks window. The Select Disks window displays the selected disks.

5. Modify the disk selection as necessary by using the **Add Disk**, **Remove Disk,** or **Remove All** buttons. Click **Next** to continue.

6. Click **Finish** to execute the selected RAID Policy.

Since processing a RAID Policy involves several steps, such as in the example of creating virtual disks, it might take several minutes to complete the task. The RAID Policy wizard automatically creates the appropriate virtual disks.

Using Custom RAID Policies

This section applies to the PV660F storage system

To Create a Custom RAID Policy

1. Right-click an array group and select **Edit RAID Policy** to launch the RAID Policy wizard.
2. Click the **New Policy** button to create a new policy.
3. Enter a name for the custom RAID policy in the **Policy Name** field. Use the **Add Disk** and **Remove Disk** buttons to add and remove disks. Click the **OK** button.
4. After creating a custom RAID Policy, use **Apply RAID Policy** to execute the policy.

To Modify a Custom RAID Policy

1. Right-click an array group and select **Edit RAID Policy** to launch the RAID Policy wizard.
2. Click the **Modify Policy** button to modify an existing custom policy.

To Delete a Custom RAID Policy

1. Right-click an array group and select **Edit RAID Policy** to launch the RAID Policy wizard.
2. Click the **Delete** button to delete an existing custom policy.

NetWare Support for the PowerVault 660F

After installation, PowerVault 660F management by NetWare is disabled by default. You can enable NetWare support for a PowerVault 660F from a Windows system by using the Array Manager Utilities. When using the Array Manager Utilities, you will select the **NetWare** tab and then enable the PV660F component.

For detailed information on enabling the PV660F component for NetWare, see "Adding and Removing Fibre Channel RAID and Enclosure Management."

For information on SAN support for NetWare, see the Dell PowerVault Systems Storage Area Network (SAN) documentation.

How the Battery Backup Unit Works

This section applies to the PV660F storage system

The battery backup unit (BBU) monitors and charges the cache backup battery for the PowerVault 660F. In order for the BBU to properly monitor and charge the cache backup battery, it first automatically performs a battery calibration and then Array Manager logs an event stating that a BBU recondition is suggested. The steps in both of these processes are identical except that the calibration process takes place automatically whereas you start the recondition process manually. Both the calibration and recondition processes consist of a battery discharge followed by a battery charge.

The BBU was designed to compare the low / high battery charges between the calibration and recondition cycles for accurate reporting. It performs the calibration cycle automatically. Dell OpenManage Array Manager logs an event stating that a BBU recondition is suggested. Upon completion of the recondition cycle, the BBU compares the capacity values for the battery. If the capacity values are within a nominal range, the recondition process is complete. However, if the capacity values exceed this nominal range, Array Manager will log another event message stating that a BBU recondition is suggested so that it can determine if the battery is operating properly.

The calibration cycle allows the BBU to determine the low and current capacity values for the battery and is required for proper operation of the write cache. The recondition cycle allows for more accurate capacity value reporting and is not required for normal write cache operation, but is recommended to verify battery capacity.

These are the events that require a battery calibration and subsequent recondition:

- Powering on the PowerVault 660F for the first time
- Reinstalling or replacing the fan pack (which houses the battery)
- Reinstalling or replacing a redundant array of independent disks (RAID) controller

Note Removing and reinserting a fan pack while the power is turned off is equivalent to installing a new battery.

If you power off or reset the PowerVault 660F during a calibration, the process restarts automatically.

If you power off or reset during a recondition, Array Manager displays an alert in the **Events** tab that lets you know the recondition has failed. Next, Array Manager logs an event stating that a BBU recondition is suggested.

It is recommended that you perform both a calibration and recondition because the BBU uses these two sets of low / high battery charge values to accurately determine the battery's actual charge status. Until then, the BBU relies on the battery manufacturer's specifications, which may give inaccurate information about the actual battery installed in the storage subsystem.

Array Manager indicates the maximum capacity value based on battery manufacturer specifications. After both the calibration and reconditioning cycles are complete, Array Manager uses this information to show actual capacity.

For more information, see the following sections:

- "Battery Calibration" for a description of the calibration cycle
- "Battery Recondition" for a description of the recondition cycle
- "Recondition" for the procedure describing how to manually initiate the recondition cycle
- "Indicators for Battery Calibration and Recondition" for information on the messages and other indicators that provide the status of the calibration and recondition cycles

Battery Calibration

This section applies to the PV660F storage system

Typically, the battery shipped with a new PowerVault 660F Storage System will not be charged to 100 percent. In other words, the first battery discharge phase during calibration could easily take far less time than the maximum 19 hours. You can expect the fast charge time to take six hours.

It is recommended that you allow sufficient time when you first start the PowerVault 660F during a storage area network (SAN) site installation. Do not interrupt the first two phases of the PowerVault 660F BBU calibration cycle (the battery discharge and fast battery charge). These two phases combined require no interruption for seven to 25 hours. If you must interrupt the process, the calibration cycle will restart from the beginning, but you have not caused damage to the device. Interruptions increase the time needed to get the PowerVault 660F Storage System's batteries running at an optimum 100 percent charge.

Note Pulling the controller or fan pack during the calibration cycle interrupts that cycle. You will know when the fast charge process finishes because Array Manager will display a message asking you to recondition the battery.

For related information, see the following sections:

- "Recondition" for the procedure describing how to manually initiate the recondition cycle

- "Battery Recondition" for a description of the recondition cycle
- "Indicators for Battery Calibration and Recondition" for information on the messages and other indicators that provide the status of the calibration and recondition cycles
- "How the Battery Backup Unit Works" for a description of the battery backup unit

Battery Recondition

This section applies to the PV660F storage system

Keep in mind, the BBU Recondition is Suggested message is informational and is not an error message.

You can perform the recondition cycle during your next scheduled maintenance. After you do this, Array Manager will be able to give you a calculated maximum power value based on information from both the calibration and recondition cycles. However, it is not a requirement for you to perform the reconditioning immediately after the calibration.

Dell OpenManage Array Manager will notify you when it is time to perform another recondition. It does this after the battery level drops 30 times below an 85 percent charge. When you receive this message, you can plan another recondition at your next scheduled maintenance. Reconditioning is not required based on elapsed time. Rather, it is based on the number of times the battery discharges below 85% as described above.

For related information, see the following sections:

- "Recondition" for the procedure describing how to manually initiate the recondition cycle
- "Battery Calibration" for a description of the calibration cycle
- "Indicators for Battery Calibration and Recondition" for information on the messages and other indicators that provide the status of the calibration and recondition cycles
- "How the Battery Backup Unit Works" for a description of the battery backup unit

Indicators for Battery Calibration and Recondition

This section applies to the PV660F storage system

This section lists the Array Manager indicators that you can observe during the calibration and reconditioning cycles.

1 - Calibration Battery Discharge Phase

- Phase starts with the PowerVault 660F storage system in conservative cache mode
- `Battery recondition has started` message may be seen in the **Events** window
- Power level displays `zero(0)` in the **Battery Properties** window
- `Recondition is Active` message displays in the **Battery Properties** window

2 - Calibration Battery Fast Charge Phase

- Power level displays current value in the **Battery Properties** window
- `Fast charge is active` message displays in the **Battery Properties** window
- `Recondition is Active` message still displays in the **Attributes** section of the **Battery Properties** window
- `Controller entered normal cache mode` message displays in the **Events** window when the power rises above the **Low Threshold** setting

3 - Recondition Battery Discharge Phase

- `Recondition Is Active` message appears in the **Battery Properties** window
- `Battery Recondition has Started` message displays in the **Events** window
- Battery power displays current power; over time this value will decrease
- `Controller entered Conservative Cache Mode` message displays in the **Events** window when the battery level drops below the Low Threshold setting

4 - Recondition Battery Fast Charge Phase

- First part of fast charge is in conservative cache mode until battery charge rises above low threshold
- `Been Reconditioned` message displays in the Status section of the **Battery Properties** window
- `Fast-charge is Active` message displays in the Attributes section of the **Battery Properties** window
- `Recondition is Active` message displays in the Attributes section of the **Battery Properties** window **Battery Charge** in the **Battery Properties** window displays actual value
- `Controller entered normal cache mode` message displays in the **Events** window when the power rises above the Low Threshold setting described earlier

For related information, see the following sections:

- "Recondition" for the procedure describing how to manually initiate the recondition cycle
- "Battery Calibration" for a description of the calibration cycle
- "Indicators for Battery Calibration and Recondition" for information on the messages and other indicators that provide the status of the calibration and recondition cycles
- "How the Battery Backup Unit Works" for a description of the battery backup unit

Remote Management 19

This chapter describes Array Manager's remote connection capabilities. With Array Manager, you can connect to remote computers and manage their storage remotely.

Understanding the Conditions for Remote Connection

In order to connect to a remote computer, you have to gain access to the computer by having an account that matches an administrator account on that computer. The connection procedure for establishing your authentication varies depending on the type of remote computer you are connecting to. However, in all cases when you attempt to connect to a remote computer, the Connect to Computer dialog box will come up.

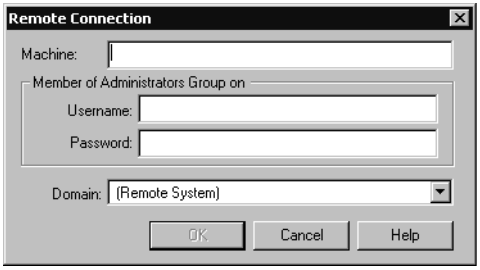
Connect to Computer Dialog Box



This dialog box allows you to identify the remote computer and gives you some options that make it more useful to connect to the computer in the future.

After you click the button in the Connect to Computer dialog box, you may get a second dialog box, the Remote Connection dialog box, which requires you to enter a user name and password and, when necessary, a domain.

Remote Connection Dialog Box



Depending on the type of computer you want to connect to and, possibly, how you have logged in to the local computer, this dialog box may not always come up. Sometimes you will be connected directly to the remote computer without having to enter any authentication information. At other times, this dialog box will not come up and you will be denied access.

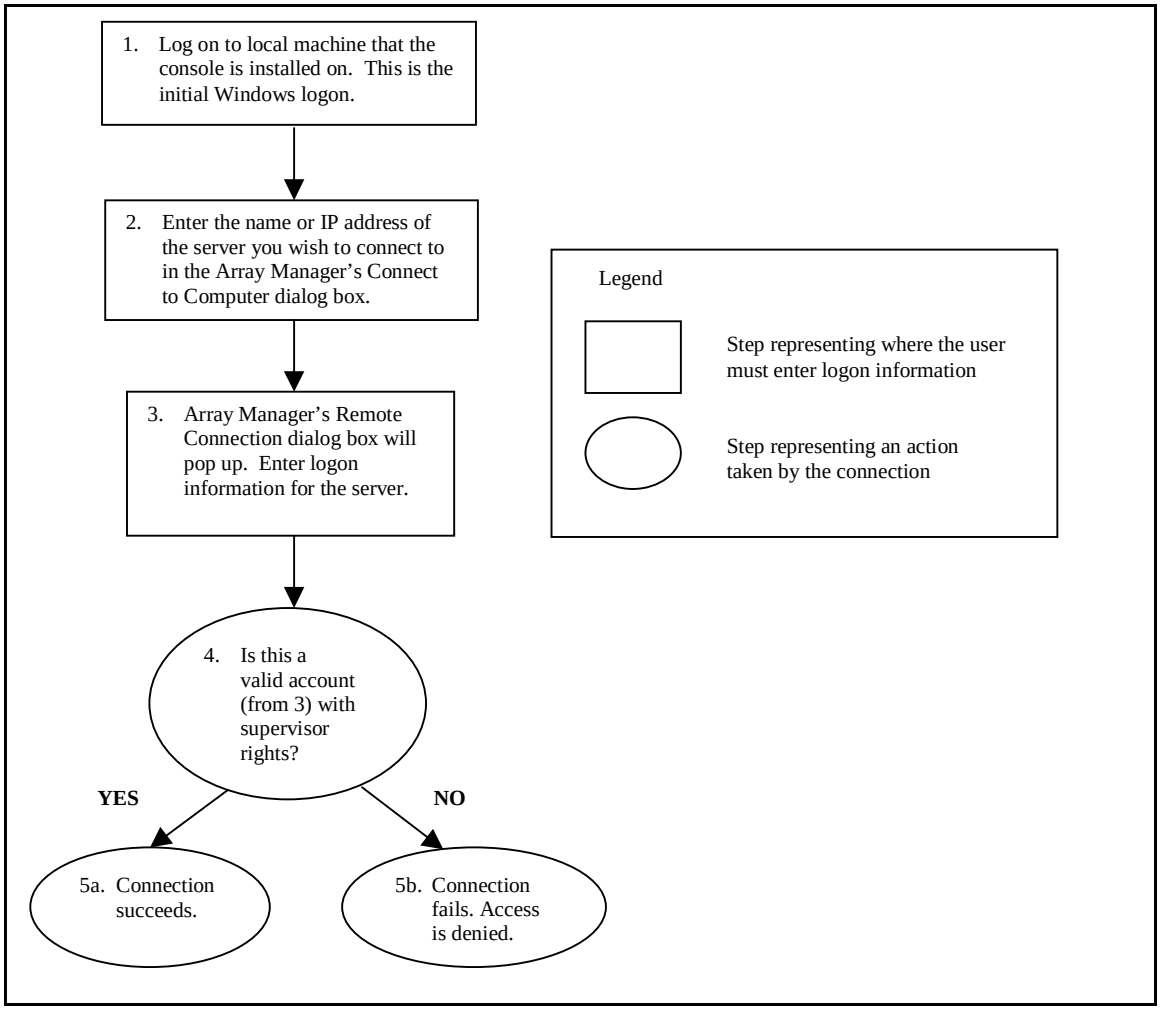
It is helpful to know that there are three different connection situations that are categorized by the type of remote computer:

- NetWare® computer. See "Remote Connection Sequence for NetWare Computers" for more information.
- Windows computer with Array Manager installed. See "Remote Connection Sequence for Windows Computers" for more information.
- Windows 2000 Disk Management computer. See "Remote Connection Sequence for Windows 2000 Server with Disk Management" for more information.

Remote Connection Sequence for NetWare Computers

The simplest connection method is for a NetWare computer. When you attempt to connect to a NetWare computer, a Remote Connection dialog box will always come up to enable you to enter a user name and password. If you enter the user name and password for an account with administrator rights on the remote computer, you will be able to connect.

The flow chart below diagrams the NetWare computer connection method:



Note If you are unable to connect to a NetWare server, see the section "You are unable to connect to a NetWare server" in the "Troubleshooting" chapter.

Note Dell does not offer NetWare in Japan.

Remote Connection Sequence for Windows Computers

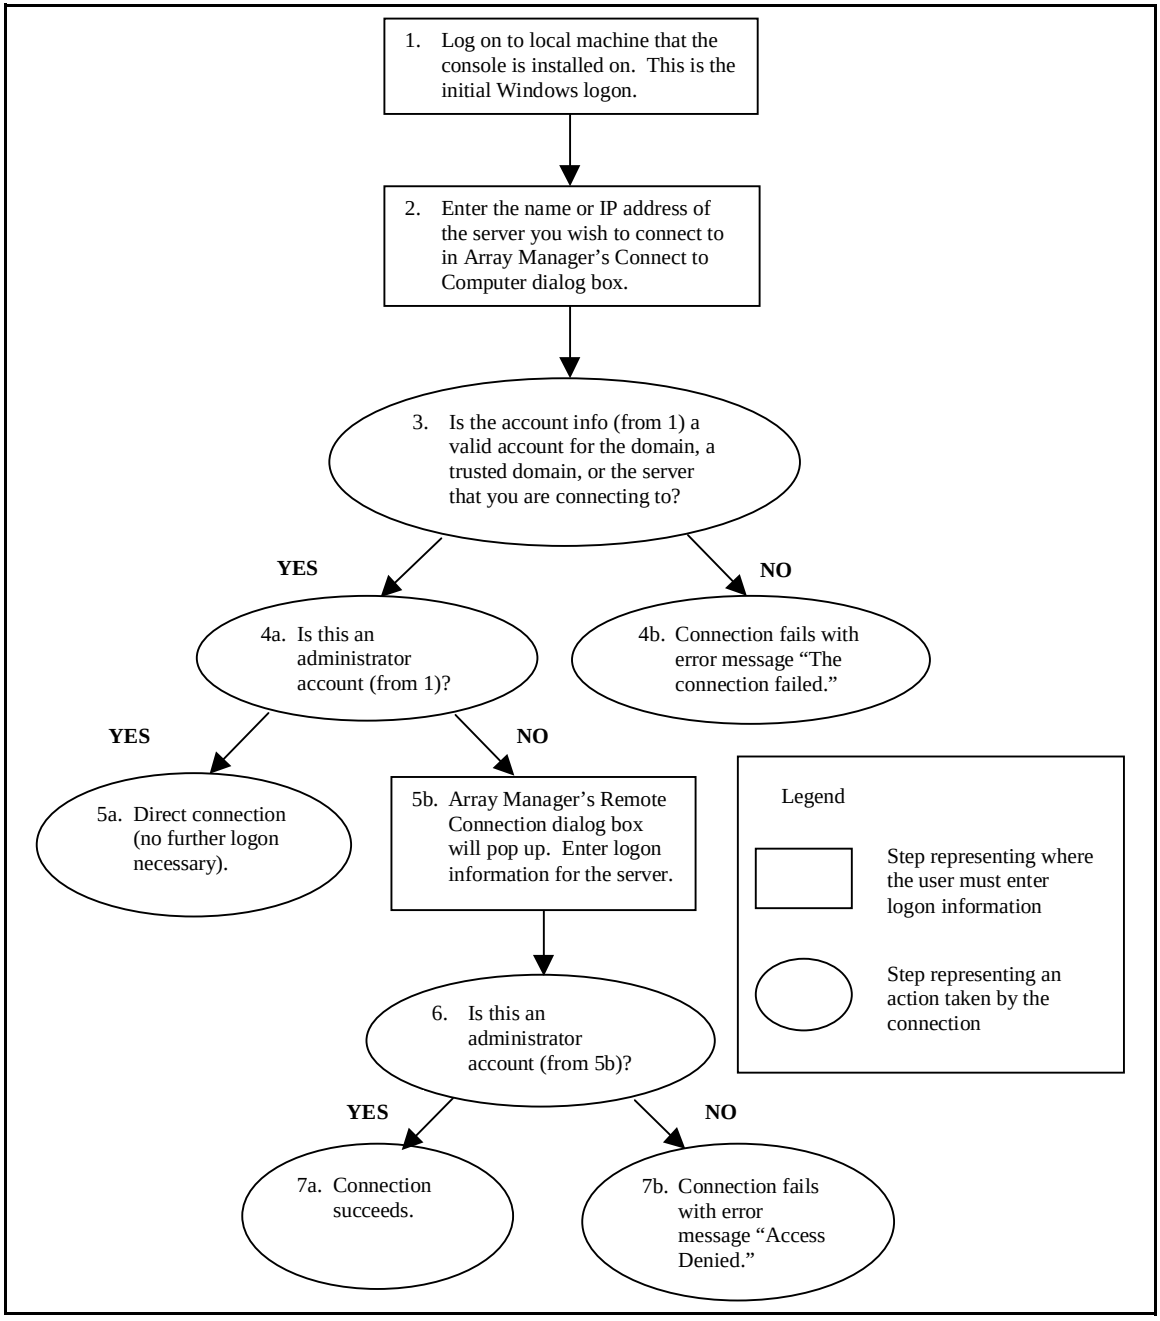
If you are attempting to connect a Windows computer with Array Manager installed, the initial login to the local computer determines whether or not you will receive the Remote Connection dialog box for entering authentication information.

The different ways you can log in initially on the local computer and their results when you attempt to connect to the remote computer are summarized as follows:

	Initial login at the local computer	Result when you try to connect to remote computer
1	You log in at the local computer with an account that matches an administrator account on the remote computer. It can be a: - Local administrator account - Domain administrator account of the remote computer's domain or a trusted domain	You are connected directly to the computer without the Remote Connection dialog box.
2	You log in at the local computer with a local user or domain user account that matches a valid account on the remote computer that is NOT in the administrator group.	The Remote Connection dialog box comes up. To connect successfully to the remote computer, you have to enter a user name and password that matches a valid administrator account on the remote computer. If you are entering a domain administrator account, you will need to indicate the name of the domain as well.
3	You log in at the local computer with an account that does NOT match any valid account on the remote computer.	The Remote Connection dialog box does not appear. Access is denied per Windows security model.

In the third case in the table, you are denied access because your initial login does not show any relationship to the remote computer. The workaround for the third case in the table is to set up an account on the local computer that matches an administrator account on the remote computer and to log in to the local computer with that account. If the local or remote computer is in a workgroup, the third case probably will apply.

The flow chart that follows diagrams the possible connections for connecting to a Windows 2000 computer with Array Manager installed:

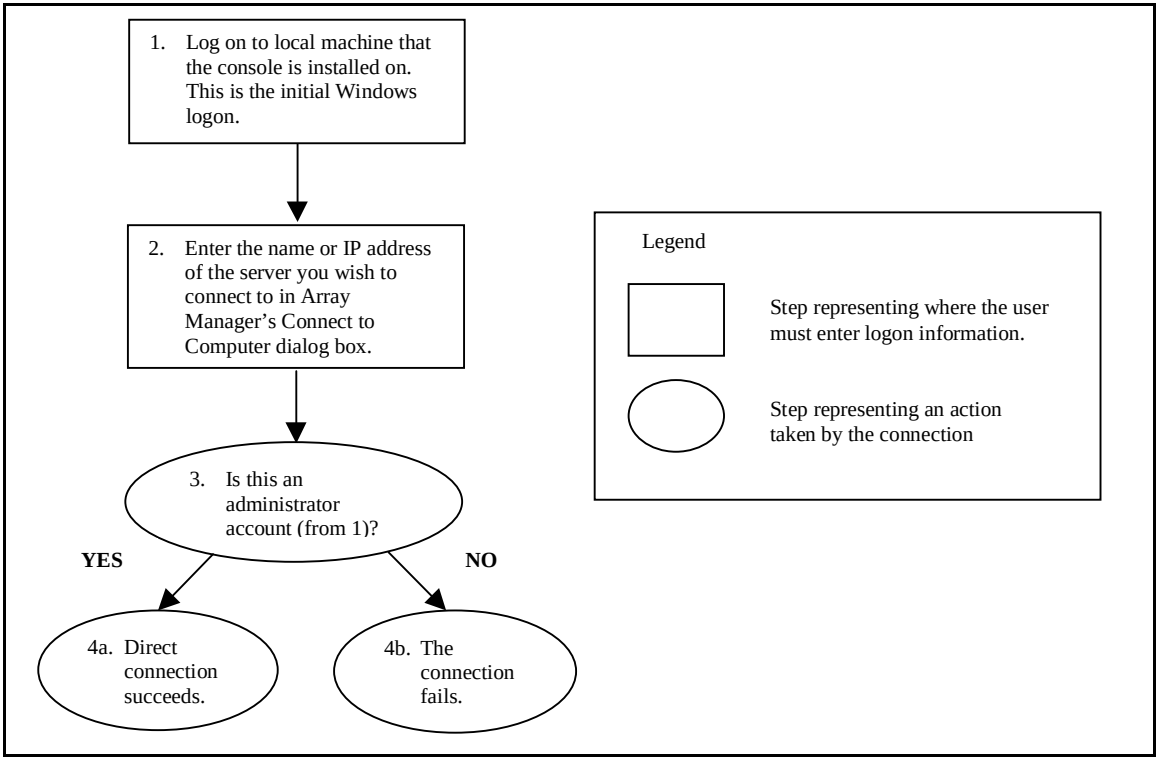

1. Log on to local machine that the console is installed on. This is the initial Windows logon.
2. Enter the name or IP address of the server you wish to connect to in Array Manager's Connect to Computer dialog box.
3. Is the account info (from 1) a valid account for the domain, a trusted domain, or the server that you are connecting to?
YES
NO
4a. Is this an administrator account (from 1)?
4b. Connection fails with error message "The connection failed."
YES
NO
5a. Direct connection (no further logon necessary).
5b. Array Manager's Remote Connection dialog box will pop up. Enter logon information for the server.
6. Is this an administrator account (from 5b)?
YES
NO
7a. Connection succeeds.
7b. Connection fails with error message "Access Denied."
Legend
Step representing where the user must enter logon information
Step representing an action taken by the connection

Remote Connection Sequence for Windows 2000 Server with Disk Management

Disk Management is the built-in disk and volume manager that comes with Windows 2000. Disk Management is also known as Logical Disk Manager, or LDM. Because Disk Management and Array Manager for Windows 2000 are related programs, you can connect to a Windows 2000 computer with Disk Management and manage its storage remotely from a computer with the Array Manager console software.

The situation for connecting to a remote Windows 2000 computer with Disk Management is fairly simple. You will never get the Remote Connection dialog box. Your authentication credentials are established at the initial login to the local computer. Thus, you must log in to the local computer with an account that matches an administrator account on the Windows 2000 remote computer with Disk Management.

The flow chart below diagrams the connection process:



If you try to connect to a Windows 2000 computer with Disk Management right after a console-only Array Manager install, you may not be able to connect. See the section "You are unable to connect to a Windows 2000 server with Disk Management after a client-only installation" in the "Troubleshooting" chapter.

Remote Connection and Windows Server 2003

When connecting to a remote Windows Server 2003 system, you must log into the remote system using an account that has administrator privileges. By default, Windows Server 2003 does not allow unauthenticated (null) connections to access the SAM user accounts. Therefore, if you are attempting to connect using an account that has a blank or "null" password, the connection may fail.

Further Considerations on Remote Connections

This section contains the following additional topics on remote connection:

- Direct Connections to Windows Systems Based on Cached Information
- Reconnecting to a Favorite Server at Startup
- "Connection Failed" or Other Error Messages
- Workgroups and Domains

Direct Connections to Windows Systems Based on Cached Information

If you connect to a remote Windows computer during an Array Manager session, disconnect from it, and then attempt to reconnect to it in the same session, Array Manager will automatically connect to it without a connection login box. The reason for this direct connection is that Array Manager keeps authentication information for all connections made during the current session. Thus, if you have already established your credentials for connecting to a computer in that session, you will not have to give them again.

If you attempt to connect to a Windows remote computer that has the same administrator user name and password as another remote computer that you connected to in the current session, you will be automatically connected without a connection login box.

If in the current Array Manager session you establish a Windows computer as a "favorite" and then disconnect from it, you can reconnect directly to it again in the current session by simply double-clicking its name in the Favorites list in the tree view. The same applies to a computer in the History list in the tree view. If you connected to a remote computer in the current session and disconnected from it, you can reconnect to it in the current session by double-clicking its name in the History list.

Reconnecting to a Favorite Server at Startup

The question is often asked about a "favorite computer" that you have designated to be reconnected when Array Manager restarts. Will that computer come up automatically in the tree view or will you get a connection login box? It depends on your initial login to the local computer.

Note Initial connection to the Array Manager server is required for connection to servers checked with "Reconnect at Next Start of Application".

In every case, Array Manager will attempt to connect to the remote computer. If your initial login establishes that you have administrator rights on the remote computer, the remote computer will appear in your tree view. However, if your initial login requires a connection login box to establish your credentials as an administrator on the remote computer, then the login connection box will appear.

For more on Favorites, see "Using Favorites" in this chapter.

"Connection Failed" or Other Error Messages

If you get a "Connection Failed" or other error message, see the section "When attempting to access a remote computer, you are denied access or get an error message" in the "Troubleshooting" chapter.

Workgroups and Domains

Windows 2000 has two types of networking organization—workgroups and domains.

The workgroup model implements a peer-to-peer network where each computer within the workgroup both accesses resources of the other computers within the workgroup and provides resources to be used by other computers in the workgroup. There is no centralized authentication within this model, so each computer must maintain its own list of users. Thus, membership within a workgroup is loosely defined.

The domain model implements centralized authentication and strictly defined membership. Instead of each server defining its own set of users, the set of users is defined only on the domain server.

Array Manager allows consoles in either model to connect to server agents. In the domain model, the user running the Array Manager console has a domain account that is a valid account on the server running the Array Manager agent.

In the workgroup model, the user running the Array Manager console has a valid local account with administrative rights on the system running the Array Manager agent. Array Manager does not restrict access by workgroup membership. In fact, you can use

the workgroup model when you want to connect between a computer in a workgroup and one in a domain. This also applies in a situation where the console and the agent are on computers in two different domains.

Array Manager System Discovery

From the Array Manager console, you can view all the computers within any connected Windows domain or in the administrator's primary Novell® Directory Services (NDS) tree that have the Array Manager software installed on them.

In Array Manager, left-click **My Network Places** and then **Microsoft Windows Network** to display all the domains and their computers attached to the network. You must log on to a computer before you can manage storage devices connected to that computer.

NetWare computers managed by Array Manager are also displayed in **My Network Places**. In Array Manager, left-click **My Network Places**, then **NetWare Services**, and then **NetWare Servers** to display all NetWare computers running Array Manager that have been found on the network.

Note It is possible that **My Network Places** will display fewer NetWare computers than are actually running Array Manager on the network. **My Network Places** may also display NetWare computers that have no past or present association with Array Manager. These conditions are due to limitations of the discovery process on NetWare.

You can connect to computers as described in "Connecting to a Remote System from the Console," "Connecting to a Remote Windows System from a Command Line," or "Connecting to NetWare Servers."

Connecting to a Remote System from the Console

From the Array Manager console, you can view all the servers within any connected Windows domain that have the Array Manager software installed on them. You need administrator-level access rights to install and use Array Manager.

You can also connect to a remote Windows or NetWare system using the Computer pull-down menu or My Network Places.

If you have not done so previously, you may also wish to view the "Understanding the Conditions for Remote Connection" section before continuing with the following procedure.

Logon Requirements for Connection to a Remote Windows System

You can remotely log on to a Windows system from Array Manager if you meet the following logon conditions:

- A remote computer with the local Administrator rights of the remote computer system.
- A remote computer in a domain with the domain administrator rights.
- A remote computer in a different domain, which has a trusted relationship, with the domain administrator rights.

Logon Requirements for Connection to a Remote NetWare System

You can remotely log on to a NetWare Array Manager server with these conditions:

- A remote computer to which you have administrator rights.

Connecting to a Remote System with the Array Manager Console

Use the following procedure to connect to a remote computer from the Array Manager console:

1. Select **Connect** from the **Computer** pull-down menu or right-click the **My Network Places** icon and select **Connect**.

   The Connect to Computer dialog box is displayed.



2. Type the name of the computer to connect to.

   For a NetWare computer, you must type in the DNS host name or IP address.

   The second check box is inactive (dimmed) unless the **Add to Favorites folder** check box is checked. When you add a remote computer to Favorites, Array Manager places the computer in the Favorites folder and will remember the logon information for accessing that computer as long as the current session of Array Manager is running.

   If you select **Add to Favorites folder**, you can also select **Reconnect at next start of application** so that Array Manager will attempt to connect to the remote computer whenever the program is started.

For more information on Favorites, see "Using Favorites" later in this chapter.

Note If when you attempt to connect to a remote computer, you get a "Connection Failed" or other error message, see the section "When attempting to access a remote computer, you are denied access or get an error message" in the "Troubleshooting" chapter.

3. Click **Connect**. The system will now connect you to the remote computer if you have already met the security logon conditions — that is, your logon name and password match the user name and password of an administrator account on the remote computer. The connection may take several minutes to complete. In this case, you are now finished with this procedure.

   If you have not met the security logon conditions, the Remote Connection dialog box is displayed. In this case, continue with the next step of this procedure.

   To understand more about why the Remote Connection dialog box is sometimes displayed and other times not, see the following sections:

   - "Understanding the Conditions for Remote Connection"
   - "When attempting to access a remote computer, you are denied access or get an error message" in the "Troubleshooting" chapter

4. Use the Remote Connection dialog box to log into the remote system by doing one of the following:

   - Type the administrator account name and password for the computer you want to connect to. You can select the domain as well from the pull-down list, but you don't have to.

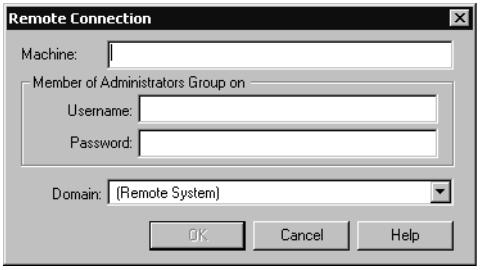
   OR

   - Type the domain administrator account name and password of the domain the computer is in (or a trusted domain to the computer). In these cases, you must select the domain from the pull-down list.

   These domain names appear under **My Network Places/Microsoft Windows Network**. The domain field is not needed when connecting to a NetWare computer running Array Manager.

Note If a domain name is not selected, then Array Manager uses a local user account.

Remote Connection Dialog Box



5. Click **Connect**.

   This connection may take several minutes to complete. The domain name is automatically updated in the window when you change it. The computer icon for the computer you connected to displays in the tree view of the Array Manager console.

Connecting to a Remote Windows System from a Command Line

To connect to a remote computer from a command line:

1. Open a DOS command shell.

2. Enter the following as a single command string. You should only press the <Enter> key after you have typed the following command string in its entirety.

   amconsole $computer:*ComputerName* user:*UserName* password:*Password* domain:*DomainName*

   where:

   - *ComputerName* is the name of the computer to which you are connecting.
   - *DomainName* is the domain to which the remote computer belongs. When connecting to a NetWare system, omit this parameter. It does not apply to NetWare.
   - *Password* is the password of a user with administrator authority. When connecting to a remote system with Array Manager, it is recommended that the password be the same on both the local and remote system. When connecting to a NetWare system, omit the *Password* parameter. NetWare will prompt you for the password.
   - *UserName* is the name of a user with administrator authority. When connecting to a remote system with Array Manager, it is recommended that the user name be the same on both the local and remote system. When connecting to a NetWare system, omit the *UserName* parameter. NetWare will prompt you for the user name.

Connecting to NetWare Servers

When you connect to a NetWare server from an Array Manager console, you need to enter a DNS name, a Hosts name, or the IP address of the Netware Server to which you are connecting. When the connection is made, you will see the NetWare Server name on the console tree.

To do NetWare discovery for NetWare servers, you must have a NetWare client running on the machine on which you are running Array Manager console.

When you click NetWare Services, you see all the remote servers running NetWare displayed underneath.

1. To log on to a displayed server, double-click it to bring up the Remote Connection dialog box.

2. Enter your full distinguished name and password and click **OK**.

   For example, "**cn=admin.0=dept**"

   This will log you on to the selected server, if you have administrator rights to that server. If you don't have administrator rights to the server, then your attempt to connect will be rejected.

   My Network Places displays all connected NetWare servers.

Note If you are unable to connect to a NetWare server, see the section "You are unable to connect to a NetWare server" in the "Troubleshooting" chapter.

Note Dell does not offer NetWare in Japan.

Identifying the Port Number for Remote Connection

Array Manager uses three ports for remote connections as described in the following table.

Port ID	Port Type and Description
TCP / UDP 161	SNMP Management — used by the Array Manager managed system (server) to receive SNMP management commands or request messages.
TCP / UDP 162	SNMP Trap — used by the Array Manager SNMP extension to receive traps from the Array Manager managed system (server). This is the port on the local management host that is used to receive SNMP traps from the Array Manager SNMP extension agent.
UDP 2148	Array Manager Console — used by Array Manager to allow remote Array Manager consoles (clients) to connect.

Note Dell does not offer NetWare in Japan.

Disconnecting from a Remote Computer

To disconnect from a remote computer

1. Right-click the computer icon you want to disconnect from, or from the **Computer** pull-down menu, select **Disconnect**.

   The remote connection ends with no message. The computer icon disappears from the left tree view.

If the connection is lost for any other reason, a message appears that indicates the connection has been terminated. Array Manager removes the icon for that computer from the left tree view. Click **OK** to close the message box.

Using History to View Recent Connections

Array Manager uses the **History** folder in the tree view to store a list of the most recent connections made from the computer. You may want to delete computer icons from the **History** list.

Deleting a Computer from the History List

- Expand the **History** folder and right-click the computer you want to delete.
- Select **Remove Computer** from the context menu that comes up.

Using Favorites

You can use the **Favorites** folder in the tree view for easy storage and retrieval of frequently used computer names. The **Favorites** folder lets you display and retrieve NetWare computer names, as well as the names of Windows computers.

Adding a Computer to Favorites

- To place computer names into **Favorites**, select a computer name in the tree view and then drag and drop it into the **Favorites** folder.
- You can also select the **Connect** command to connect to a remote computer and click the option **Add to Favorites folder** in the initial dialog box that comes up. Right-clicking the **Favorites** folder in the tree view also brings up a menu that has the **Connect** command.

Removing a Computer from Favorites

- To remove a computer from the list, right-click the computer icon you want to remove and select **Remove Computer**.

Connecting to a “Favorite” Computer

- To open any computer in the **Favorites** folder, double-click the computer’s icon. If the Remote Connection dialog box appears, enter your name and password as appropriate to connect to this computer.
- The **Add to Favorites** folder option is always paired with the option **Reconnect at next start of application**. If you select this option, Array Manager will attempt to connect to the “Favorite” remote computers when you restart Array Manager.

Note If you are logged on to the computer initially as a domain administrator, the "Favorites" computers should come up automatically and appear in the tree view (if you selected to reconnect at startup and the remote computers are online).

Event Monitoring 20

Array Manager's event management features let you monitor the health of storage resources such as controllers, channels, array disks, virtual disks, and volumes.

Event Monitoring and Logging

The Disk Management Service performs event monitoring and logging on Array Manager managed systems. By default, the Disk Management Service starts when the managed system starts up. If you stop the Disk Management Service, then event monitoring and logging stops.

- Updates the status of the storage object that generated the event.
- Propagates the storage object's status to all the related higher objects in the storage hierarchy. For example, if the tree view is collapsed, the computer object will show the alert icon.
- Logs an event into the Array Manager log and Windows application event log.
- Sends an SNMP trap if the operating system's SNMP service is installed and enabled.
- Changes the color of the Array Manager task bar icon (displayed in the upper-left corner of the Array Manager window) to red. The task bar icon remains red until the alerting event displayed on the Events tab is opened.

Note A critical error will cause a pop-up "critical event" message to display even if the console is not running.

Note Array Manager does not log events regarding the data I/O path. These events are logged by the respective RAID drivers in the system event log.

Viewing Event Listings

Array Manager generates events that are added to the Windows application event log. You can view these events in the Array Manager Events tab or in the Microsoft Windows Event Viewer.

In the right pane of the Array Manager console, click the **Events** tab to display the events listings. All events that have been logged appear on this tab. You can also display the **Events** tab by clicking the **View Events** button on the **Quick Access Window**.

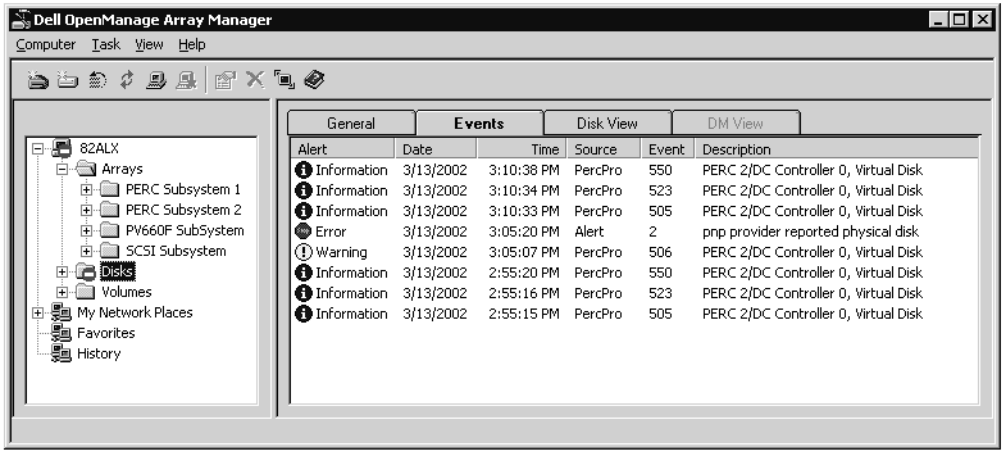
Every event consists of the following:

- **Alert** — Shows the severity of alert.
- **Date** — Date when Array Manager logged the event.
- **Time** — Time when Array Manager logged the event.
- **Source** — Array Manager component that logged the event.
- **Event** — Displays the event ID or number. You can use this number when locating the event descriptions in the "Event Messages" section.
- **Description** — A brief description of the event.

Note The Array Manager Event Log file cannot be viewed through a text editor.

Note If you are receiving an error message on your Windows system stating that the event log is full, you may need to clear the event log in order for incoming events to be displayed in the event log. See "Clearing Events" for information on clearing the Array Manager event log. For information on the Windows event log settings, see your Windows documentation.

Events Tab



Viewing Event Details

You can view specific information on each event listed on the Events tab by double-clicking the event to display the Event Details dialog box. The Event Details dialog box displays the same information as the Events tab, as well as additional information such as the computer on which the event was generated.

Viewing SNMP Traps

SNMP traps that are generated by Array Manager can be viewed in any standard SNMP-compatible enterprise management console. These traps are defined in the Array Manager management information base (MIB). These traps correspond to the events documented in the "Event Messages" section. For more information on the MIB and its structure, as well as a change history of the SNMP traps, see the PDF document titled *Dell OpenManage Array Manager MIB Structure*. The file name for this document is amMIBstructure.pdf. The default location for this file is in the Program Files\Dell\OpenManage\Array Manager\MIB directory.

For more information on configuring SNMP, see "SNMP Support."

Viewing Traps with HP OpenView Network Node Manager

To view trap events with strings inserted with HP OpenView Network Node Manager (NNM), perform the following steps:

1. Load the Array Manager MIB (arymgr.mib) into the HP OpenView NNM. The arymgr.mib file is located in the Management Information Base (MIB) directory of your Array Manager installation directory.

2. Select **OK** When prompted to Load a **TRAP-TYPE/NOTIFICATION-TYPE macro**. The system loads all of the trap definitions into HP OpenView.

For more information on loading MIBs into HP OpenView, see the HP OpenView documentation.

Managing the Event Log

This section describes how you can manage the Array Manager Event Log.

Note If you uninstall Array Manager or modify, repair, or remove installation components, you may notice that events generated before the uninstall, modify, repair, or remove no longer display an event description. This situation occurs when the component that generated the event is no longer installed on the system. If you wish to preserve the events and their descriptions, you should back up or export the event log before you uninstall Array Manager, or modify, repair, or remove installation components.

Backing Up and Restoring Events

You can choose to back up the Array Manager Event Log. Backing up saves the file to a specific, predefined location that is not user-selectable. This file cannot be viewed through a text editor.

To Back Up Events

1. Right-click in the Events tab and select **Backup Event Log** from the context menu that appears.

2. The Event Log is stored locally.

To Restore Events

You can restore the most recently saved Event Log.

1. Right-click in the **Events** tab area and select **Restore Event Log** from the context menu that comes up**.**

2. The saved events reappear in the **Events** tab view.

Clearing Events

You can clear all Array Manager events stored in the Event Log. This will remove all entries displayed under the **Events** tab. You can restore the entries by using the **Restore Event Log** option (if the Event Log was backed up). This restores the most recently backed-up Event Log entries.

To Clear Events

1. Right-click in the **Events** tab area and select **Clear All Events** from the context menu that comes up.

2. Click **Yes** to confirm. This clears the Event Log.

Exporting Events to a Text File

You can export the events in the Array Manager Event Log to a text file. You might do this for archival purposes before clearing the event log. Exporting to a text file also lets you search the events using a text editor.

To Export Events to a Text File

1. Right-click in the **Events** tab area and select **Export All Events to Text File** from the context menu.

2. Specify the text file name and location in the **Save As** dialog box and click **Save**.

Setting the Event Log Size

You can set the size of the Array Manager Event Log to a specific value, so that it will take only a certain amount of space on the disk. The default log size is 10 MB.

If the Event Log is full, every new system event generates a pop-up message indicating that the log file is full. The message also requests that you clear the log file.

You can clear the Event Log or you can continue to let events be written into it. This file will take log entries as long as there is writable space on the disk. Therefore, clearing the file is not required unless your disk space is extremely limited. You may wish to increase the size of your Event Log file to eliminate the error messages that are generated each time a log entry appears.

To Set the Event Log Size

1. Right-click in the **Events** tab area and select **Set Event Log Size** from the context menu that comes up.
2. The **Set Event Log Size** dialog box appears, showing the current size of the log.
3. Enter a new size for the file and press **OK**. The **Set Event Log Size** dialog box closes automatically.

Event Type or Severity Levels

Each event message in the Array Manager event log has a severity level. The severity level is displayed in the Alert field of the event message. The severity level indicates the nature of the event.

The event severity levels are as follows:

- **Information** — The event is provided for informational purposes and does not indicate an error condition. An Information event indicates the normal start or stop of an operation. For example, an information event may indicate that an operation such as a consistency check has started or come to completion.
- **Warning** — A Warning event indicates a potential problem, but does not necessarily mean that the system has currently lost data or is nonfunctional. For example, a Warning event may indicate that a component (such as a temperature probe in an enclosure) has crossed a warning threshold.
- **Error** — An Error event indicates a serious problem such as a loss of data or a loss of function. For example, an Error event may indicate that an array disk has failed.

Event Source

Each event message in the event log displays the source of the event. The following table shows the subsystem terminology (Array Manager console) and the event source names as displayed in the Array Manager event log and the Windows event log. The subsystem terminology does not apply to events that have fsys as their source.

Array Manager Console	Array Manager Event Log	Windows Event Log
PERC Subsystem 1	PercPro	VxSvc_PercPro
PERC Subsystem 2	Perc2Pro	VxSvc_Perc2Pro
SCSI Subsystem	AspiPro	VxSvc_AspiPro
N / A	fsys	VxSvc_fsys
PV660F SubSystem	Mylex	VxSvc_Mylex

Variable Substitution in the Event Messages

When you view an event in the Array Manager event log, the event identifies the specific components such as the controller name or the virtual disk name to which the event applies. In an actual operating environment, a storage system can have many combinations of controllers and disks as well as user-defined names for virtual disks and other components. Because each environment is unique in its storage configuration and user-defined names, an accurate event message requires that Array Manager be able to insert the environment-specific names of storage components into an event message.

In order to provide environment-specific information in the event message, Array Manager uses variables for the event message text. When a controller or other component generates an event, Array Manager substitutes each variable with the controller name and other specifics. The modified event containing the environment-specific information is then displayed in the Array Manager event log.

This online help provides a definition of each event. When defining the event message, this online help indicates which portion of the message is constructed from a variable, and which components are substituted for the variable when Array Manager applies the environment-specific information. In the event definition, variables are represented by the % character (for example, %s, %1, %2, and so forth). The SNMP trap variables indicate how the % is substituted to provide environment-specific information.

Consider the following example event:

Description: Virtual Disk check consistency started

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_TASK_CHECK_CONSISTENCY

Severity: Information

Source: PercPro (PERC Subsystem 1)

Text: %1, Virtual Disk (%2) consistency check started.

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

In this example, the **Text** field indicates that two variables (%1, %2) are used to construct the event message displayed in the Array Manager event log. The **SNMP Trap Variables** field indicates that a controller name (controllerNameEv) and a virtual disk name (virtualDiskNameEv) substitute for the %1 and %2 variables. Therefore, if a PERC 3/Di controller number zero generated the event because a virtual disk number one with a user-defined name of HRdata has started a consistency check, then the resulting event message displayed in the event log is as follows:

PERC 3/Di Controller 0, Virtual Disk (HRdata 1) consistency check started.

Event Messages

The following sections describe events generated by the fibre channel RAID controllers in the PowerVault 660F storage system. The events are displayed in the Events tab of the Array Manager console or through Windows Event Viewer. These events can also be forwarded as SNMP traps to other applications.

SNMP traps are generated for the events listed in the following sections. These traps are included in the Array Manager management information base (MIB). For more information on the MIB and its structure, as well as a change history of the SNMP traps, see the PDF document titled *Dell OpenManage Array Manager MIB Structure*. The file name for this document is amMIBstructure.pdf. The default location for this file is in the Program Files\Dell\OpenManage\Array Manager\MIB directory. For more information on Array Manager's SNMP support, see "SNMP Support."

Select one of the following to view controller specific events:

- Events Generated by the PERC 2/SC, 3/SC, 2/DC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, 4/IM, 4e/DC, 4e/Si, 4e/Di, CERC ATA100/4ch and the PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s Controllers — These events correspond to event numbers 500 through 598.
- Events Generated by the PowerVault 660F Storage System — These events correspond to event numbers 700 through 927.
- Volume Management Events — These events correspond to event numbers 10804 through 10805.

Events Generated by the PERC 2/SC, 3/SC, 2/DC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, 4/IM, 4e/DC, 4e/Si, 4e/Di, CERC ATA100/4ch and the PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s Controllers

The following sections describe events generated by the following controllers:

- PERC 2/SC
- PERC 3/SC
- PERC 2/DC
- PERC 3/DCL
- PERC 3/DC
- PERC 3/QC
- PERC 4/SC
- PERC 4/DC
- 4e/DC
- PERC 4/Di
- PERC 4/IM
- PERC 4e/Si
- PERC 4e/Di
- CERC ATA100/4ch
- PERC 2
- PERC 2/Si
- PERC 3/Si
- PERC 3/Di
- CERC SATA1.5/6ch
- CERC SATA1.5/2s

To locate an event, scroll through these sections to find the event number displayed on the Events tab of the Array Manager console or use the Search function in this online help file to search for the event message text.

500

Description: **Array Disk Failed**

Cause: A physical disk in the array failed. The failed disk may have been identified by the controller or channel. Performing a consistency check can also identify a failed disk.

Action: Replace the failed array disk. You can identify which disk has failed by locating the disk with a red "X" in the Array Manager console's tree view. Perform a rescan after replacing the disk.

Name: PERC_ALERT_ARRAY_DISK_FAILURE or PERC2_ALERT_ARRAY_DISK_FAILURE or ASPI_ALERT_ARRAY_DISK_FAILURE

Severity: ErrorPercPro (PERC Subsystem 1)

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %s, %s failed.),

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

Note In certain circumstances, this event is not generated when an array disk fails. See Array Disk State, Event Messages and Hot Spares on PERC Subsystem 1 Controllers for more information.

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

501

Description: **Array Disk removed**

Cause: A physical disk has been removed from the array. This event can also be caused by loose or defective cables or by problems with the enclosure.

Action: If a physical disk was removed from the array, either replace the disk or restore the original disk. You can identify which disk has been removed by locating the disk with a red "X" in the Array Manager console's tree view. Perform a rescan after replacing or restoring the disk. If a disk has not been removed from the array, then check for problems

with the cables. See the "Cables attached correctly" section for more information on checking the cables. Make sure that the enclosure is powered on. If the problem persists, check the enclosure documentation for further diagnostic information.

Name: PERC_ALERT_ARRAY_DISK_REMOVED or PERC2_ALERT_ARRAY_DISK_REMOVED or ASPI_ALERT_ARRAY_DISK_REMOVED

Severity: Warning

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %s, %s removed.

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

502

Description: Array Disk offline

Cause: A physical disk in the array is offline. A disk can be made offline during a Prepare to Remove operation or because a user manually put the disk offline.

Action: Perform a rescan. You can also select the offline disk and perform a Make Online operation.

Name: PERC_ALERT_ARRAY_DISK_OFFLINE or PERC2_ALERT_ARRAY_DISK_OFFLINE or ASPI_ALERT_ARRAY_DISK_OFFLINE

Severity: Warning

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %s, %s offline.

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

503

Description: **Array Disk degraded**

Cause: A physical disk in the array is degraded. The degraded disk may have been identified by the controller or channel. Performing a consistency check can also identify a degraded disk.

Action: Replace the degraded array disk. You can identify which disk is degraded by locating the disk with a red "X" in the Array Manager console's tree view. Perform a rescan after replacing the disk.

Name: PERC_ALERT_ARRAY_DISK_DEGRADED or PERC2_ALERT_ARRAY_DISK_DEGRADED or ASPI_ALERT_ARRAY_DISK_DEGRADED

Severity: Warning

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %s, %s degraded.

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

504

Description: Array Disk inserted

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_NEW_DISK_INSERTED or PERC2_ALERT_NEW_DISK_INSERTED or ASPI_ALERT_NEW_DISK_INSERTED

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %s, %s inserted.

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

505

Description: Virtual Disk created

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_VIRTUAL_DISK_CREATED or PERC2_ALERT_VIRTUAL_DISK_CREATED

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %s, virtual disk (%s) created.

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

506

Description: **Virtual Disk deleted**

Cause: A virtual disk has been deleted. Performing a Reset Configuration may detect that a virtual disk has been deleted and generate this event.

Action: None.

Name: PERC_ALERT_VIRTUAL_DISK_DELETED or PERC2_ALERT_VIRTUAL_DISK_DELETED

Severity: Warning

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %s, virtual disk (%s) deleted.

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

507

Description: **Virtual Disk created**

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_VIRTUAL_DISK_CONFIGURATION_CHANGED or PERC2_ALERT_VIRTUAL_DISK_CONFIGURATION_CHANGED

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %1, Virtual Disk (%2) configuration changed from %3 to %4.

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv, oldVDConfigEv, newVDConfigEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

508

Description: Virtual Disk failed

Cause: One or more physical disks included in the virtual disk have failed. If the virtual disk is non-redundant (does not use mirrored or parity data), then the failure of a single physical disk can cause the virtual disk to fail. If the virtual disk is redundant, then more physical disks have failed than can be rebuilt using mirrored or parity information.

Action: Create a new virtual disk and restore from a backup.

Name: PERC_ALERT_VIRTUAL_DISK_FAILED or PERC2_ALERT_VIRTUAL_DISK_FAILED

Severity: Error

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %1, Virtual Disk (%2) failed.

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

509

Description: **Virtual Disk degraded**

Cause 1: This event message occurs when a physical disk included in a redundant virtual disk fails. Because the virtual disk is redundant (uses mirrored or parity information) and only one physical disk has failed, the virtual disk can be rebuilt.

Action 1: Configure a hot spare for the virtual disk if one is not already configured. Rebuild the virtual disk. When using a PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, or CERC ATA100/4ch controller, rebuild the virtual disk by first configuring a hot spare for the disk, and then initiating a write operation to the disk. The write operation will initiate a rebuild of the disk.

Cause 2: A physical disk in the array has been removed.

Action 2: If a physical disk was removed from the array, either replace the disk or restore the original disk. You can identify which disk has been removed by locating the disk with a red "X" in the Array Manager console's tree view. Perform a rescan after replacing the disk.

Name: PERC_ALERT_VIRTUAL_DISK_DEGRADED or PERC2_ALERT_VIRTUAL_DISK_DEGRADED

Severity: Warning

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %1, Virtual Disk (%2) degraded.

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

510

Description: Virtual disk is no longer redundant

Cause: A physical disk included in a redundant virtual disk has failed. Due to the failed drive or drives, the virtual disk is no longer maintaining redundant (mirrored or parity) data. The failure of an additional drive will result in lost data.

Action: Replace the failed drives and rebuild the virtual disk.

Name: PERC2_ALERT_VIRTUAL_DISK_FAILED_REDUNDANCY

Severity: Warning

Source: Perc2Pro (PERC Subsystem 2)

Text: %1, Virtual Disk (%2) failed.

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

520

Description: Virtual Disk check consistency started

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_TASK_CHECK_CONSISTENCY or PERC2_ALERT_TASK_CHECK_CONSISTENCY

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %1, Virtual Disk (%2) consistency check started.

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

521

Description: **Virtual Disk format started**

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_TASK_VD_FORMAT or PERC2_ALERT_TASK_VD_FORMAT

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %1, Virtual Disk (%2) format started.

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

522

Description: **Array Disk format started**

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_TASK_DISK_FORMAT or PERC2_ALERT_TASK_DISK_FORMAT or ASPI_ALERT_TASK_DISK_FORMAT

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 format started.

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

523

Description: Virtual Disk initialization started

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_TASK_VD_INITIALIZE or PERC2_ALERT_TASK_VD_INITIALIZE

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %1, Virtual Disk (%2) initialization started.

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

524

Description: **Array Disk initialize started**

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_TASK_DISK_INITIALIZE or PERC2_ALERT_TASK_DISK_INITIALIZE or ASPI_ALERT_TASK_DISK_INITIALIZE

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 initialization started.

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

525

Description: **Virtual Disk Reconfiguration started**

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_TASK_RECONFIGURATION or PERC2_ALERT_TASK_RECONFIGURATION

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %1, Virtual Disk (%2) reconfiguration from %3 stripe size to %4 stripe size started.

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv, oldVDConfigEv, newVDConfigEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

526

Description: Virtual Disk rebuild started

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_TASK_VD_REBUILD or PERC2_ALERT_TASK_VD_REBUILD

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %1, Virtual Disk (%2) rebuild started.

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

527

Description: Array Disk rebuild started

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_TASK_DISK_REBUILD or PERC2_ALERT_TASK_DISK_REBUILD or ASPI_ALERT_TASK_DISK_REBUILD

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 rebuild started.

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

528

Description: Array Disk diagnostics started

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_TASK_DIAGNOSE or PERC2_ALERT_TASK_DIAGNOSE or ASPI_ALERT_TASK_DIAGNOSE

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 diagnostics started.

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

529

Description: Virtual Disk check consistency cancelled

Cause: The check consistency operation cancelled because a physical disk in the array has failed or because a user cancelled the check consistency operation.

Action: If the physical disk failed, then replace the physical disk. You can identify which disk failed by locating the disk with a red "X" in the Array Manager console's tree view. Perform a rescan after replacing the disk. When performing a consistency check, be aware that the consistency check can take a long time. The time it takes depends on the size of the physical disk or the virtual disk.

Name: PERC_ALERT_CANCEL_CHECK_CONSISTENCY or PERC2_ALERT_CANCEL_CHECK_CONSISTENCY

Severity: Warning

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %1, Virtual Disk (%2) consistency check cancelled.

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

530

Description: Virtual Disk format cancelled

Cause: The format virtual disk operation cancelled because a physical disk included in the virtual disk has failed or because a user cancelled the format virtual disk operation.

Action: If a physical disk failed, then replace the physical disk. You can identify which disk has failed by locating the disk with a red "X" in the Array Manager console's tree view. Perform a rescan after replacing the disk. Restart the format virtual disk operation.

Name: PERC_ALERT_CANCEL_VD_FORMAT or PERC2_ALERT_CANCEL_VD_FORMAT

Severity: Warning

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %1, Virtual Disk (%2) format cancelled.

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

531

Description: **Array Disk format cancelled**

Cause: The format array disk operation cancelled because a physical disk in the array has failed or because a user cancelled the format array disk operation.

Action: If a physical disk failed, then replace the physical disk. You can identify which disk has failed by locating the disk with a red "X" in the Array Manager console's tree view. Perform a rescan after replacing the disk. Restart the format array disk operation.

Name: PERC_ALERT_CANCEL_DISK_FORMAT or PERC2_ALERT_CANCEL_DISK_FORMAT or ASPI_ALERT_CANCEL_DISK_FORMAT

Severity: Warning

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 format cancelled.

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

532

Description: Virtual Disk initialization cancelled

Cause: The virtual disk initialization cancelled because a physical disk included in the virtual disk has failed or because a user cancelled the virtual disk initialization.

Action: If a physical disk failed, then replace the physical disk. You can identify which disk has failed by locating the disk with a red "X" in the Array Manager console's tree view. Perform a rescan after replacing the disk. Restart the format array disk operation. Restart the virtual disk initialization.

Name: PERC_ALERT_CANCEL_VD_INITIALIZE or PERC2_ALERT_CANCEL_VD_INITIALIZE

Severity: Warning

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %1, Virtual Disk (%2) initialization cancelled.

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

533

Description: Array Disk initialize cancelled

Cause: The array disk initialization cancelled because a physical disk in the array has failed or because a user cancelled the array disk initialization.

Action: If a physical disk failed, then replace the physical disk. You can identify which disk has failed by locating the disk with a red "X" in the Array Manager console's tree view. Perform a rescan after replacing the disk. Restart the array disk initialization.

Name: PERC_ALERT_CANCEL_DISK_INITIALIZE or PERC2_ALERT_CANCEL_DISK_INITIALIZE or ASPI_ALERT_CANCEL_DISK_INITIALIZE

Severity: Warning

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 initialization cancelled.

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

534

Description: **Virtual Disk Reconfiguration cancelled**

Cause: A physical disk in the array has failed. A Transform Type, Add Member, or Reconfigure operation can identify that a disk has failed and generate this event.

Action: Replace the failed array disk. You can identify which disk has failed by locating the disk with a red "X" in the Array Manager console's tree view. Perform a rescan after replacing the disk. Restart the Transform Type, Add Member, or the Reconfigure operation.

Name: PERC_ALERT_CANCEL_RECONFIGURATION or PERC2_ALERT_CANCEL_RECONFIGURATION

Severity: Warning

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %1, Virtual Disk (%2) reconfiguration cancelled.

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

535

Description: **Virtual Disk rebuild cancelled**

Cause: A user has cancelled the rebuild operation.

Action: Restart the rebuild operation.

Name: PERC_ALERT_CANCEL_VD_REBUILD or PERC2_ALERT_CANCEL_VD_REBUILD

Severity: Warning

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %1, Virtual Disk (%2) rebuild cancelled.

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

536

Description: **Array Disk rebuild cancelled**

Cause: A user has cancelled the rebuild operation.

Action: Restart the rebuild operation.

Name: PERC_ALERT_CANCEL_DISK_REBUILD or PERC2_ALERT_CANCEL_DISK_REBUILD or ASPI_ALERT_CANCEL_DISK_REBUILD

Severity: Warning

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 rebuild cancelled.

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

537

Description: **Array Disk diagnostics cancelled**

Cause: A user cancelled the array disk diagnostics operation.

Action: Restart the array disk diagnostics operation.

Name: PERC_ALERT_CANCEL_DIAGNOSE or PERC2_ALERT_CANCEL_DIAGNOSE or ASPI_ALERT_CANCEL_DIAGNOSE

Severity: Warning

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 diagnostics cancelled.

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

538

Description: **Virtual Disk check consistency failed**

Cause: An array disk included in the virtual disk failed or there is an error in the parity information. A failed array disk can cause errors in parity information.

Action: Replace the failed array disk. You can identify which disk has failed by locating the disk with a red "X" in the Array Manager console's tree view. Rebuild the array disk. When finished, restart the check consistency operation.

Name: PERC_ALERT_FAIL_CHECK_CONSISTENCY or PERC2_ALERT_FAIL_CHECK_CONSISTENCY

Severity: Error

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %1, Virtual Disk (%2) consistency check failed.

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

539

Description: Virtual Disk format failed

Cause: An array disk included in the virtual disk failed.

Action: Replace the failed array disk. You can identify which disk has failed by locating the disk with a red "X" in the Array Manager console's tree view. Rebuild the array disk. When finished, restart the format operation.

Name: PERC_ALERT_FAIL_VD_FORMAT or PERC2_ALERT_FAIL_VD_FORMAT

Severity: Error

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %1, Virtual Disk (%2) format failed.

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

540

Description: **Array Disk format failed**

Cause: The array disk has failed.

Action: Replace the failed array disk. You can identify which disk has failed by locating the disk with a red "X" in the Array Manager console's tree view. If the array disk is part of a redundant array, then rebuild the array disk.

Name: PERC_ALERT_FAIL_DISK_FORMAT or PERC2_ALERT_FAIL_DISK_FORMAT or ASPI_ALERT_FAIL_DISK_FORMAT

Severity: Error

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 format failed.

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

541

Description: **Virtual Disk initialization Failed**

Cause: An array disk included in the virtual disk has failed or a user has cancelled the initialization.

Action: If an array disk has failed, then replace the array disk.

Name: PERC_ALERT_FAIL_VD_INITIALIZE or PERC2_ALERT_FAIL_VD_INITIALIZE

Severity: Error

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %1, Virtual Disk (%2) initialization failed.

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

542

Description: **Array Disk initialize failed**

Cause: The array disk has failed or is corrupt.

Action: Replace the failed or corrupt disk. You can identify a disk that has failed by locating the disk with a red "X" in the Array Manager console's tree view. Restart the initialization.

Name: PERC_ALERT_FAIL_DISK_INITIALIZE or PERC2_ALERT_FAIL_DISK_INITIALIZE or ASPI_ALERT_FAIL_DISK_INITIALIZE

Severity: Error

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 initialization failed.

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

543

Description: **Virtual Disk Reconfiguration failed**

Cause: An array disk included in the virtual disk has failed or is corrupt. A user may also have cancelled the reconfiguration.

Action: Replace the failed or corrupt disk. You can identify a disk that has failed by locating the disk with a red "X" in the Array Manager console's tree view. If the array disk is part of a redundant array, then rebuild the array disk. When finished, restart the reconfiguration.

Name: PERC_ALERT_FAIL_RECONFIGURATION or PERC2_ALERT_FAIL_RECONFIGURATION

Severity: Error

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %1, Virtual Disk (%2) reconfiguration failed.

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

544

Description: **Virtual Disk rebuild failed**

Cause: An array disk included in the virtual disk has failed or is corrupt. A user may also have cancelled the rebuild.

Action: Replace the failed or corrupt disk. You can identify a disk that has failed by locating the disk with a red "X" in the Array Manager console's tree view. Restart the virtual disk rebuild.

Name: PERC_ALERT_FAIL_VD_REBUILD or PERC2_ALERT_FAIL_VD_REBUILD

Severity: Error

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %1, Virtual Disk (%2) rebuild failed.

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

545

Description: **Rebuild failed.**

Cause: A rebuild was initiated and it failed to successfully complete.

Action: The array disk that was rebuilding should be replaced so that another rebuild can occur.

Name: PERC_ALERT_FAIL_DISK_REBUILD or PERC2_ALERT_FAIL_DISK_REBUILD

Severity: Error

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %1, %2 rebuild failed.

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

546

Description: **Array Disk diagnostics failed**

Cause: The array disk disk has failed or is corrupt. A user may also have cancelled the diagnostics operation.

Action: Replace the failed or corrupt disk. You can identify a disk that has failed by locating the disk with a red "X" in the Array Manager console's tree view. If the disk is part of a redundant array, then rebuild the disk. Restart the diagnostics operation.

Name: PERC_ALERT_FAIL_DIAGNOSE or PERC2_ALERT_FAIL_DIAGNOSE or ASPI_ALERT_FAIL_DIAGNOSE

Severity: Error

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 diagnostics failed.

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

547

Description: Virtual Disk check consistency completed

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_COMPLETE_CHECK_CONSISTENCY or PERC2_ALERT_COMPLETE_CHECK_CONSISTENCY

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %1, Virtual Disk (%2) consistency check completed.

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

548

Description: Virtual Disk format completed

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_COMPLETE_VD_FORMAT or PERC2_ALERT_COMPLETE_VD_FORMAT

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %1, Virtual Disk (%2) format completed.

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

549

Description: Array Disk format completed

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_COMPLETE_DISK_FORMAT or PERC2_ALERT_COMPLETE_DISK_FORMAT or ASPI_ALERT_COMPLETE_DISK_FORMAT

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 format completed.

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

550

Description: **Virtual Disk initialization completed**

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_COMPLETE_VD_INITIALIZE or PERC2_ALERT_COMPLETE_VD_INITIALIZE

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %1, Virtual Disk (%2) initialization completed.

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

551

Description: **Array Disk initialize completed**

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_COMPLETE_DISK_INITIALIZE or PERC2_ALERT_COMPLETE_DISK_INITIALIZE or ASPI_ALERT_COMPLETE_DISK_INITIALIZE

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 initialization completed.

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

552

Description: Virtual Disk Reconfiguration completed

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_COMPLETE_RECONFIGURATION or PERC2_ALERT_COMPLETE_RECONFIGURATION

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %1, Virtual Disk (%2) reconfiguration completed.

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

553

Description: **Virtual Disk rebuild completed**

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_COMPLETE_VD_REBUILD or PERC2_ALERT_COMPLETE_VD_REBUILD

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %1, Virtual Disk (%2) rebuild completed.

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

554

Description: **Array Disk rebuild completed**

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_COMPLETE_DISK_REBUILD or PERC2_ALERT_COMPLETE_DISK_REBUILD ASPI_ALERT_COMPLETE_DISK_REBUILD

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 rebuild completed.

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

555

Description: Array Disk diagnostics completed

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_COMPLETE_DIAGNOSE or PERC2_ALERT_COMPLETE_DIAGNOSE or ASPI_ALERT_COMPLETE_DIAGNOSE

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 diagnostics completed.

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

570

Description: Predictive Failure reported. If this disk is part of a redundant virtual disk, select the 'Offline' option and then replace the disk. Then configure a hot spare and it will start the rebuild automatically. If this disk is a hot spare, select the 'Prepare to

Remove' option and then replace the disk. If this disk is part of a non-redundant disk, you should back up your data immediately. If the disk fails, you will not be able to recover the data.

Cause: The array disk is predicted to fail. Many array disks contain Self Monitoring Analysis and Reporting Technology (S.M.A.R.T.). When enabled, SMART monitors the health of the disk based on indications such as the number of write operations that have been performed on the disk.

Action: Replace the array disk. Even though the disk may not have failed yet, it is strongly recommended that you replace the disk. Review the message text for additional information.

Name: PERC_ALERT_PREDICTIVE_FAILURE or ASPI_ALERT_PREDICTIVE_FAILURE

Severity: Warning

Source: PercPro (PERC Subsystem 1), AspiPro (SCSI Subsystem)

Text: %1, %2 reported predictive failure.

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

571

Description: **Sense Key = X, Sense Code = X, Sense Qualifier = X. Please refer to the Array Manager Documentation for details.**

Cause: This event indicates that a SCSI device has reported a condition requiring user attention.

Action: Drives experiencing valid failures will normally report status information to the system logs. Refer to the tables below to aid in response to the reported conditions. **Table A** should be used as a guideline on when to replace a drive without the need for additional troubleshooting. Drives reporting the errors in **Table A** should be scheduled for replacement. Error codes reported in **Table B** may require drive reseating or additional troubleshooting.

Table A: Sense data to justify replacing the SCSI HDD.

Sense Key (SK)	Sense Code (ASC)	Qualifier (ASCQ)	Comments
02	04	00	Drive not ready. For ASC and ASCQ other than “04-00“, additional trouble shooting should be performed.
03	XX	XX	Medium error. See notes.
04	XX	XX	Unrecovered hardware error.
09	XX	XX	Vendor-specific error.
XX	5D	XX	Predictive error reported. (SMART error.)

Note Message SK 03, ASC 11, ASCQ 00 indicates medium error was encountered. The RAID controllers will reassign the defective media block when encountered. Replacement should not be performed unless excessive defects are identified on a single drive.

Note Message SK 03, ASC 31 indicates the low level drive format has been corrupted. If a format via SCSI controller is not possible, drive replacement is recommended.

Table B: Sense data to justify additional troubleshooting/reseating required.

Sense Key (SK)	Sense Code (ASC)	Qualifier (ASCQ)	Comments
01	XX	XX	Recovered error.
02	XX	XX	Drive not ready. Replace only if 02-04-00.
05	XX	XX	Illegal request.
06	XX	XX	Unit Attention.
08	5D	XX	Aborted command.

Name: PERC_ALERT_SCSI_SENSE_DATA or ASPI_ALERT_SCSI_SENSE_DATA

Severity: Warning

Source: PercPro (PERC Subsystem 1),AspiPro (SCSI Subsystem)

Text: %1, %2 Sense Key = %3, Sense Code = %4, Sense Qualifier = %5.

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. (When forwarded as an SNMP trap, this event only has two variables.) For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

572

Description: I/O paused. I/O will automatically resume after the allotted time.

Cause: I/O operations are paused due to a user action. The I/O operations will automatically resume after the indicated period of time.

Action: Wait for the I/O operations to resume automatically.

Name: PERC_ALERT_PAUSE_IO or ASPI_ALERT_PAUSE_IO

Severity: Warning

Source: PercPro (PERC Subsystem 1), AspiPro (SCSI Subsystem)

Text: %1, Channel %2 I/O paused, I/O will automatically resume after the allotted time (%3 seconds).

SNMP Trap Variables: controllerNameEv, channelNumberEv, timeEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

573

Description: I/O resumed

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_RESUME_IO or ASPI_ALERT_RESUME_IO

Severity: Information

Source: PercPro (PERC Subsystem 1), AspiPro (SCSI Subsystem)

Text: %1, Channel %2 I/O resumed.

SNMP Trap Variables: controllerNameEv, channelNumberEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

574

Description: Hot spare assigned

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_HOTSPARE_ASSIGN or PERC2_ALERT_HOTSPARE_ASSIGN

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %1, assigned %2 as Hot Spare.

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

575

Description: **Hot spare unassigned**

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_HOTSPARE_UNASSIGN or PERC2_ALERT_HOTSPARE_UNASSIGN

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %1, unassigned %2 as Hot Spare.

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

579

Description: **Controller battery needs reconditioning.**

Cause: The controller battery charge is low.

Action: Recondition the controller battery.

Name: PERC2_ALERT_BATTERY_NEEDS_RECONDITIONING

Severity: Warning

Source: Perc2Pro (PERC Subsystem 2)

Text: %1, battery needs reconditioning.

SNMP Trap Variables: controllerNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

580

Description: Controller battery low.

Cause: The controller battery charge is low.

Action: Recondition the battery.

Name: PERC2_ALERT_BATTERY_NEEDS_RECONDITIONING

Severity: Warning

Source: Perc2Pro (PERC Subsystem 2)

Text: %1, battery low.

SNMP Trap Variables: controllerNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

581

Description: Controller battery is reconditioning.

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC2_ALERT_TASK_BATTERY_RECONDITION

Severity: Information

Source: Perc2Pro (PERC Subsystem 2)

Text: %1, battery is reconditioning.

SNMP Trap Variables: controllerNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

582

Description: **Controller battery recondition is completed.**

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC2_ALERT_COMPLETE_BATTERY_RECONDITION

Severity: Information

Source: Perc2Pro (PERC Subsystem 2)

Text: %1, battery recondition completed.

SNMP Trap Variables: controllerNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

583

Description: **Controller I/O paused.**

Cause: The controller I/O operations are paused due to a user action. The controller I/O will automatically resume after the indicated period of time.

Action: None.

Name: PERC2_ALERT_PAUSE_IO

Severity: Warning.

Source: Perc2Pro (PERC Subsystem 2)

Text: %1, I/O paused, I/O will automatically resume after the allotted time (%2 seconds).

SNMP Trap Variables: controllerNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

584

Description: Controller I/O resumed.

Cause: The controller I/O has resumed.

Action: None.

Name: PERC2_ALERT_RESUME_IO

Severity: Information.

Source: Perc2Pro (PERC Subsystem 2)

Text: %1, I/O resumed.

SNMP Trap Variables: controllerNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

585

Description: **Smart FPT exceeded.**

Cause: A disk on the specified controller has received a SMART alert (predictive failure) indicating that the disk is likely to fail in the near future.

Action: Replace the disk that has received the SMART alert. If the array disk is a member of a non-redundant virtual disk, then back up the data before replacing the disk. Removing an array disk that is included in a non-redundant virtual disk will cause the virtual disk to fail and may cause data loss.

Name: PERC2_ALERT_SMART_FPT_EXCEEDED

Severity: Warning

Source: Perc2Pro (PERC Subsystem 2)

Text: %1, SMART:FPT_EXCEEDED on %2.

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

586

Description: **Smart configuration change.**

Cause: A disk has received a SMART alert (predictive failure) after a configuration change. The disk is likely to fail in the near future.

Action: Replace the disk that has received the SMART alert. If the array disk is a member of a non-redundant virtual disk, then back up the data before replacing the disk. Removing an array disk that is included in a non-redundant virtual disk will cause the virtual disk to fail and may cause data loss.

Name: PERC2_ALERT_SMART_CONFIG_CHANGE

Severity: Warning

Source: Perc2Pro (PERC Subsystem 2)

Text: %1, SMART_CONFIG_CHANGE on %2.

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

587

Description: **Smart warning.**

Cause: A disk has received a SMART alert (predictive failure). The disk is likely to fail in the near future.

Action: Replace the disk that has received the SMART alert. If the array disk is a member of a non-redundant virtual disk, then back up the data before replacing the disk. Removing an array disk that is included in a non-redundant virtual disk will cause the virtual disk to fail and may cause data loss.

Name: PERC2_ALERT_SMART_WARNING

Severity: Warning

Source: Perc2Pro (PERC Subsystem 2)

Text: %1, SMART_WARNING on %2.

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

588

Description: **Smart warning temperature.**

Cause: A disk has reached an unacceptable temperature and received a SMART alert (predictive failure). The disk is likely to fail in the near future.

First Action: Determine why the array disk has reached an unacceptable temperature. A variety of factors can cause the excessive temperature. For example, a fan may have failed, the thermostat may be set too high, or the room temperature may be too hot or cold. Verify that the fans in the server or enclosure are working. If the array disk is in an enclosure, you should check the thermostat settings and examine whether the enclosure is located near a heat source. Make sure the enclosure has enough ventilation and that the room temperature is not too hot. Refer to the enclosure documentation for more diagnostic information.

Second Action: If you cannot identify why the disk has reached an unacceptable temperature, then replace the disk. If the array disk is a member of a non-redundant virtual disk, then back up the data before replacing the disk. Removing an array disk that is included in a non-redundant virtual disk will cause the virtual disk to fail and may cause data loss.

Name: PERC2_ALERT_SMART_WARNING_TEMPERATURE

Severity: Warning

Source: Perc2Pro (PERC Subsystem 2)

Text: %1, SMART_WARNING_TEMPERATURE on %2.

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

589

Description: **Smart warning degraded.**

Cause: A disk is degraded and has received a SMART alert (predictive failure). The disk is likely to fail in the near future.

Action: Replace the disk that has received the SMART alert. If the array disk is a member of a non-redundant virtual disk, then back up the data before replacing the disk. Removing an array disk that is included in a non-redundant virtual disk will cause the virtual disk to fail and may cause data loss.

Name: PERC2_ALERT_SMART_WARNING_DEGRADED

Severity: Warning

Source: Perc2Pro (PERC Subsystem 2)

Text: %1, SMART_WARNING_DEGRADED on %2.

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

590

Description: **Failure prediction threshold exceeded due to test – No action needed.**

Cause: A disk has received a SMART alert (predictive failure) due to test conditions.

Action: None.

Name: PERC2_ALERT_SMART_FPT_EXCEEDED_TEST

Severity: Warning

Source: Perc2Pro (PERC Subsystem 2)

Text: %1, On %2 the failure prediction threshold exceeded due to test-No action needed.

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

591

Description: **Temperature exceeded the maximum warning threshold.**

Cause: The array disk enclosure is too hot. A variety of factors can cause the excessive temperature. For example, a fan may have failed, the thermostat may be set too high, or the room temperature may be too hot.

Action: Check for factors that may cause overheating. For example, verify that the enclosure fan is working. You should also check the thermostat settings and examine whether the enclosure is located near a heat source. Make sure the enclosure has enough ventilation and that the room temperature is not too hot. Refer to the enclosure documentation for more diagnostic information.

Name: PERC_ALERT_TEMP_WARN_MAX or PERC2_ALERT_TEMP_WARN_MAX or ASPI_ALERT_TEMP_WARN_MAX

Severity: Warning

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 has exceded the maximum warning threshold.

SNMP Trap Variables: enclosureNameEv, unitNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

592

Description: **Temperature dropped below the minimum warning threshold.**

Cause: The array disk enclosure is too cool.

Action: Check whether the thermostat setting is too low and whether the room temperature is too cool.

Name: PERC_ALERT_TEMP_WARN_MIN or PERC2_ALERT_TEMP_WARN_MIN or ASPI_ALERT_TEMP_WARN_MIN

Severity: Warning

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 has exceded the minimum warning threshold.

SNMP Trap Variables: enclosureNameEv, unitNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

593

Description: Temperature exceeded the maximum error threshold.

Cause: The array disk enclosure is too hot. A variety of factors can cause the excessive temperature. For example, a fan may have failed, the thermostat may be set too high, or the room temperature may be too hot.

Action: Check for factors that may cause overheating. For example, verify that the enclosure fan is working. You should also check the thermostat settings and examine whether the enclosure is located near a heat source. Make sure the enclosure has enough ventilation and that the room temperature is not too hot. Refer to the enclosure documentation for more diagnostic information.

Name: PERC_ALERT_TEMP_ERR_MAX or PERC2_ALERT_TEMP_ERR_MAX or ASPI_ALERT_TEMP_ERR_MAX

Severity: Error

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 has exceded the maximum error threshold.

SNMP Trap Variables: enclosureNameEv, unitNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

594

Description: **Temperature dropped below the minimum error threshold.**

Cause: The array disk enclosure is too cool.

Action: Check whether the thermostat setting is too low and whether the room temperature is too cool.

Name: PERC_ALERT_TEMP_ERR_MIN or PERC2_ALERT_TEMP_ERR_MIN or ASPI_ALERT_TEMP_ERR_MIN

Severity: Error

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 has exceded the minimum error threshold.

SNMP Trap Variables: enclosureNameEv, unitNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

595

Description: **Enclosure failed.**

Cause: The enclosure, or any component within the enclosure, has failed.

Action: Identify and replace the failed component. To identify the failed component, expand the enclosure object in the Array Manager tree view to display the components (such as fans, disks, and so forth) contained in the enclosure. Any failed component will be identified with a red X. See the enclosure documentation for information on replacing enclosure components and for other diagnostic information.

Name: PERC_E_GENERIC_FAILED or PERC2_E_GENERIC_FAILED or ASPI_E_GENERIC_FAILED

Severity: Error

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 Failed.

SNMP Trap Variables: enclosureNameEv, unitNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

596

Description: Enclosure is offline.

Cause: The enclosure, or any component within the enclosure, has been removed.

Action: Identify and replace the removed component. To identify the removed component, expand the enclosure object in the Array Manager tree view to display the components (such as fans, disks, and so forth) contained in the enclosure. Any removed component will be identified with a yellow exclamation point. See the enclosure documentation for information on replacing enclosure components and for other diagnostic information.

Name: PERC_E_GENERIC_OFFLINE or PERC2_E_GENERIC_OFFLINE or ASPI_E_GENERIC_OFFLINE

Severity: Warning

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 is Offline.

SNMP Trap Variables: enclosureNameEv, unitNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

597

Description: Enclosure is in an unknown state.

Cause: There is an error in the enclosure firmware.

Action: Upgrade the enclosure firmware.

Name: PERC_E_GENERIC_UNKNOWN or PERC2_E_GENERIC_UNKNOWN or ASPI_E_GENERIC_UNKNOWN

Severity: Error

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 is in an unknown state.

SNMP Trap Variables: enclosureNameEv, unitNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

598

Description: Enclosure is in a warning state.

Cause: The array disk enclosure is either hotter or cooler than the desired temperature range.

Action: Check for factors that may cause overheating or excessive cooling. For example, verify that the enclosure fan is working. You should also check the thermostat settings and examine whether the enclosure is located near a heat source. Make sure the enclosure has enough ventilation and that the room temperature is not too hot or too cold. Refer to the enclosure documentation for more diagnostic information.

Name: PERC2_E_GENERIC_WARNING

Severity: Warning

Source: Perc2Pro (PERC Subsystem 2)

Text: %1, %2 is in a warning state.

SNMP Trap Variables: enclosureNameEv, unitNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

600

Description: Enclosure is in a degraded state.

Cause: The enclosure, or any component within the enclosure, is in a degraded state.

Action: Identify and replace the degraded component. To identify the degraded component, expand the enclosure object in the Array Manager tree view to display the components (such as fans, disks, and so forth) contained in the enclosure. Any degraded component will be identified with a yellow exclamation point. See the enclosure documentation for information on replacing enclosure components and for other diagnostic information.

Name: PERC_E_GENERIC_DEGRADED or PERC2_E_GENERIC_DEGRADED or ASPI_E_GENERIC_DEGRADED

Severity: Warning

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 is degraded.

SNMP Trap Variables: controllerNameEv, enclosureUnitNamesEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

602

Description: Enclosure was shut down.

Cause: The array disk enclosure is either hotter or cooler than the maximum allowable temperature range.

Action: Check for factors that may cause overheating or excessive cooling. For example, verify that the enclosure fan is working. You should also check the thermostat settings and examine whether the enclosure is located near a heat source. Make sure the enclosure has enough ventilation and that the room temperature is not too hot or too cold. Refer to the enclosure documentation for more diagnostic information.

Name: PERC_ALERT_SHUTDOWN_ENCLOSURE or PERC2_ALERT_SHUTDOWN_ENCLOSURE or ASPI_ALERT_SHUTDOWN_ENCLOSURE

Severity: Error

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 Shutting down.

SNMP Trap Variables: controllerNameEv, enclosureNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

603

Description: Server was shut down.

Cause: The array disk enclosure is either hotter or cooler than the maximum allowable temperature range.

Action: Check for factors that may cause overheating or excessive cooling. For example, verify that the enclosure fan is working. You should also check the thermostat settings and examine whether the enclosure is located near a heat source. Make sure the enclosure has enough ventilation and that the room temperature is not too hot or too cold. Refer to the enclosure documentation for more diagnostic information.

Name: PERC_ALERT_SHUTDOWN_SERVER or PERC2_ALERT_SHUTDOWN_SERVER or ASPI_ALERT_SHUTDOWN_SERVER

Severity: Error

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 Shutting down the server.

SNMP Trap Variables: controllerNameEv, enclosureNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

604

Description: A consistency check on a virtual disk has been paused (suspended).

Cause: The check consistency operation on a virtual disk was paused by a user.

Action: To resume the check consistency operation, right-click the virtual disk in the Array Manager tree view and select **Resume Check Consistency**.

Name: PERC2_ALERT_PAUSED_CHECK_CONSISTENCY

Severity: Warning

Source: Perc2Pro (PERC Subsystem 2)

Text: %1, Virtual Disk (%2) consistency check paused.

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

605

Description: A consistency check on a virtual disk has been resumed.

Cause: The check consistency operation on a virtual disk has resumed processing after being paused by a user.

Action: This event is provided for informational purposes.

Name: PERC2_ALERT_RESUMED_CHECK_CONSISTENCY

Severity: Information

Source: Perc2Pro (PERC Subsystem 2)

Text: %1, Virtual Disk (%2) consistency check resumed.

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

606

Description: **A virtual disk and its mirror have been split.**

Cause: A user has caused a mirrored virtual disk to be split. When a virtual disk is mirrored, its data is copied to another virtual disk in order to maintain redundancy. After being split, both virtual disks retain a copy of the data, although because the mirror is no longer intact, updates to the data are no longer copied to the mirror.

Action: This event is provided for informational purposes.

Name: PERC2_ALERT_VIRTUAL_DISK_SPLIT_MIRROR

Severity: Information

Source: Perc2Pro (PERC Subsystem 2)

Text: %1, Virtual Disk (%2) split mirror.

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

607

Description: **A mirrored virtual disk has been unmirrored.**

Cause: A user has caused a mirrored virtual disk to be unmirrored. When a virtual disk is mirrored, its data is copied to another virtual disk in order to maintain redundancy. After being unmirrored, the disk formerly used as the mirror returns to being an array disk and becomes available for inclusion in another virtual disk.

Action: This event is provided for informational purposes.

Name: PERC2_ALERT_VIRTUAL_DISK_UNMIRROR

Severity: Information

Source: Perc2Pro (PERC Subsystem 2)

Text: %1, Virtual Disk (%2) unmirrored.

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

608

Description: **A virtual disk has been renamed.**

Cause: A user has changed the name of a virtual disk.

Action: This event is provided for informational purposes.

Name: PERC_ALERT_RENAME_VD or PERC2_ALERT_RENAME_VD

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %1, %2 renamed to "%3"

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv, virtualDiskNameNewEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

609

Description: **An object is in a ready state.**

Cause: This event is generated when an object transitions to ready state.

Action: No action needed. This is an informational message.

Name: PERC_E_GENERIC_READY or PERC2_E_GENERIC_READY or ASPI_E_GENERIC_READY

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 is in a ready state.

SNMP Trap Variables: controllerNameEv, device1NameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

610

Description: **Communication has timed out.**

Cause: The controller cannot communicate with an attached device.

Action: Make sure that the device is powered on and that the cables are securely attached. See the "Cables attached correctly" section for more information on checking the cables. If communication is not restored, replace the SCSI cables.

Name: PERC_ALERT_COMM_TIMEOUT or PERC2_ALERT_COMM_TIMEOUT or ASPI_ALERT_COMM_TIMEOUT

Severity: Warning

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 SES communication timeout

SNMP Trap Variables: controllerNameEv, device1NameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

611

Description: Communication is lost.

Cause: The controller cannot communicate with an attached device.

Action: Make sure that the device is powered on and that the cables are securely attached. See the "Cables attached correctly" section for more information on checking the cables. If communication is not restored, replace the SCSI cables.

Name: PERC_ALERT_COMM_FAILURE or PERC2_ALERT_COMM_FAILURE or ASPI_ALERT_COMM_FAILURE

Severity: Warning

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 SES communication lost

SNMP Trap Variables: controllerNameEv, device1NameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

612

Description: Communication restored.

Cause: The controller has restored communication with an attached device.

Action: This event is provided for informational purposes.

Name: PERC_ALERT_COMM_FAILURE or PERC2_ALERT_COMM_FAILURE ASPI_ALERT_COMM_FAILURE

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 SES communication restored

SNMP Trap Variables: controllerNameEv, device1NameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

668

Description: **Array Manager has lost communication with this RAID controller and subsystem. A reboot is STRONGLY suggested at this point to avoid further problems. If this reboot does not resolve the problem, then this may be an indication of hardware failure.**

Cause: Array Manager has lost communication with a device. There may be faulty hardware or loose or defective cables.

Action: Reboot the system. If the problem is not resolved, check for hardware failures. Array Manager identifies failed components with a red "X" in the tree view. Any failed component must be replaced. Make sure the cables are attached securely. See the "Cables attached correctly" section for more information on checking the cables. Refer to the hardware documentation for more diagnostics information.

Name: PERC_ALERT_MEGALIB_TIMEOUT

Severity: Error

Source: PercPro (PERC Subsystem 1)

Text: %1, Array Manager has lost communication with this RAID controller and subsystem. A reboot is STRONGLY suggested at this point to avoid further problems. If this reboot does not resolve the problem, then this may be an indication of hardware failure.

SNMP Trap Variables: controllerNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on event source(s), see "Event Source."

670

Description: **A format unit operation was interrupted prior to completion. This will require a format unit operation to be run until completion for recovery of this disk.**

Cause: A low-level disk format operation was interrupted. The disk is currently in an unusable state.

Action: Run a full low-level format on the disk. The format operation must run to completion.

Name: PERC_ALERT_SCSI_SENSE_FORMAT_FAIL

Severity: Warning

Source: PercPro (PERC Subsystem 1)

Text: %1, %2 Sense Key = %3, Sense Code = %4, Sense Qualifier = %5. A format unit operation was interrupted prior to completion. This will require a format unit operation to be run until completion for recovery of this disk.

SNMP Trap Variables: controllerNameEv, diskNameEv, senseKeyEv, senseCodeEv, senseQualifierEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

671

Description: **If this disk is part of a non-redundant virtual disk, the data for this block cannot be recovered. The disk will require replacement and data restore. If this disk is part of a redundant virtual disk, the data in this block will be reallocated.**

Cause: A sector of the disk is corrupted and data cannot be maintained on this portion of the disk.

Action: If the disk is not part of a redundant virtual disk, then replace the disk. Any data residing on the corrupt portion of the disk may be lost and you may need to restore from backup. If the disk is part of a redundant virtual disk, then any data residing on the corrupt portion of the disk will be reallocated elsewhere in the virtual disk.

Name: PERC_ALERT_SCSI_SENSE_SECTOR_REASSIGN

Severity: Warning

Source: PercPro (PERC Subsystem 1)

Text: %1, %2 Sense Key = %3, Sense Code = %4, Sense Qualifier = %5. If this disk is part of a non-redundant virtual disk, the data for this block cannot be recovered. The disk will require replacement and data restore. If this disk is part of a redundant virtual disk, the data in this block will be reallocated.

SNMP Trap Variables: controllerNameEv, diskNameEv, senseKeyEv, senseCodeEv, senseQualifierEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

672

Description: EMM firmware mismatch. This is an unsupported configuration. Please download the firmware to this enclosure again to correct this condition.

Cause: The firmware on the EMM modules is not the same version. It is required that both modules have the same version of the firmware.

Action: Download the same version of the firmware to both EMM modules. For more information on downloading firmware, see "Download Firmware."

Name: PERC_ALERT_EMM_FW_MISMATCH, or PERC2_ALERT_EMM_FW_MISMATCH or ASPI_ALERT_EMM_FW_MISMATCH

Severity: Warning

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 EMM firmware mismatch (%3 / %4). This is an unsupported configuration. Please download the firmware to this enclosure again to correct this condition.

SNMP Trap Variables: controllerNameEv, enclosureNameEv, eMMFWVersion0Ev, eMMFWVersion1Ev

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

673

Description: Conservative cache mode enabled.

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_CONSERVE_CACHE_MODE_ENABLE or PERC2_ALERT_CONSERVE_CACHE_MODE_ENABLE or ASPI_ALERT_CONSERVE_CACHE_MODE_ENABLE

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 conservative cache mode enabled.

Event Variables: Controller name, enclosure name

SNMP Trap Variables: controllerNameEv, enclosureNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

674

Description: Conservative cache mode disabled.

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_CONSERVE_CACHE_MODE_DISABLE or PERC2_ALERT_CONSERVE_CACHE_MODE_DISABLE or ASPI_ALERT_CONSERVE_CACHE_MODE_DISABLE

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 conservative cache mode disabled.

Event Variables: Controller name, enclosure name

SNMP Trap Variables: controllerNameEv, enclosureNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

675

Description: Firmware downloaded successfully.

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_ENCLOSURE_FW_DOWNLOAD or PERC2_ALERT_ENCLOSURE_FW_DOWNLOAD or ASPI_ALERT_ENCLOSURE_FW_DOWNLOAD

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 firmware version %3 downloaded successfully.

Event Variables: Controller name, enclosure name, firmware version

SNMP Trap Variables: controllerNameEv, enclosureNameEv, eMMFWVersion0Ev

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

676

Description: Enclosure alarm enabled.

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_ENCLOSURE_ALARM_ENABLE or PERC2_ALERT_ENCLOSURE_ALARM_ENABLE or ASPI_ALERT_ENCLOSURE_ALARM_ENABLE

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 alarm enabled.

Event Variables: Controller name, enclosure name

SNMP Trap Variables: controllerNameEv, enclosureNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

677

Description: Enclosure alarm disabled.

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_ENCLOSURE_ALARM_DISABLE or PERC2_ALERT_ENCLOSURE_ALARM_DISABLE or ASPI_ALERT_ENCLOSURE_ALARM_DISABLE

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 alarm disabled.

Event Variables: Controller name, enclosure name

SNMP Trap Variables: controllerNameEv, enclosureNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

678

Description: **Controller alarm enabled.**

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_CONTROLLER_ALARM_ENABLE or PERC2_ALERT_CONTROLLER_ALARM_ENABLE

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %1, alarm enabled.

Event Variables: Controller name

SNMP Trap Variables: controllerNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

679

Description: Controller alarm disabled.

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_CONTROLLER_ALARM_DISABLE or PERC2_ALERT_CONTROLLER_ALARM_DISABLE

Severity: Information

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %1, alarm disabled.

Event Variables: Controller name

SNMP Trap Variables: controllerNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

680

Description: Controller rebuild rate has changed.

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_CONTROLLER_REBUILD_RATE

Severity: Information

Source: PercPro (PERC Subsystem 1)

Text: %1, rebuild rate changed to %2.

Event Variables: Controller name, new rebuild rate

SNMP Trap Variables: controllerNameEv, rebuildRateEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

681

Description: **Array disk forced online.**

Cause: This event is provided for informational purposes.

Action: None.

Name: PERC_ALERT_ARRAY_DISK_FORCE_ONLINE

Severity: Information

Source: PercPro (PERC Subsystem 1)

Text: %1, %2 forced online.

Event Variables: Controller name, array disk name

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

682

Description: **Array disk forced offline.**

Cause: A user has manually forced a physical disk in the array offline.

Action: If the array disk is part of a non-redundant virtual disk, force the array disk online. Right-click the array disk in the Array Manager tree view and select the **Online** option. Be aware that if data was written to the non-redundant virtual disk while the array disk was offline, this data may be corrupt. If the array disk is part of a redundant virtual disk, then manually remove the array disk and reinsert it. Removing and reinserting the

array disk causes the controller to automatically rebuild data from the redundant virtual disk. In this case, any data written to the redundant virtual disk while the array disk was offline will be restored. If the array disk is not displayed in the tree view after making it online, perform an I/O operation or a rescan.

Name: PERC_ALERT_ARRAY_DISK_FORCE_OFFLINE

Severity: Warning

Source: PercPro (PERC Subsystem 1)

Text: %1, %2 forced offline.

Event Variables: Controller name, array disk name

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

683

Description: Virtual Disk background initialization started.

Cause: Controllers in the PERC subystem 1 family support background initialization. This event indicates that background initialization has started for this virtual disk object.

Action: No action needed. This is an informational message.

Name: PERC_ALERT_TASK_BGI

Severity: Information

Source: PercPro (PERC Subsystem 1)

Text: %1, Virtual Disk (%2) background initialization started.

Event Variables:

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

684

Description: **Virtual disk background initialization cancelled.**

Cause: The background initialization that was previously started has been cancelled. This can happen manually—upon user request, or automatically by initiating certain commands, like Create Virtual Disk.

Action: The controller firmware automatically restarts the background initialization within approximately five minutes.

Name: PERC_ALERT_CANCEL_BGI

Severity: Warning

Source: PercPro (PERC Subsystem 1)

Text: %1, Virtual Disk (%2) background initialization cancelled.

Event Variables: Controller name, virtual disk name

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

685

Description: **Virtual disk background initialization failed.**

Cause: The background initialization that was previously started has abruptly ended unsuccessfully.

Action: The controller firmware automatically restarts the background initialization within approximately five minutes. If the background initialization fails repeatedly, check the event log to identify error conditions that may cause this failure. For example, an array disk with bad disk sectors may cause background initialization to fail. Refer to the controller documentation for more diagnostic information.

Name: PERC_ALERT_FAIL_BGI

Severity: Error

Source: PercPro (PERC Subsystem 1)

Text: %1, Virtual Disk (%2) background initialization failed.

Event Variables: Controller name, virtual disk name

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

686

Description: Virtual disk background initialization completed.

Cause: The background initialization that was previously started has successfully completed.

Action: No action needed. This is an informational message.

Name: PERC_ALERT_COMPLETE_BGI

Severity: Information

Source: PercPro (PERC Subsystem 1)

Text: %1, Virtual Disk (%2) background initialization completed.

Event Variables: Controller name, virtual disk name

SNMP Trap Variables: controllerNameEv, virtualDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

687

Description: An enclosure or one of its components is not installed

Cause: The system has attempted an operation on an enclosure or an enclosure component that is not installed. For example, the enclosure may not have a fan or temperature probe.

Action: This event is provided for informational purposes.

Name: PERC_E_GENERIC_NOT_INSTALLED or PERC2_E_GENERIC_NOT_INSTALLED or ASPI_E_GENERIC_NOT_INSTALLED

Severity: Information.

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2), AspiPro (SCSI Subsystem)

Text: %1, %2 is not installed.

Event Variables: Controller name, enclosure name and enclosure component name

SNMP Trap Variables: controllerNameEv, enclosureUnitNamesEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

688

Communication timeout.

Cause: Array Manager is unable to Perc2Pro (Perc2Pro (PERC Subsystem 2))communicate with an enclosure. The communication failure can be due to any faulty component that connects Array Manager to the enclosure. These components can include drivers, firmware, cables, or the enclosure.

Action: Examine all components that connect Array Manager to the enclosure. You may need to upgrade the firmware or drivers. Use the Array Manager Utilities to verify that the drivers and firmware are at a supported level. See "Array Manager Utilities" for more information. You can download drivers and firmware from http://support.dell.com. For information on checking the cables, see "Cables attached correctly." Make sure that the enclosure is powered on and functioning properly. You may need to check the enclosure documentation for further diagnostic information.

Name: PERC_ALERT_GENERIC_COMM_TIMEOUT

Severity: Warning

Source: PercPro (PERC Subsystem 1)

Text: %1, %2, %3 communication timeout. rc = %4, scsi_stat = %5

Event Variables: Controller name, enclosure name, type of communication problem, return code, SCSI status

SNMP Trap Variables: controllerNameEv, enclosureNameEv, commErrorTypeEv, returnCodeEv, scsiStatusEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

689

Communication lost.

Cause: Array Manager is unable to communicate with an enclosure. The communication failure can be due to any faulty component that connects Array Manager to the enclosure. These components can include drivers, firmware, cables, or the enclosure.

Action: Examine all components that connect Array Manager to the enclosure. You may need to upgrade the firmware or drivers. Use the Array Manager Utilities to verify that the drivers and firmware are at a supported level. See "Array Manager Utilities" for more information. You can download drivers and firmware from http://support.dell.com. For information on checking the cables, see "Cables attached correctly." Make sure that the enclosure is powered on and functioning properly. You may need to check the enclosure documentation for further diagnostic information.

Name: PERC_ALERT_GENERIC_COMM_LOST

Severity: Warning

Source: PercPro (PERC Subsystem 1)

Text: %1, %2, %3 communication lost. rc = %4, scsi_stat = %5

Event Variables: Controller name, enclosure name, type of communication problem, return code, SCSI status

SNMP Trap Variables: controllerNameEv, enclosureNameEv, commErrorTypeEv, returnCodeEv, scsiStatusEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

690

Rebuild completed with error(s).

Cause: See A Rebuild Completes with Errors on PERC Subsystem 1 Controllers for more information.

Action: See A Rebuild Completes with Errors on PERC Subsystem 1 Controllers for more information.

Name: PERC_ALERT_COMPLETE_W_ERROR_DISK_REBUILD

Severity: Error

Source: PercPro (PERC Subsystem 1)

Text: %1, %2 rebuild completed with error(s)

Event Variables: Controller name, array disk name

SNMP Trap Variables: controllerNameEv, arrayDiskNameEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see "Variable Substitution in the Event Messages."

For information on the event type or severity level, see "Event Type or Severity Levels."

For information on event source(s), see "Event Source."

691

"BBR Error" or "Error with Sense," or "Medium Error"

Cause: See Receive Event 691 with "BBR Error" or "Error with Sense," or "Medium Error" Message for more information.

Action: See Receive Event 691 with "BBR Error" or "Error with Sense," or "Medium Error" Message for more information.

Name: PERC2_ALERT_ADAPTER_MISC_ERROR

Severity: Error

Source: PercPro (PERC Subsystem 1), Perc2Pro (PERC Subsystem 2)

Text: %1; %2 “BBR Error” or “Error with Sense” or “Medium Error” followed by message text provided by the controller followed by “Please see the Array Manager documentation for diagnostic procedure.”

Event Variables: Controller name, additional messages

SNMP Trap Variables: controllerNameEv, additionalInfoEv

The event message text displayed in the event log is constructed from environment-specific information which is substituted for the variables shown in the **Text** field and the **SNMP Trap Variables** field. For an explanation of how a readable event message is constructed from these variables, see “Variable Substitution in the Event Messages.”

For information on the event type or severity level, see “Event Type or Severity Levels.”

For information on event source(s), see “Event Source.”

Events Generated by the PowerVault 660F Storage System

The following sections describe events generated by the fibre channel RAID controller in the PowerVault 660F storage system.

To locate an event, scroll through these sections to find the event number displayed on the Events tab of the Array Manager console or use the Search function in this online help file to search for the event message text.

700

Description: A physical disk has been placed online.

Cause: One of the following has occurred: 1) A rebuild was completed 2) A physical disk was configured 3) a manual online was done.

Action: None.

Address: ctl: 0 chn: 1 tgt: 2

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_ONLINE

For information on event source(s), see "Event Source."

702

Description: **Physical disk error found.**

Cause: A bad sector was found on the physical disk. Mechanical failure on the physical disk. Host SCSI device detected illegal instruction. Target device generated unknown phase sequence.

Action: If the problem occurs frequently, replace the physical disk. If the problem does not improve, contact your service representative.

Address: ctl: 0 chn: 1 tgt: 2

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_HARD_ERROR

For more information on resolving this event, see "Event 702, Hard disk error found."

For information on event source(s), see "Event Source."

703

Description: **Physical disk PFA condition found; this disk may fail soon.**

Cause: Physical disk predicted some future failure. External RAID logical device may have become critical.

Action: Refer to the enclosure hardware documentation.

Address: ctl: 0 chn: 1 tgt: 2

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_PFA

For more information on resolving this event, see "Event 703, Hard disk PFA condition found."

For information on event source(s), see "Event Source."

704

Description: An automatic rebuild has started.

Cause: A physical disk failed and a spare was available. A physical disk failed and no spare was available. A spare was added.

Action: None.

Address: ctl: 0 chn: 1 tgt: 2

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_AUTO_REBUILD_START

For information on event source(s), see "Event Source."

705

Description: A rebuild has started.

Cause: Client started the rebuild on user's request. User replaced the failed device and the controller started a rebuild.

Action: None.

Address: ctl: 0 chn: 1 tgt: 2

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_MANUAL_REBUILD_START

For information on event source(s), see "Event Source."

706

Description: Rebuild is over.

Cause: Rebuild completed successfully.

Action: None.

Address: ctl: 0 chn: 1 tgt: 2

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_REBUILD_DONE

For information on event source(s), see "Event Source."

707

Description: Rebuild is cancelled.

Cause: User cancelled the rebuild.

Action: Restart the rebuild if required.

Address: ctl: 0 chn: 1 tgt: 2

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_REBUILD_CANCELED

For more information on resolving this event, see "Event 707, Rebuild is cancelled."

For information on event source(s), see "Event Source."

708

Description: Rebuild stopped with error.

Cause: Due to some unknown error on the controller, rebuild failed.

Action: Try rebuild again.

Address: ctl: 0 chn: 1 tgt: 2

Severity: ERROR

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_REBUILD_ERROR

For more information on resolving this event, see "Event 708, Rebuild stopped with error."

For information on event source(s), see "Event Source."

709

Description: Rebuild stopped with error. New physical disk failed.

Cause: New physical disk failed. New physical disk may not be compatible with MDAC hardware/firmware.

Action: Replace the physical disk.

Address: ctl: 0 chn: 1 tgt: 2

Severity: ERROR

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_REBUILD_NEWDEV_FAILED

For more information on resolving this event, see "Event 709, Rebuild stopped with error: New device failed."

For information on event source(s), see "Event Source."

710

Description: Rebuild stopped because logical drive failed.

Cause: At least one or more physical disks failed in the array. Bad data table overflow.

Action: It may not be possible to recover from this error. Contact your service representative.

Address: ctl: 0 chn: 1 tgt: 2

Severity: ERROR

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_REBUILD_SYSDEV_FAILED

For more information on resolving this event, see "Event 710, Rebuild stopped because logical drive failed."

For information on event source(s), see "Event Source."

711

Description: A physical disk has failed.

Cause: A physical disk failed. A user action caused the physical disk to fail.

Action: Replace the physical disk.

Address: ctl: 0 chn: 1 tgt: 2

Severity: SERIOUS

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_DEAD

For more information on resolving this event, see "Event 711, A hard disk has failed."

For information on event source(s), see "Event Source."

712

Description: A new physical disk has been found.

Cause: A new physical disk has been added to the enclosure. A PowerVault 660F controller was added to the enclosure. The PowerVault 660F array was rebooted.

Action: None.

Address: ctl: 0 chn: 1 tgt: 2

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_FOUND

For information on event source(s), see "Event Source."

713

Description: A physical disk has been removed.

Cause: User removed an unconfigured physical disk. An unconfigured physical disk failed. A controller was removed. A controller powered off.

Action: Replace the device if needed.

Address: ctl: 0 chn: 1 tgt: 2

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_GONE

For information on event source(s), see "Event Source."

718

Description: SCSI command timeout on physical device.

Cause: Physical device has been removed. Physical device failed. Command time out value is not correct.

Action: Refer to the enclosure hardware documentation.

Address: ctl: 0 chn: 1 tgt: 2

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_COMMAND_TIMEOUT

For more information on resolving this event, see "Event 718, SCSI command timeout on hard device."

For information on event source(s), see "Event Source."

719

Description: SCSI command abort on physical disk.

Cause: User may have requested to abort the command. Firmware may have aborted the command to recover from error. The physical disk may have aborted the command.

Action: None.

Address: ctl: 0 chn: 1 tgt: 2

Severity: CRITICAL

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_COMMAND_ABORT

For information on event source(s), see "Event Source."

720

Description: SCSI command retried on physical disk.

Cause: The command may have timed out. Internal bus reset may have occurred. Device reset may have occurred.

Action: None.

Address: ctl: 0 chn: 1 tgt: 2

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_COMMAND_RETRIED

For information on event source(s), see "Event Source."

721

Description: Parity error found.

Cause: A physical device did not generate proper parity. The controller failed or did not check parity properly. Another physical device interfered. Some outside environment affected the data on the cable (for example, a radio frequency signal). Terminator is not connected. Improper termination.

Action: Refer to the enclosure hardware documentation.

Address: ctl: 0 chn: 1 tgt: 2

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_PARITY_ERROR

For more information on resolving this event, see "Event 721, Parity error found."

For information on event source(s), see "Event Source."

722

Description: Soft error found.

Cause: An error was detected by physical device and data was recovered.

Action: Run consistency check. If problem occurs frequently, replace the physical device.

Address: ctl: 0 chn: 1 tgt: 2

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_SOFT_ERROR

For information on event source(s), see "Event Source."

723

Description: Misc error found.

Cause: A physical device reported some error which does not fit in any category. Read / Write command time out. Data over run. Physical device was busy when host attempted to send command.

Action: If problem occurs frequently, replace the physical device.

Address: ctl: 0 chn: 1 tgt: 2

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_MISC_ERROR

For information on event source(s), see "Event Source."

724

Description: SCSI device reset.

Cause: Firmware has done a reset to recover from an error. User has done a reset.

Action: None.

Address: ctl: 0 chn: 1 tgt: 2

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_RESET

For information on event source(s), see "Event Source."

725

Description: Active spare found.

Cause: Physical disk was configured. Manual active spare was done. Automatic active spare was done.

Action: None.

Address: ctl: 0 chn: 1 tgt: 2

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_ACTIVESPARE

For information on event source(s), see "Event Source."

726

Description: **Warm spare found.**

Cause: Physical disk was configured. Manual warm spare was done. Automatic warm spare was done.

Action: None.

Address: ctl: 0 chn: 1 tgt: 2

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_WARMSPARE

For information on event source(s), see "Event Source."

731

Description: **Initialization canceled.**

Cause: User cancelled the operation

Action: Physical disk must be initialized again or the physical disk cannot be used.

Address: ctl: 0 chn: 1 tgt: 2

Severity: SERIOUS

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_INIT_CANCELED

For information on event source(s), see "Event Source."

749

Description: **Physical disk status changed to offline.**

Cause: A user has made an array disk offline or an array disk has failed.

Action: None.

Address: ctl: 0 chn: 1 tgt: 2

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_OFFLINE

For more information on resolving this event, see "Event 749, Physical device status changed to Offline."

For information on event source(s), see "Event Source."

750

Description: Physical disk status changed to hot spare.

Cause: A user has assigned an array disk as a hot spare.

Action: None.

Address: ctl: 0 chn: 1 tgt: 2

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_STANDBY

For information on event source(s), see "Event Source."

751

Description: Physical disk status changed to rebuild.

Cause: A hot spare has started to rebuild.

Action: None.

Address: ctl: 0 chn: 1 tgt: 2

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_REBUILD

For information on event source(s), see "Event Source."

752

Description: **Physical device ID did not match.**

Action: None.

Address: ctl: 0 chn: 1 tgt: 2

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_ID_MISMATCH

For information on event source(s), see "Event Source."

753

Description: **Physical disk failed to start.**

Cause: An array disk is not properly seated, properly inserted, or has been removed.

Action: Reseat or reinsert the array disk.

Address: ctl: 0 chn: 1 tgt: 2

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_FAILED_START

For more information on resolving this event, see "Event 753, Physical device failed to start."

For information on event source(s), see "Event Source."

754

Description: **Physical disk negotiated different offset than config.**

Action: None.

Address: ctl: 0 chn: 1 tgt: 2

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_OFFSET_SET

For information on event source(s), see “Event Source.”

755

Description: Physical disk negotiated different bus width than config.

Action: None.

Address: ctl: 0 chn: 1 tgt: 2

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: MLXEV_PHYSDEV_SET_BUS_WIDTH

For information on event source(s), see “Event Source.”

756

Description: Physical disk missing on startup.

Cause: Physical disk missing.

Action: Replace the physical disk or power on all enclosures.

Address: ctl: 0 chn: 1 tgt: 2

Severity: SERIOUS

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_MISSING_ONSTARTUP

For more information on resolving this event, see “Event 756, Physical drive missing on startup.”

For information on event source(s), see “Event Source.”

758

Description: Physical disk is switching from one channel to the other channel.

Cause: Physical disk removed or channel failed.

Action: Check FC loop; replace physical disk if necessary.

Address: ctl: 0 chn: 1 tgt: 2

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_MOVING_TO_OTHER_CHN

For more information on resolving this event, see "Event 758, Physical drive is switching from one channel to another."

For information on event source(s), see "Event Source."

761

Description: **Device loop ID conflict (soft addressing) detected.**

Cause: Device loop ID conflict detected on disk channel resulting in soft addressing. Potential data corruption.

Action: Change index selector to enable hard addressing per enclosure manufacturer's specification.

Address: ctl: 0 chn: 1

Severity: SERIOUS

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_FIBREDEV_LOOPID_SOFTADDR_OCCURRED

For information on event source(s), see "Event Source."

762

Description: **Consistency check is started.**

Cause: User started a consistency check. The controller started a consistency check.

Action: None.

Address: ctl: 0 logdrv: 1

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SYSDEV_CHECK_START

For information on event source(s), see "Event Source."

763

Description: Consistency check is finished.

Cause: Consistency check completed successfully without detecting any errors.

Action: None.

Address: ctl: 0 logdrv: 1

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SYSDEV_CHECK_DONE

For information on event source(s), see "Event Source."

764

Description: Consistency check is cancelled.

Cause: User cancelled the consistency check.

Action: Restart consistency check, if required.

Address: ctl: 0 logdrv: 1

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SYSDEV_CHECK_CANCELED

For information on event source(s), see "Event Source."

765

Description: Consistency check on logical drive error.

Cause: Inconsistent data was found. Bad sectors were found. A physical disk reliability problem.

Action: See bad block and request sense table for more information.

Address: ctl: 0 logdrv: 1

Severity: ERROR

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SYSDEV_CHECK_ERROR

For more information on resolving this event, see "Event 765, Consistency check on logical drive error."

For information on event source(s), see "Event Source."

766

Description: **Consistency check on logical drive failed.**

Cause: A logical device became critical. A logical device failed.

Action: See request sense data for more information.

Address: ctl: 0 logdrv: 1

Severity: ERROR

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SYSDEV_CHECK_SYSDEV_FAILED

For more information on resolving this event, see "Event 766, Consistency check on logical drive failed."

For information on event source(s), see "Event Source."

767

Description: **Consistency check failed due to physical disk failure.**

Cause: A physical disk failed.

Action: See request sense data for more information.

Address: ctl: 0 logdrv: 1

Severity: SERIOUS

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SYSDEV_CHECK_PHYSDEV_FAILED

For more information on resolving this event, see "Event 767, Consistency check failed due to physical device failure."

For information on event source(s), see "Event Source."

768

Description: Logical drive has been made offline.

Cause: One / multiple physical disk(s) failed.

Action: It may not be possible to recover from this error. Contact your service representative.

Address: ctl: 0 logdrv: 1

Severity: SERIOUS

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SYSDEV_OFFLINE

For more information on resolving this event, see "Event 768, Logical drive has been made offline."

For information on event source(s), see "Event Source."

769

Description: Logical drive is critical.

Cause: One physical device failed.

Action: Replace the physical device. Start the rebuild, if required.

Address: ctl: 0 logdrv: 1

Severity: ERROR

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SYSDEV_CRITICAL

For more information on resolving this event, see "Event 769, Logical drive is Critical."

For information on event source(s), see "Event Source."

770

Description: Logical drive has been placed online.

Cause: Rebuild completed. User set the physical disk online. New configuration was added.

Action: None.

Address: ctl: 0 logdrv: 1

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SYSDEV_ONLINE

For information on event source(s), see "Event Source."

778

Description: Logical drive initialization started.

Cause: User started the initialization. Any previous data is lost.

Action: None.

Address: ctl: 0 logdrv: 1

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SYSDEV_INIT_STARTED

For information on event source(s), see "Event Source."

779

Description: Logical drive initialization done.

Cause: Initialize operation completed successfully.

Action: None.

Address: ctl: 0 logdrv: 1

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SYSDEV_INIT_DONE

For information on event source(s), see "Event Source."

780

Description: Logical drive initialization canceled.

Cause: User cancelled the initialization.

Action: Restart initialization if required.

Address: ctl: 0 logdrv: 1

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SYSDEV_INIT_CANCELED

For information on event source(s), see "Event Source."

781

Description: Logical drive initialization failed.

Cause: One / multiple physical device(s) failed. Controller has been removed. Controller has been powered off.

Action: Refer to the device failure event.

Address: ctl: 0 logdrv: 1

Severity: ERROR

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SYSDEV_INIT_FAILED

For more information on resolving this event, see "Event 781, Logical drive initialization failed."

For information on event source(s), see "Event Source."

784

Description: Expand capacity started.

Cause: User started either an Expand Capacity or Add Virtual Disk operation on a disk group.

Action: None.

Address: ctl: 0 logdrv: 1

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SYSDEV_EXPANDCAPACITY_START

For information on event source(s), see "Event Source."

785

Description: **Expand capacity completed.**

Cause: An Expand Capacity or Add Virtual Disk operation on a disk group completed.

Action: None.

Address: ctl: 0 logdrv: 1

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SYSDEV_EXPANDCAPACITY_DONE

For information on event source(s), see "Event Source."

786

Description: **Expand capacity stopped with error.**

Cause: Multiple physical disks failed during an Expand Capacity or Add Virtual Disk operation.

Action: It may not be possible to recover from this error. Contact your service representative.

Address: ctl: 0 logdrv: 1

Severity: ERROR

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SYSDEV_EXPANDCAPACITY_ERROR

For more information on resolving this event, see "Event 786, Expand Capacity stopped with error."

For information on event source(s), see "Event Source."

787

Description: Bad Blocks found.

Cause: Bad sector was found on an array disk during a Consistency Check, Rebuild, Expand Capacity or Add Virtual Disk operation.

Action: Run a consistency check with the restore option. Restore data from a back up.

Address: ctl: 0 logdrv: 1

Severity: CRITICAL

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SYSDEV_BADBLOCK

For more information on resolving this event, see "Event 787, Bad blocks found."

For information on event source(s), see "Event Source."

788

Description: Logical drive size changed.

Cause: A new configuration has been added. Expand Capacity has added extra capacity to the virtual disk.

Action: None.

Address: ctl: 0 logdrv: 1

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SYSDEV_SIZECHANGED

For information on event source(s), see "Event Source."

789

Description: Logical drive type changed.

Cause: A new configuration has been added. RAID migration completed. Expand Capacity or Add Virtual Disk was completed on RAID 1.

Action: None.

Address: ctl: 0 logdrv: 1

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SYSDEV_TYPECHANGED

For information on event source(s), see "Event Source."

791

Description: **Logical drive LUN mapping has been written to config.**

Cause: A new virtual disk was created and made available to a server. This could occur either during the creation of virtual disk, or when the virtual disk is assigned to a server by way of Dell OpenManage Storage Consolidation.

Action: This is a normal operation when a virtual disk is created or assigned.

Address: ctl: 0

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SYSDEV_WR_LUN_MAP

For information on event source(s), see "Event Source."

792

Description: **Attempt to read data from block that is marked in Bad Data Table.**

Cause: Attempt to read from block that is already marked bad in Bad Data Table. Potential data loss.

Action: Restore data from a back up.

Address: ctl: 0 logdrv: 1

Severity: SERIOUS

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SYSDEV_DATAREAD_FROM_BLOCK_IN_BDT

For information on event source(s), see "Event Source."

793

Description: Data for Disk Block has been lost due to Logical Drive problem.

Cause: Data retained in RAID Cache for a Write-back Logical Drive cannot be stored to the physical medium be**Cause:** of Logical Drive problem. The Logical Drive problem could be be**Cause:** of multiple physical medium error or multiple physical devices offline or other reasons.

Action: Insure that all the physical disks related to the Logical Drive, disk channel, enclosure or cabling are functional and accessible. Repair or replace as necessary.

Address: ctl: 0 logdrv: 1

Severity: ERROR

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SYSDEV_DATA_FOR_BLOCK_LOST

For information on event source(s), see “Event Source.”

799

Description: Fan failure.

Cause: Cable connection is broken. Fan failure.

Action: Replace fan.

Address: ctl: 0 chn: 1 tgt: 2 Fan: 3

Severity: SERIOUS

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_AEMI_FAN_FAILED

For information on event source(s), see “Event Source.”

803

Description: Power supply failure.

Cause: Cable connection is broken. Power supply failure.

Action: Replace power supply.

Address: ctl: 0 chn: 1 tgt: 2 Power Supply: 3

Severity: SERIOUS

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_AEMI_PWR_SUPPLY_FAILED

For information on event source(s), see "Event Source."

804

Refer to volume management event "10804."

805

Refer to volume management event "10805."

808

Description: **Over temperature.**

Cause: Room temperature is too high. Fan failure.

Action: Turn off the system and allow it to cool down. Adjust the room temperature.

Address: ctl: 0 chn: 1 tgt: 2 Heat Sensor: 3

Severity: SERIOUS

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_AEMI_OVER_TEMPERATURE

For information on event source(s), see "Event Source."

818

Description: **Fan failure.**

Cause: Cable connection is broken. Fan failure.

Action: Replace fan.

Address: ctl: 0 enclosure: 1 unit: 2

Severity: SERIOUS

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_ENCLFAN_FAILED

For more information on resolving this event, see “Event 818 and 820, Fan failure or Fan is not present.”

For information on event source(s), see “Event Source.”

819

Description: Fan has been restored.

Cause: Faulty fan has been replaced. Cable is connected properly.

Action: None.

Address: ctl: 0 enclosure: 1 unit: 2

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_ENCLFAN_OK

For information on event source(s), see “Event Source.”

820

Description: Fan is not present.

Cause: Enclosure management connection is broken. Management hardware failure. Fan is not present.

Action: Refer to the enclosure hardware documentation.

Address: ctl: 0 enclosure: 1 unit: 2

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_ENCLFAN_NOTPRESENT

For more information on resolving this event, see “Event 818 and 820, Fan failure or Fan is not present.”

For information on event source(s), see “Event Source.”

821

Description: **Power supply failure.**

Cause: Cable connection is broken. Power supply failure.

Action: Refer to the enclosure hardware documentation.

Address: ctl: 0 enclosure: 1 unit: 2

Severity: SERIOUS

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_ENCLPOWER_FAILED

For more information on resolving this event, see "Event 821 and 823, Power supply failure or Power supply is not present."

For information on event source(s), see "Event Source."

822

Description: **Power supply has been restored.**

Cause: Faulty power supply has been replaced.

Action: None.

Address: ctl: 0 enclosure: 1 unit: 2

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

For information on event source(s), see "Event Source."

Name: DAM_MLXEV_ENCLPOWER_OK

823

Description: **Power supply is not present.**

Cause: Management connection is broken. Management hardware failure. Power supply is not present.

Action: Refer to the enclosure hardware documentation.

Address: ctl: 0 enclosure: 1 unit: 2

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_ENCLPOWER_NOTPRESENT

For more information on resolving this event, see "Event 821 and 823, Power supply failure or Power supply is not present."

For information on event source(s), see "Event Source."

824

Description: Temperature is over safe limit. Failure imminent.

Cause: Room temperature is too high. Fan failure. Sensor failure.

Action: Turn off the system and allow it to cool down. Adjust the room temperature.

Address: ctl: 0 enclosure: 1 unit: 2

Severity: SERIOUS

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_ENCLHEAT_BAD

For information on event source(s), see "Event Source."

825

Description: Temperature is above working limit.

Cause: Room temperature is too high. Fan failure.

Action: Adjust the room temperature. Replace fan.

Address: ctl: 0 enclosure: 1 unit: 2

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_ENCLHEAT_CRITICAL

For more information on resolving this event, see "Event 825, Temperature is above working limit."

For information on event source(s), see "Event Source."

826

Description: **Normal temperature has been restored.**

Cause: Faulty fan has been replaced. Room temperature was reduced.

Action: None.

Address: ctl: 0 enclosure: 1 unit: 2

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_ENCLHEAT_OK

For information on event source(s), see "Event Source."

827

Description: **Temperature sensor is not present.**

Cause: Enclosure management connection is broken. Management hardware failure. Sensor is not present.

Action: Refer to the enclosure hardware documentation.

Address: ctl: 0 enclosure: 1 unit: 2

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_ENCLHEAT_NOTPRESENT

For information on event source(s), see "Event Source."

828

Description: **Enclosure access critical.**

Cause: Problems were encountered. Internal communication to enclosure management module was not successful. Possible enclosure management hardware failure.

Action: Refer to the enclosure hardware documentation.

Address: ctl: 0 enclosure: 1

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_ENCLACCESS_CRITICAL

For more information on resolving this event, see "Event 828, Enclosure access is Critical."

For information on event source(s), see "Event Source."

829

Description: Enclosure access has been restored.

Cause: Enclosure has been fixed or replaced.

Action: None.

Address: ctl: 0 enclosure: 1

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_ENCLACCESS_OK

For information on event source(s), see "Event Source."

830

Description: Enclosure access is offline.

Cause: Problems were encountered with internal communication to the enclosure management module.

Action: Follow enclosure management vendor's diagnostics and repair procedures.

Address: ctl: 0 enclosure: 1 unit: 2

Severity: SERIOUS

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_ENCLACCESS_OFFLINE

For information on event source(s), see "Event Source."

831

Description: Enclosure soft addressing detected.

Cause: Enclosure has duplicate loop IDs (soft addressing). Potential data corruption.

Action: Refer to the enclosure hardware documentation.

Address: ctl: 0 enclosure: 1

Severity: SERIOUS

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_ENCLSES_SOFTADDR_OCCURRED

For more information on resolving this event, see "Event 831, Enclosure Soft Addressing detected."

For information on event source(s), see "Event Source."

832

Description: Enclosure services ready.

Cause: The enclosure is available for management. This event usually occurs after reboot or bootup of the PowerVault 660F.

Action: None.

Address: ctl: 0

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_ENCLACCESS_READY

For information on event source(s), see "Event Source."

833

Description: Access to temperature sensor has been lost.

Cause: Module containing temperature sensor has been removed.

Action: None.

Address: ctl: 0 enclosure: 1 unit: 2

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_ENCLHEAT_UNKNOWN

For information on event source(s), see "Event Source."

834

Description: Access to power supply status information has been lost.

Cause: Module establishing connectivity has been removed.

Action: None.

Address: ctl: 0 enclosure: 1 unit: 2

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_ENCLPOWER_UNKNOWN

For information on event source(s), see "Event Source."

835

Description: Access to fan status information has been lost.

Cause: Module establishing connectivity has been removed.

Action: None.

Address: ctl: 0 enclosure: 1 unit: 2

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_ENCLFAN_UNKNOWN

For information on event source(s), see "Event Source."

836

Description: Array management server software started successfully.

Cause: The server system (or array management utility server) started.

Action: If you did not expect a system reboot, investigate.

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SYSTEM_STARTED

For information on event source(s), see "Event Source."

837

Description: **Write back error.**

Cause: Data cache write failed.

Action: The data may have been lost. Restore the data from a backup.

Address: ctl: 0

Severity: ERROR

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_WRITEBACK_ERROR

For information on event source(s), see "Event Source."

839

Description: **Controller is dead. System is disconnecting from this controller.**

Cause: Hardware failure.

Action: Contact your service representative.

Address: ctl: 0

Severity: CRITICAL

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_DEAD

For information on event source(s), see "Event Source."

840

Description: **Controller has been reset.**

Cause: Controller received a reset command.

Action: If this was an unexpected event, refer to the enclosure hardware documentation.

Address: ctl: 0

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_RESET

For more information on resolving this event, see “Event 840, Controller has been reset.”

For information on event source(s), see “Event Source.”

841

Description: Controller is found.

Cause: New controller has been installed. Management software has restarted. System has rebooted.

Action: None.

Address: ctl: 0

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_FOUND

For information on event source(s), see “Event Source.”

842

Description: Controller is gone. System is disconnecting from this controller.

Cause: Controller has been powered off or has been reset. Controller has been removed from the system.

Action: None.

Address: ctl: 0

Severity: CRITICAL

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_GONE

For information on event source(s), see “Event Source.”

843

Description: Battery present.

Cause: A battery module has been detected and is now operational.

Action: None.

Address: ctl: 0

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_BBU_FOUND

For information on event source(s), see "Event Source."

844

Description: Battery power low.

Cause: Battery power is low. This normally happens during a battery calibration or recondition cycle. If this is not the reason, then investigate further.

Action: If this message occurs without a power failure or battery recondition / calibration, replace the battery.

Address: ctl: 0

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_BBU_POWER_LOW

Note While the battery is low, the controller write cache will be disabled.

For more information on resolving this event, see "Event 844, BBU power low."

For information on event source(s), see "Event Source."

845

Description: Battery power OK.

Cause: Battery has enough power to enable the write data cache.

Action: None.

Address: ctl: 0

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_BBU_POWER_OK

For information on event source(s), see "Event Source."

846

Description: Controller is gone. System is disconnecting from this controller.

Cause: The connection to the controller has been lost.

Action: None.

Address: ctl: 0

Severity: CRITICAL

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_POWER_OFF

For information on event source(s), see “Event Source.”

847

Description: Controller powered on

Cause: Controller was removed from the system. Controller has been powered off.

Action: None.

Address: ctl: 0

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_POWER_ON

For information on event source(s), see “Event Source.”

848

Description: Controller is online.

Cause: New controller has been installed.

Action: None.

Address: ctl: 0

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_ONLINE

For information on event source(s), see "Event Source."

849

Description: **Controller is gone. System is disconnecting from this controller.**

Cause: Controller was set online.

Action: None.

Address: ctl: 0

Severity: CRITICAL

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_OFFLINE

For information on event source(s), see "Event Source."

850

Description: **Controller's partner is gone, controller is in failover mode now.**

Cause: Controller was set offline.

Action: If you did not expect this, investigate.

Address: ctl: 0

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_CRITICAL

For information on event source(s), see "Event Source."

851

Description: **Battery reconditioning is started.**

Cause: User started a battery reconditioning.

Action: None.

Address: ctl: 0

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_BBU_RECOND_START

For information on event source(s), see “Event Source.”

852

Description: Battery reconditioning is finished.

Cause: Battery reconditioning completed successfully.

Action: None.

Address: ctl: 0

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_BBU_RECOND_DONE

For information on event source(s), see “Event Source.”

853

Description: Battery reconditioning is canceled.

Cause: User cancelled the battery reconditioning.

Action: Restart the battery reconditioning, if required.

Address: ctl: 0

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_BBU_RECOND_ABORT

For information on event source(s), see “Event Source.”

855

Description: Controller firmware mismatch.

Cause: Replacement controller with downlevel firmware installed.

Action: Reload controller firmware.

Address: ctl: 0 chn: 1 tgt: 2

Severity: SERIOUS

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_FIRMWARE_MISMATCH

For information on event source(s), see "Event Source."

856

Description: Battery removed.

Cause: Battery physically removed.

Action: Reinstall the battery.

Address: ctl: 0

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_BBU_NORESPONSE

For information on event source(s), see "Event Source."

857

Description: WARM BOOT failed.

Cause: Memory error detected during WARM boot scan. Possible data loss.

Action: Restore data from a back up.

Address: ctl: 0

Severity: SERIOUS

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_WARM_BOOT_ERROR

For more information on resolving this event, see "Event 857, Warm boot failed."

For information on event source(s), see "Event Source."

858

Description: Controller entered conservative cache mode.

Cause: User changed the setting. Controller preparing for shutdown. Battery low. RAID expansion in process. Over-temp condition. No redundant power. Power supply failure. UPS ac failed. UPS battery low.

Action: If this was an unexpected event, refer to the enclosure hardware documentation.

Address: ctl: 0 param: 0x0001

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_CONSERV_CACHE_MODE

Paramtype bitfield

Params

1: **Name:** _FAULT_MANAGEMENT: “Fault Management”

2: **Name:** _INPUT_SIG: “Hardware Input Signal”

4: **Name:** _SHUT_DOWN: “Controller Shut Down”

8: **Name:** _BBU_LOW: “BBU Battery Low”

16: **Name:** _MORE: “More Active”

32: **Name:** _OVER_TEMP: “Over Temperature”

64: **Name:** _NON_REDUN_POWER: “Non-redundant Power”

128: **Name:** _POWER_SUPPLY_FAIL: “Power Supply Failed”

For more information on resolving this event, see “Event 858, Controller entered Conservative Cache mode.”

For information on event source(s), see “Event Source.”

859

Description: Controller entered normal cache mode.

Cause: Controller operation occurred which could have changed the cache mode. This may or may not indicate a change in controller cache mode.

Action: None.

Address: ctl: 0

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_NORMAL_CACHE_MODE

For information on event source(s), see "Event Source."

860

Description: **Controller device start complete.**

Action: None.

Address: ctl: 0

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_DEV_START_CMPLT

For information on event source(s), see "Event Source."

861

Description: **Soft ECC error corrected.**

Cause: Faulty memory module.

Action: Replace memory module.

Address: ctl: 0

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_SOFT_ECC_CORRECTED

For information on event source(s), see "Event Source."

862

Description: **Hard ECC error corrected.**

Cause: Faulty memory module.

Action: Replace memory module.

Address: ctl: 0 chn: 1 tgt: 2

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_HARD_ECC_CORRECTED

For information on event source(s), see "Event Source."

863

Description: Battery recondition suggested.

Cause: The battery in the BBU may need to be reconditioned to improve full capacity accuracy.

Action: Recondition the battery backup unit. See "Recondition" for more information.

Address: ctl: 0

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_BBU_RECOND_NEEDED

For more information on resolving this event, see "Event 863, BBU Recondition suggested."

For information on event source(s), see "Event Source."

864

Description: Controller's partner has been removed.

Cause: The controller's partner has been disabled. This may occur when performing an expand capacity operation.

Action: None.

Address: ctl: 0

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_REMOVED_PTNR

For information on event source(s), see "Event Source."

866

Description: Updated partner's status.

Cause: The controller's partner has changed its operational state.

Action: If the new state is unexpected, view earlier events to determine why it changed state.

Address: ctl: 0 param: 0x0001

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_UPDATE_PTNR_STATUS

Paramtype enum

Params

0x0000: **Name:** _UNKNOWN: "Unknown"

0x0100: **Name:** _BOOTING: "Booting"

0x0200: **Name:** _ACTIVE: "Active"

0x0300: **Name:** _FAILED: "Failed"

0x0400: **Name:** _REMOVED: "Removed"

0x0500: **Name:** _INSERTED: "Inserted"

For information on event source(s), see "Event Source."

867

Description: Relinquished partner.

Cause: A user has enabled the controller's partner.

Action: None.

Address: ctl: 0

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_RELINQUISH_PTNR

For information on event source(s), see "Event Source."

868

Description: Inserted partner.

Cause: A user has inserted a partner controller into the enclosure.

Action: None.

Address: ctl: 0

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_INSERTED_PTNR

For information on event source(s), see "Event Source."

869

Description: Dual controllers enabled.

Cause: The controller has entered duplex mode.

Action: None.

Address: ctl: 0

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_DUAL_ENABLED

For information on event source(s), see "Event Source."

870

Description: Disabled partner.

Cause: The partner controller has been manually disabled by way of Array Manager. A failure occured causing one of the controllers to terminate. See param code for details on cause.

Action: Refer to the enclosure hardware documentation.

Address: ctl: 0 param: 0x0001

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_KILL_PTNR

Paramtype enum

Params

0x0300: **Name:** _PING_TIMEOUT: “ping timeout”

0x0301: **Name:** _NEGOTIATE_GETCHUNK: “couldn’t allocate chunk of memory”

0x0302: **Name:** _NEGOTIATE_CABLE: “SCSI communication failed or cables are wrong”

0x0303: **Name:** _NEGOTIATE_HOSTID: “host ID jumper mismatch”

0x0304: **Name:** _NEGOTIATE_JUMPERS: “SLIP / DIFFL / FBR jumper mismatch”

0x0305: **Name:** _NEGOTIATE_DISKCHNS: “number of disk channels present mismatch”

0x0306: **Name:** _NEGOTIATE_HOSTCHNS: “number of host channels present mismatch”

0x0307: **Name:** _NEGOTIATE_FWVER: “FW version mismatch”

0x0302: **Name:** _NEGOTIATE_FWTYPE: “FW header type mismatch”

0x0309: **Name:** _NEGOTIATE_MEMSIZE: “memory size mismatch”

0x030a: **Name:** _NEGOTIATE_MEMRD: “memory read of partner failed”

0x030b: **Name:** _NEGOTIATE_CACHEMEMSIZE: “Cache memory size mismatch”

0x030c: **Name:** _COMMAND: “received Kill partner command”

0x030d: **Name:** _DURING_FAILBACK: “failure handing over TID”

0x030e: **Name:** _ENTERING_NEXUS: “negotiation finished, but nexus not entered in time”

0x030f: **Name:** _UNKNOWN: “failed for unknown reason”

0x0310: **Name:** _WBSYNC_CHAN0: “write-back sync to partner - chan 0”

0x0311: **Name:** _WBSYNC_CHAN1: “write-back sync to partner - chan 1”

0x0312: **Name:** _WBSYNC_CHAN2: “write-back sync to partner - chan 2”

0x0313: **Name:** _WBSYNC_CHAN3: “write-back sync to partner - chan 3”

0x0314: **Name:** _WBSYNC_CHAN4: “write-back sync to partner - chan 4”

0x0315: **Name:** _WBSYNC_CHAN5: “write-back sync to partner - chan 5”

0x0316: **Name:** _NEGOTIATE_FWBUILD: “FW build mismatch”

0x0317: **Name:** _NEGOTIATE_CABLE_CROSS: “device cables are crossed”

0x0320: **Name:** _FAILED_HOTPULL: “partner removal detect while nexus active”

0x0321: **Name:** _ABSENT_AT_BOOT: “partner missing at negotiation time”

0x0322: **Name:** _POWERFAIL_DURING_FAIL: “BBU - we power failed before fail-over finished”

0x0323: **Name:** _POWERFAIL_DURING_RELINQ: “BBU - we power failed before relinquish finished”

0x0340: **Name:** _LOCK_TIMEOUT: “Lock Timeout”

0x0341: **Name:** _LOCK_SCSI_FAIL: “Lock SCSI Failed”

0x0342: **Name:** _FOR_SIMPLEX_MIGRATION: “kill partner so we can do migrate recovery”

0x0343: **Name:** _FAILED_CTOC: “General ctoc message failure”

For more information on resolving this event, see “Event 870, Killed partner.”

For information on event source(s), see “Event Source.”

871

Description: Dual controllers entered nexus.

Cause: The array was booted up. This event means that the two controllers in a PowerVault 660F are now in sync with each other.

Action: None.

Address: ctl: 0

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_NEXUS

For information on event source(s), see “Event Source.”

872

Description: Controller boot ROM image needs to be reloaded.

Cause: Wrong firmware image file downloaded. MAC address changed.

Action: Contact your service representative to determine the failure.

Address: ctl: 0

Severity: SERIOUS

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_BAD_BOOTROM_IMAGE

For more information on resolving this event, see "Event 872, Controller boot ROM image needs to be reloaded."

For information on event source(s), see "Event Source."

873

Description: **Controller is using default non-unique world-wide Name.**

Cause: MAC

Address: lost or not set.

Action: Contact your service representative to determine the failure.

Address: ctl: 0

Severity: CRITICAL

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_BAD_MAC_ADDRESS

For more information on resolving this event, see "Event 873, Controller is using default nonunique world-wide name."

For information on event source(s), see "Event Source."

881

Description: **Lost connection to server, or server is down.**

Cause: Lost network connection to server. Server shutdown.

Action: None.

Severity: SERIOUS

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SYSTEM_DEAD

For information on event source(s), see "Event Source."

882

Description: Automatic reboot count has changed.

Cause: Controller has rebooted. Automatic reboot has rearmed itself or was reconfigured.

Action: None.

Address: ctl: 0 param: 0x0001

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_AUTOBOOT_CHANGED

Paramtype value

For information on event source(s), see “Event Source.”

883

Description: Channel Failed.

Cause: Cable disconnected.

Action: Plug in cable.

Address: ctl: 0 chn: 1

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CHANNEL_FAILED

For information on event source(s), see “Event Source.”

884

Description: Channel Online.

Cause: Cable reconnected.

Action: None.

Address: ctl: 0 chn: 1

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CHANNEL_OK

For information on event source(s), see "Event Source."

887

Description: Back end fibre dead.

Cause: Lost access to data on Fibre Channel.

Action: Refer to the enclosure hardware documentation.

Address: ctl: 0 chn: 1

Severity: SERIOUS

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CHANNEL_FIBER_DEAD

For more information on resolving this event, see "Event 887, Back End Fibre Dead."

For information on event source(s), see "Event Source."

888

Description: Back end fibre alive.

Cause: Regained access to data on internal fibre channel.

Action: None.

Address: ctl: 0 chn: 1

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CHANNEL_FIBER_ALIVE

For information on event source(s), see "Event Source."

889

Description: Event log empty.

Cause: Tried to read past last entry of controller's event log.

Action: None.

Address: ctl: 0

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_LOG_EMPTY

For information on event source(s), see “Event Source.”

890

Description: Event log entries lost.

Cause: Tried to read an entry that does not exist in the event log from the controller.

Action: None.

Address: ctl: 0

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_LOG_OUT_SYNC

For information on event source(s), see “Event Source.”

891

Description: Request Sense

Cause: A physical disk has generated an error.

Action: Interpret the Key / ASC / ASCQ and take appropriate action.

Address: ctl: 0 chn: 1 tgt: 2

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_LOG_REQUEST_SENSE

For information on event source(s), see “Event Source.”

892

Description: **Set real time clock.**

Cause: Real time clock was set.

Action: None.

Address: ctl: 0 chn: 1

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_LOG_SET_RTC

For information on event source(s), see "Event Source."

893

Description: **New configuration received.**

Cause: The configuration on the controller changed. Deleting or creating a virtual disk can cause this event to be logged.

Action: None.

Address: ctl: 0

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CFG_NEW

For information on event source(s), see "Event Source."

894

Description: **Configuration cleared.**

Cause: The controller configuration was cleared. All virtual disks and hot spares were removed. This is logged when the clear configuration operation is done.

Action: None.

Address: ctl: 0

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CFG_CLEAR

For information on event source(s), see “Event Source.”

895

Description: Configuration invalid.

Cause: The controller found an invalid configuration.

Action: None.

Address: ctl: 0

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CFG_INVALID

For information on event source(s), see “Event Source.”

896

Description: Configuration on disk access error.

Cause: The controller could not read the configuration off of the disk.

Action: None.

Address: ctl: 0

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CFG_COD_ACCESS_ERROR

For information on event source(s), see “Event Source.”

897

Description: Configuration on disk converted.

Cause: The controller converted a down level configuration on disk.

Action: None.

Address: ctl: 0

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CFG_COD_CONVERTED

For information on event source(s), see "Event Source."

898

Description: **Configuration on disk import failed.**

Cause: The controller could not import the configuration.

Action: None.

Address: ctl: 0

Severity: WARNING

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CFG_COD_IMPORT_FAILED

For information on event source(s), see "Event Source."

900

Description: **A debug dump exists on this system.**

Cause: The partner controller aborted and created debug dump information.

Action: Contact field support for assistance in retrieving the data.

Address: ctl: 0 Abort Code: 0x0102

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_DEBUG_DUMP_GENERATED_PARTNER

For information on event source(s), see "Event Source."

905

Description: Logical drive background initialization started.

Cause: User may have started background initialization. Firmware may have automatically started background initialization.

Action: None.

Address: ctl: 0 logdrv: 1

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SYSDEV_BG_INIT_STARTED

For information on event source(s), see "Event Source."

906

Description: Logical drive background initialization stopped.

Cause: User may have stopped background initialization. Firmware may have automatically stopped background initialization.

Action: None.

Address: ctl: 0 logdrv: 1

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SYSDEV_BG_INIT_STOPPED

For information on event source(s), see "Event Source."

907

Description: Logical drive background initialization paused.

Cause: Background initialization paused due to higher priority operation.

Action: None.

Address: ctl: 0 logdrv: 1

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SYSDEV_BG_INIT_PAUSED

For information on event source(s), see "Event Source."

908

Description: Logical drive background initialization restarted.

Cause: Background initialization restarted after being paused.

Action: None.

Address: ctl: 0 logdrv: 1

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SYSDEV_BG_INIT_RESTARTED

For information on event source(s), see "Event Source."

909

Description: Logical drive background initialization failed.

Cause: Background initialization failed.

Action: None.

Address: ctl: 0 logdrv: 1

Severity: ERROR

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SYSDEV_BG_INIT_FAILED

For information on event source(s), see "Event Source."

910

Description: Logical drive background initialization completed.

Cause: Background initialization completed successfully.

Action: None.

Address: ctl: 0 logdrv: 1

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SYSDEV_BG_INIT_COMPLETED

For information on event source(s), see “Event Source.”

911

Description: Battery calibration cycle started.

Cause: New battery detected.

Action: None.

Address: ctl: 0

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_BBU_CALIBRATE_START

For information on event source(s), see “Event Source.”

912

Description: Battery calibration cycle finished.

Cause: Battery calibration completed successfully.

Action: None.

Address: ctl: 0

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_BBU_CALIBRATE_DONE

For information on event source(s), see “Event Source.”

913

Description: Battery calibration cycle was cancelled.

Cause: User cancelled the battery calibration cycle.

Action: None.

Address: ctl: 0

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_BBU_CALIBRATE_ABORT

For information on event source(s), see "Event Source."

914

Description: **Battery is not present.**

Cause: The battery electronics are present but a battery was not detected.

Action: Install or connect the battery.

Address: ctl: 0

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_BBU_NO_BATTERY

For information on event source(s), see "Event Source."

919

Description: **Dual-active automatic flash of replacement controller.**

Cause: Replacement controller is being flashed with survivor's firmware.

Action: None.

Address: ctl: 0

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_AUTOMATIC_FLASH_STARTED

For information on event source(s), see "Event Source."

920

Description: **Dual-active negotiation failed jumpers.**

Cause: Jumper settings do not match.

Action: Replace, adjust or add components so that the controllers have the same jumper settings.

Address: ctl: 0

Severity: ERROR

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_NEGOTIATION_FAILED_JUMPERS

For information on event source(s), see “Event Source.”

921

Description: **Dual-active negotiation failed IDs.**

Cause: Controller IDs not unique. Both controllers have same ID.

Action: Replace, adjust or add components so that the controllers have the correct IDs.

Address: ctl: 0

Severity: ERROR

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_NEGOTIATION_SAME_ID

For information on event source(s), see “Event Source.”

922

Description: **Dual-active negotiation failed board types.**

Cause: Board types differ. Both controllers must be of the same type.

Action: Replace appropriate controller so that the board types are the same.

Address: ctl: 0

Severity: ERROR

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_NEGOTIATION_BOARD_TYPE

For information on event source(s), see "Event Source."

923

Description: **Dual-active negotiation failed disk channels.**

Cause: Number of disk channels differ. Both controllers must have same number of disk channels.

Action: Replace, adjust or add components so that the controllers have the same number of disk channels.

Address: ctl: 0

Severity: ERROR

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_NEGOTIATION_DISK_CHANNELS

For information on event source(s), see "Event Source."

924

Description: **Dual-active negotiation failed host ports.**

Cause: Host ports differ. Both controllers must have same number of host ports.

Action: Replace, adjust or add components so that the controllers have the same number of host ports.

Address: ctl: 0

Severity: ERROR

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_NEGOTIATION_HOST_CHANNELS

For information on event source(s), see "Event Source."

925

Description: **Dual-active negotiation failed memory size.**

Cause: Memory sizes differ. Both controllers must have same control store size.

Action: Replace, adjust or add components so that the controllers have the same control store size.

Address: ctl: 0

Severity: ERROR

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_NEGOTIATION_MEMORY_SIZE

For information on event source(s), see "Event Source."

926

Description: Dual-active negotiation failed cache memory size.

Cause: Cache memory sizes differ. Both controllers must have same cache memory size.

Action: Replace, adjust or add components so that the controllers have the same cache memory size.

Address: ctl: 0

Severity: ERROR

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_CTLDEV_NEGOTIATION_CACHE_SIZE

For information on event source(s), see "Event Source."

927

Description: Hot spare replaced with a smaller capacity physical disk.

Cause: The new hot spare may have a smaller physical capacity than the physical disk it replaced. The controller's coercion setting may have reduced the configurable size of the new hot spare.

Action: None.

Address: ctl: 0 chn: 1 tgt: 2

Severity: INFORMATION

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_PHYSDEV_HOT_SPARE_SMALLER

For information on event source(s), see "Event Source."

980

Description: **Controller failed, SES not operational. Try a rescan....**

Cause: When Array Manager scanned the fibre channel loop / fabic, it found a PowerVault 660F array that wasn't fully initialized (may be still booting up). Therefore, Array Manager wasn't able to obtain configuration data from the array. This array will not be included in the Array Manager console for management.

Action: Wait a few minutes, and then do a Rescan from the PowerVault 660F folder object. If the array was booting up when the error occurred, it may be fully initialized during this rescan.

Address: None

Severity: ERROR

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SES_ERR

For information on event source(s), see "Event Source."

981

Description: **Unable to get enclosure status. Doing a Mylex (PV660F SubSystem) rescan may resolve this problem.**

Cause: When Array Manager scanned the fibre channel loop / fabic, it found a PowerVault 660F array that wasn't fully initialized (may be still booting up). Therefore, Array Manager wasn't able to obtain configuration data from the array. This array will not be included in the Array Manager console for management.

Action: Wait a few minutes, and then do a Rescan from the PowerVault 660F folder object. If the array was booting up when the error occurred, it may be fully initialized during this rescan.

Address: ctl: 0

Severity: SERIOUS

Source: Mylex (PV660F SubSystem)

Name: DAM_MLXEV_SES_ERR

For information on event source(s), see "Event Source."

Volume Management Events

The following sections describe events related to volume management.

To locate an event, scroll through these sections to find the event number displayed on the Events tab of the Array Manager console or use the Search function in this online help file to search for the event message text.

10804

Description: Free disk space has reached warning threshold on volume C:

Cause: The amount of used disk space has reached the warning threshold set in the Windows registry. Array Manager generates this event when capacity monitoring is turned on and the used disk space reaches a threshold.

Action: Either remove non-critical files to free up disk space or extend the volume to provide additional space.

Severity: Information

Source: fsys

Event and SNMP Trap Number: This event is listed as event 804 in the Array Manager event log. The SNMP trap number is 10804.

For more information on volume capacity monitoring, see "Volume Capacity Monitoring."

For information on event source(s), see "Event Source."

10805

Description: Free disk space has reached error threshold on volume C:

Cause: The amount of used disk space has reached the error threshold set in the Windows registry. Array Manager generates this event when capacity monitoring is turned on and the used disk space reaches a threshold.

Action: Either remove non-critical files to free up disk space or extend the volume to provide additional space.

Severity: Information

Source: fsys

Event and SNMP Trap Number: This event is listed as event 805 in the Array Manager event log. The SNMP trap number is 10805.

For more information on volume capacity monitoring, see "Volume Capacity Monitoring."

For information on event source(s), see "Event Source."

SNMP Support

By default, Array Manager installs SNMP trap forwarding support. For this support to function, you should have SNMP installed on the managed system prior to installing Array Manager. (For a full description of Array Manager installation requirements, see "Installation Requirements.")

The events displayed in the Array Manager event log are forwarded to the Windows application event log. If you have SNMP installed on the managed system (and the SNMP service is running), the Array Manager events in the Windows application event log will be forwarded as SNMP traps. In order for these traps to be viewable, however, a target system or application must be configured to receive these traps. SNMP traps that are generated by Array Manager can be viewed in any standard SNMP-compatible enterprise management console.

Microsoft Windows and NetWare SNMP services must be configured to forward the SNMP traps to the target system or application. When forwarding to an application, the application should also be configured to receive the SNMP traps. The IT Assistant application is already configured to receive the SNMP traps generated by Array Manager.

Refer to your Windows or NetWare operating system documentation for information on configuring the operating system to forward SNMP traps. This information may be located under such topics as "setting up SNMP" or "SNMP traps." When configuring SNMP for Windows, be sure that the SNMP traps are forwarded to the correct server. For information on configuring an application to receive SNMP traps, refer to the documentation for that application.

The Array Manager management information base (MIB) defines the SNMP traps that Array Manager generates. These traps correspond to the events documented in the "Event Messages" section. The MIB is located in the Array Manager installation directory. The default location for the MIB is Program Files\Dell\OpenManage\Array Manager\MIB. For more information on the MIB and its structure, as well as a change history of the SNMP traps, see the PDF document titled *Dell OpenManage Array Manager MIB Structure*. The file name for this document is amMIBstructure.pdf. The default location for this file is in the Program Files\Dell\OpenManage\Array Manager\MIB directory.

Note Array Manager supports trap forwarding on both 32-bit and 64-bit operating systems.

Array Manager Utilities 21

The Array Manager Utilities allows you to easily view and modify your Array Manager configuration on a local or remote Windows managed system or a remote NetWare® managed system.

You may find that some of the options in the Array Manager Utilities are disabled after you connect to a specific managed system. This usually occurs when specific requirements with regard to that option have not been met on that managed system. Some options are also not available with earlier versions of Array Manager.

Starting and Stopping the Array Manager Service

To open the Array Manager Utilities, click **Start | Programs | Dell OpenManage Applications | Array Manager** and select **Array Manager Utilities**.

The Array Manager service is what allows communication with storage hardware in your system. Stopping the service will therefore stop this communication and render any console monitoring inoperative. The Array Manager service has been designed to run at all times for constant storage integrity.

To stop the service for any reason, press the **Stop** button. Use the **Start** button to start the service.

Connecting to Remote Systems with Array Manager Utilities

You can connect to a remote Windows or NetWare system using the **Server**, **Username**, and **Password** text boxes in the **Authentication** section of the **Windows** or **NetWare** tab.

Note If you are already authenticated on the NetWare system (such as through a login or mapped drive) then you may observe that any value is accepted in the **Password** text box. This is because you are already authenticated and the password information is cached.

Command Line Options for NetWare

The Array Manager Utilities options for NetWare do not include the ability to stop the Array Manager service remotely. However, these options are available from the NetWare server console. To enable this feature, you must load **vxcmds.nlm**. If you checked the "Automatically load Array Manager when NetWare is started" check box during installation, this feature is already enabled. Array Manager's NetWare command-line options include:

vxstart	Starts the Array Manager service
vxstop (-f)	Stops the Array Manager service (-f forces the stop, even if clients are currently connected)
vxproviders	Displays a list of the currently loaded providers

Note On Windows, the Array Manager service is a system service. On NetWare, the Array Manager service is an NLM.

Note Dell does not offer NetWare in Japan.

Adding and Removing PowerEdge RAID Controller Configuration

You can add or remove components for PERC RAID controller configuration from the Array Manager Utilities.

To add or remove PERC RAID controller configuration components:

1. Launch the Array Manager Utilities by clicking **Start | Programs | Dell OpenManage Applications | Array Manager** and selecting **Array Manager Utilities**.

2. Select the **Windows** tab if you are adding or removing components from a Windows system. Select the **NetWare** tab if you are adding or removing components from a NetWare system.

   Both the **Windows** and **NetWare** tabs have a **PowerEdge RAID Controller Configuration** section with radio buttons for adding and removing these components.

3. Select or deselect the appropriate radio button for adding or removing the required components.

4. Click **Apply** and then **Close** to exit the Array Manager Utilities. On a Windows system you are now finished with adding or removing a component. On NetWare, you need to reboot the NetWare system in order for the added or removed component to take effect.

Adding and Removing Fibre Channel RAID and Enclosure Management

You can add or remove components for fibre channel RAID enclosure management from the Array Manager Utilities.

To add or remove fibre channel RAID and enclosure management components:

1. Launch the Array Manager Utilities by clicking **Start | Programs | Dell OpenManage Applications | Array Manager** and selecting **Array Manager Utilities**.

2. Select the **Windows** tab if you are adding or removing components from a Windows system. Select the **NetWare** tab if you are adding or removing components from a NetWare system.

   Both the **Windows** and **NetWare** tabs have a **Fibre Channel RAID and Enclosure Management** section with radio buttons for adding and removing these components.

3. Select or deselect the appropriate radio button for adding or removing the required components. On the **Windows** tab, you can add or remove fibre channel RAID and enclosure management components for both a PV660F and a direct-attached environment. Because NetWare does not support a direct-attached environment, the **NetWare** tab only enables you to add or remove PV660F components.

4. Click **Apply** and then **Close** to exit the Array Manager Utilities.

When using Array Manager to manage a PowerVault 660F storage system in a SAN configuration, you must set the managed system on only one server per fiber zone to do PowerVault 660F management. The managed system is capable of managing multiple PowerVault 660F arrays, but each PowerVault 660F array must not have multiple Array Manager systems managing it. Having multiple managed systems managing the same PowerVault 660F storage subsystem will cause Array Manager to malfunction.

If the managed system is configured to do PowerVault 660F storage system management, you will need to also select direct-attach mode if the system is directly connected (HBA to controller) to the PowerVault 660F storage system. Do not select this option if the PowerVault 660F storage system will be shared by multiple Windows or Linux servers (that is, in a SAN configuration). Doing so may result in data corruption.

Caution You can change the PV660F setting only when there are no virtual disks configured on the PowerVault 660F storage system.

Local and Remote Server Management

By default, Array Manager Utilities will attempt to manage the local system. Remote server management requires authentication first for security purposes. To connect to a remote system:

1. Launch the Array Manager Utilities by clicking **Start | Programs | Dell OpenManage Applications | Array Manager** and selecting **Array Manager Utilities**.

2. Select the **Windows** tab if you are connecting to a remote Windows system. Select the **NetWare** tab if you are connecting to a remote NetWare system.

   Both the **Windows** and **NetWare** tabs have a **Server Type** section with radio buttons for specifying a local or remote server. The default setting is **Local Server**.

3. Select the **Remote Server** option. This action activates the fields in the **Authentication** section of the tab.

4. Enter your server name, user name, and password.

5. Click **Connect!**

If the connection is successful, you should see information about the configuration of the server.

Note Novell's Client 32 version 4.6 or later is required to make a connection to a NetWare server.

Volume Capacity Monitoring

When volume capacity monitoring is enabled, Array Manager monitors the NTFS volumes on the local server for the amount of space used. When the space used reaches 90%, a warning event is logged in the Array Manager event log and the Windows application log. When the space used reaches 98%, an error event is logged.

When enabled, volume capacity monitoring can degrade system performance due to the increased CPU usage. For this reason, volume capacity monitoring is disabled by default. Your operating system may also have its own volume capacity monitoring function. If so, see your operating system documentation for more information.

Note Volume capacity monitoring is not supported on NetWare or Windows Server 2003. For more information see "Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux."

To enable volume capacity monitoring:

1. Launch the Array Manager Utilities by clicking **Start | Programs | Dell OpenManage Applications | Array Manager** and selecting **Array Manager Utilities**.
2. Select the **Volume Capacity Monitoring** check box on the **Windows** tab.
3. Click **Apply** and then **Close** to exit the Array Manager Utilities.

For information related to volume capacity monitoring, see the following:

- Volume Management Events
- System Performance Problems
- Volume Troubleshooting Procedures

SNMP Configuration on NetWare

To add or remove SNMP configuration components on NetWare:

1. Launch the Array Manager Utilities by clicking **Start | Programs | Dell OpenManage Applications | Array Manager** and selecting **Array Manager Utilities**.
2. Select the **NetWare** tab.
3. Select or deselect the **Simple Network Management Protocol (SNMP)** check box.
4. Click **Apply** and then **Close** to exit the Array Manager Utilities.

Message Notification

You can specify whether or not the Array Manager console displays pop-up messages. Pop-up messages are displayed in response to system and user errors. By default, pop-up messages are displayed. You can disable or enable the display of pop-up messages by using the **Display Messages** option.

To disable or enable the pop-up messages:

1. Launch the Array Manager Utilities by clicking **Start | Programs | Dell OpenManage Applications | Array Manager** and selecting **Array Manager Utilities**.
2. Select or deselect the **Display Messages** check box on the **Windows** tab. When selected, the Array Manager console displays pop-up messages.
3. Click **Apply** and then **Close** to exit the Array Manager Utilities.

Array Manager also displays event messages in the event log and generates SNMP traps. for more information on event log messages and SNMP traps, see the following:

- Event Monitoring
- Viewing SNMP Traps

Version Checking

The Array Manager Utilities automatically checks the versions of required or recommended hardware and software on your local system. This includes:

- Internet Explorer
- FAST Application Status
- Drivers and Firmware

Internet Explorer

Internet Explorer is required for viewing the online help. For more information on the version requirements, see “Online Help Requirements.”

FAST Application Status

Some controllers have their own management application that may not be compatible with Array Manager. Such is the case with the Flexible Array Storage Tool (FAST). Array Manager cannot be installed on a system that has FAST installed. The Array Manager Utilities check whether FAST is installed on the system.

Dell OpenManage Server Agent

On NetWare, Array Manager requires that the Dell OpenManage Server Agent (OMSA) 4.0.0 or later be installed on the managed system to provide thermal shutdown (when an enclosure becomes too hot or cold) for the Dell PowerVault 20xS, 21xS, and 22xS SCSI enclosures. If OMSA is not installed, Array Manager can provide other features of enclosure management, but not thermal shutdown without the risk of data loss.

Drivers and Firmware

The Versions tab in the Array Manager Utilities displays the drivers and firmware currently installed on your system. The Versions tab also indicates whether the drivers and firmware on your system meet the requirements of the version of Array Manager installed on your system.

During Array Manager installation, you may need to reboot the system in order for Array Manager to detect the currently installed drivers and firmware. The Array Manager installation process prompts you to reboot your system at the appropriate time.

See the Readme file for driver and firmware requirements, as well as other system requirements. For any updates that occur between software releases, refer to the Dell support site at http://support.dell.com.

MMC Snap-in 22

Under Windows 2000 and Windows Server 2003, you can use Array Manager from the Microsoft Management Console (MMC) instead of the Array Manager console. For more information on the MMC, refer to the MMC online help.

You can launch the MMC snap-in from the Windows Start menu.

1. Click the **Start** menu.

2. Select **Start | Programs | Dell OpenManage Applications | Array Manager | MMC Snap-in | Array Manager Snap-in for MMC**.

Array Manager Command Line Interface 23

Array Manager provides the following commands for initiating management functions from the command line:

- amcli Command
- amconsole Command

amcli Command

Using the Array Manager amcli command, you can initiate a battery recondition and consistency check from a command prompt. You can also check the status of controller batteries and consistency check operations.

- Scheduling amcli Commands
- amcli Command Options

Scheduling amcli Commands

You can use the Windows Task Scheduler or Add Scheduled Tasks wizard to automatically initiate the amcli commands. For example, you may want to schedule regular consistency checks or battery reconditions. When using the Task Scheduler or Add Scheduled Tasks wizard, you need to specify the path to the amcli executable. The amcli executable is located in the Array Manager installation directory, and the file name is amcli.exe.

When specifying the path name to the amcli executable, you may need to use quotation marks around the path name. The parameters for the amcli command should follow the path name and not be included in the quotation marks. The following is example syntax for scheduling a check consistency on virtual disk 1.

"C:*PathName*\amcli.exe" /c1

where *PathName* is the path to the amcli executable.

amcli Command Options

To view the amcli command options, display the amcli help information by entering the following at a command prompt:

amcli /?

The amcli command options are as follows:

amcli Command Options	Comments
amcli /?	Displays the amcli help.
amcli /da where the **d** option indicates "display" and the **a** option indicates "adapter."	Displays the status of the battery of the RAID system controllers (adapters). The possible status values are: ◆ **None** — The controller does not have a battery. ◆ **Ready** — The battery is functioning normally and can be recharged. ◆ **Charging** — The battery is recharging. ◆ **Auto** — The battery is a self-charging battery. A battery recharge does not need to be initiated for this type of battery. ◆ **Degraded** — The battery charge is too low. Initiate a recharge immediately. ◆ **Low** — The battery charge is low. Initiate a recharge soon. ◆ **Failed** — The battery needs to be replaced. **Note** Non-RAID controllers, such as SCSI Subsystem controllers (adapters), do not have batteries. For this reason, the **amcli /da** command only displays the battery status of RAID controllers.
amcli /dp where the **d** option indicates "display" and the **p** option indicates "provider" or subsystem.	Displays the subsystems that can be enabled or disabled for Array Manager. The subsystems are displayed in a numbered list. Use this command to identify the number of a subsystem you want to enable or disable with the **amcli /in** and **amcli /un** commands.

amcli Command Options	Comments
amcli /dv where the **d** option indicates "display" and the **v** option indicates "virtual disk."	Displays the status of the virtual disks. The possible status values are: ◆ **Ready** — The virtual disk is operational. ◆ **Not Ready** — The virtual disk is not operational. For example, the virtual disk may have a SMART alert or one if its member array disks may have been removed. ◆ **Resynching** — A consistency check is being performed on the virtual disk. ◆ **Failed Redundancy** — An operation to create a redundant RAID level on the virtual disk has failed. This status only applies to virtual disks on the following controllers: PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s.
amcli /c*n* where the **c** option indicates "consistency check" and *n* is the number of a virtual disk as displayed in the Array Manager tree view.	Initiates a check consistency of the specified virtual disk. Unlike the Array Manager GUI (tree view), the amcli command numbers virtual disks starting at 1 and not 0. This means that virtual disk 0 in the tree view would be specified as **amcli /c1**. Virtual disk 1 in the tree view would be specified as **amcli /c2**, and so on.
amcli /a*n* where the **a** option indicates "abort" and *n* is the number of a virtual disk as displayed in the Array Manager tree view.	Stops a check consistency operation that is already in progress for the specified virtual disk. Unlike the Array Manager GUI (tree view), the amcli command numbers virtual disks starting at 1 and not 0. This means that virtual disk 0 in the tree view would be specified as **amcli /a1**. Virtual disk 1 in the tree view would be specified as **amcli /a2**, and so on.

amcli Command Options	Comments
amcli /r*n* where the **r** option indicates "recondition" and *n* is the number of the controller as displayed in the Array Manager tree view.	Initiates a battery recondition for the specified controller.
amcli /i*n* where *n* is the number of the subsystem you wish to enable.	Enables a subsystem. When enabled, the subsystem and all of its controllers are displayed in the Array Manager tree view. A subsystem must be enabled for it to be managed with Array Manager. Use the **amcli /dp** command to identify the number of the subsystem you wish to enable.
amcli /u*n* where *n* is the number of the subsystem you wish to disable.	Disables a subsystem. When disabled, the subsystem and its controllers are not displayed in the Array Manager tree view. A disabled subsystem cannot be managed with Array Manager. Use the **amcli /dp** command to identify the number of the subsystem you wish to disable.
amcli /l <*FileName*> where the **l** option indicates "log file" and <*FileName*> is the path and file name you specify for the log file.	Saves configuration information to the log file that you specify. If you do not specify a path or a file name, then the log file resides in the current directory and is named Amconfig.txt. When specifying a path, be sure to include the file name.

amcli Command Options	Comments
amcli /e*n* where *n* is the number of a PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, or CERC SATA1.5/2s controller as displayed in the Array Manager tree view.	Saves the log file information from a PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, or CERC SATA1.5/2s controller to the log_date.log file in the Array Manager install directory.
amcli /v <*FileName*> where <*FileName*> is the path and file name you specify for the log file.	Saves the Array Manager events from the system application log to a flat text log file that you specify. If you do not specify a path or a file name, then the log file resides in the current directory and is named amevent.log.

amconsole Command

The **amconsole** command launches Array Manager on a Windows system. This section provides the **amconsole** command's syntax along with its parameters. This section is provided for reference purposes. Other sections provide specific uses of the **amconsole** command and appropriate syntax examples. You may wish to refer to one of these sections:

- Starting Array Manager from the Command Line
- Starting Array Manager with a Controller Highlighted
- Starting Array Manager with a Remote Connection

To launch Array Manager on the local system, you can simply enter:

amconsole

without including any of the **amconsole** command parameters.

The **amconsole** command syntax and parameters are as follows:

amconsole $computer:*ComputerName* user:*UserName* password:*Password* domain:*DomainName* device:"*ControllerName*"

where:

- *ComputerName* is the name of the remote computer. When you specify a remote computer, the controllers attached to the remote computer are displayed in the Array Manager tree view. This parameter is only necessary when connecting to a remote computer.
- *ControllerName* is the controller that you wish to be highlighted in the Array Manager tree view. This parameter is only necessary if you wish to highlight a particular controller.
- *DomainName* is the domain to which the remote computer belongs. This parameter is only necessary when connecting to a remote computer.
- *Password* is the password of the administrator. Array Manager requires administrator authority. When connecting to a remote system with Array Manager, it is recommended that the password be the same on both the local and remote system.
- *UserName* is the user name of an administrator. Array Manager requires administrator authority. When connecting to a remote system with Array Manager, it is recommended that the user name be the same on both the local and remote system.

The **amconsole** command and its parameters should be entered on a single line without a line break. Elements of the controller name should be separated by a space. When using spaces, enclose the item containing spaces in quotation marks. For example, to start Array Manager with a PERC 2 /Si controller selected in the tree view, enter:

amconsole $device:“perc 2/si controller 0” user:*UserName* password:*Password*

Troubleshooting 24

This chapter contains status message information, troubleshooting procedures, and common problems and solutions. It also has a separate section for troubleshooting the Dell PowerVault™ 660F and 224F storage systems.

Disks and Volumes

If a disk or volume fails, it is important to repair the disk or volume as quickly as possible to avoid data loss. Because time is critical, Array Manager makes it easy for you to locate problems quickly. In the Status column of the list view, you can view the status of a disk or volume. The status also appears in the graphical view of each disk or volume. If the status is not **Healthy** for volumes or **Online** for disks, use the status information to determine the problem and then fix it.

There are also various troubleshooting procedures for disks, volumes, and arrays.

Topics include:

- Operating System Disk Status Descriptions
- Array and Hard Disk Status Information
- Disk Troubleshooting Procedures
- Volume Status Descriptions
- Volume Troubleshooting Procedures

Operating System Disk Status Descriptions

On a Windows system, you can expand the Disks object in the Array Manager tree view to display the operating system disks recognized by Windows Disk Management. When an operating system disk is selected, the following disk status descriptions display in the Status column in the right pane of the console window.

If there is a problem with an operating system disk, you can use this troubleshooting chart to diagnose and correct the problem.

.

Status	Meaning
Online	The disk is accessible and has no known problems. This is the normal disk status. No user action is required. Both dynamic disks and basic disks display the **Online** status.
Online (Errors)	This status indicates that the disk is in an error state or that I/O errors have been detected on a region of the disk. All the volumes on the disk will display **Failed** or **Failed Redundancy** status, and you may not be able to create new volumes on the disk. Only dynamic disks display this status. Right-click the failed disk and select **Reactivate Disk** to bring the disk to an **Online** status and bring all the volumes to a **Healthy** status.
Offline	The disk is not accessible. The disk may be corrupted or intermittently unavailable. An error icon appears on the offline disk. Only dynamic disks display the **Offline** status. If the disk status is **Offline** and a separate corresponding icon titled **Missing Disk** appears, the disk was recently available on the system but can no longer be located or identified. The **Missing** disk may be corrupted, powered down, or disconnected, or the disk may be a virtual disk that has been deleted.
Unreadable	The disk is not accessible. The disk may have experienced hardware failure, corruption, or I/O errors. The disk's copy of the system's disk configuration database may be corrupted. An error icon appears on the **Unreadable** disk. Both dynamic and basic disks display the **Unreadable** status. Disks may display the **Unreadable** status while they are spinning up or when Array Manager is rescanning all the disks on the system. In some cases, an **Unreadable** disk has failed and is not recoverable. For dynamic disks, the **Unreadable** status usually results from corruption or I/O errors on part of the disk, rather than failure of the entire disk. You can rescan the disks (using the **Rescan Disks** command) or reboot the computer to see if the disk status changes.
Unrecognized	The disk has an original equipment manufacturer's (OEM) signature and Array Manager will not allow you to use this disk. For example, a disk from a UNIX system displays the **Unrecognized** status. Only **Unknown** disk types display the **Unrecognized** status.

Status	Meaning
Foreign Disk	The disk has been moved to your computer from another Windows® 2000 computer and has not been set up for use. Only dynamic disks display this status. To add the disk so that it can be used, right-click the disk and select **Merge Foreign Disk**. All existing volumes on the disk will be visible and accessible. Because a volume can span more than one disk (e.g., a mirrored volume), it is important that you first verify your disk configurations and then move the entire disk set that the volume is on. If only part of the disk set is moved, some of the volumes will show the **Failed Redundancy** or **Failed** error condition.

Array and Hard Disk Status Information

These definitions appear in the Status line and indicate the condition of array and hard disks. Array disks are attached to a RAID controller and are capable of being used in a virtual disk.

Status line entry	Status indication
Unknown	May signify a problem or indicate a transitional state. Additionally, a new disk that had previously been formatted or initialized by another type of RAID controller may show this state.
Ready	Means the array disk is operational. For PERC 2/SC, 3/SC, 2/DC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers, Ready status applies to operational array disks that are not part of a virtual disk. For the PERC 2, PERC 2/Si, PERC 3/Si, PERC 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controllers, operational array disks display Ready status regardless of whether they are a part of a virtual disk or not.

Failed	Not operational. A disk needs repair, has been removed, or has another problem that prevents operation.
Online	The array disk is part of a virtual disk and is functioning normally.
Offline	The drive is not available to the RAID controller.
Degraded	Refers to a fault-tolerant array / virtual disk that has a failed disk. This state definition may also appear when resynching the array / virtual disk, since the array / virtual disk is not fault-tolerant during the resynchronization.
Recovering	Refers to state of recovering from bad blocks on disks.
Removed	Indicates that an array disk has been removed. This state may occur when the disk was part of a virtual disk but has lost communication with the controller. A user may have removed the disk or it may have lost power. A disk that is in Removed state will no longer be displayed by Array Manager after a system reboot or stopping and restarting the Array Manager service.
Resynching	This state definition appears during the following types of disk operations: Transform Type, Reconfiguration, and Check Consistency.
Rebuilding	Refers to part of a virtual disk being rebuilt.
No Media	CD-ROM or removable disk has no media.
Diagnostics	Indicates that diagnostics are running.
Reconstructing	The configuration of a virtual disk has been changed. The individual array disks within the virtual disk are being modified to support the changes. The data on the virtual disk will be saved. You cannot cancel a virtual disk reconstruction.

Initializing	Applies only to virtual disks on PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers. This prepares the virtual disk for use by Array Manager by deleting the configuration information on this virtual disk. The data on the virtual disk will be lost.

Disk Troubleshooting Procedures

The following sections describe disk troubleshooting procedures:

- Rescan to Update Information
- Maintain integrity of redundant (mirrored and parity) information
- Reactivate a Disk
- Bring a Dynamic RAID-5 or Mirrored Volume Back Online
- Replacing a Failed Disk
- Recovering from Removing the Wrong Drive
- Rebuild does not work
- A Rebuild Completes with Errors on PERC Subsystem 1 Controllers
- Rebuilding a Virtual Disk on the CERC SATA1.5/2s Controller
- Cannot create a virtual disk (option is inactive)
- When expanding the Disks object, error icons appear
- Missing Disk displays error icon
- To bring a disk that is Offline and Missing back online
- To bring a disk that is Offline (not Missing) and is still named Disk # back online
- A disk on a PERC 4/Di, 4e/Si, or 4e/Di controller does not return online after a Prepare to Remove
- A disk is marked as Foreign
- A Disk is Marked as Offline or Foreign after Upgrading to Dynamic with a PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, or CERC ATA100/4ch Controller
- A Disk's Functions become Inactive or It is Marked as Foreign or Offline after Upgrading to Dynamic with a PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s Controller

- During reboot, a message displays about a "corrupt drive," suggesting that you run autocheck
- After creating a virtual disk with a PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, or CERC ATA100/4ch controller, the virtual disk does not appear under the Disks storage object

See also the following sections for these and other troubleshooting procedures:

- Common Troubleshooting Procedures
- Problem Situations and Solutions
- Dell PowerVault 660F and 224F Storage Systems Troubleshooting

Volume Status Descriptions

One of the following volume status descriptions will always appear in the graphical view of the volume and in the Status column of the volume in list view. If there is a problem with a volume, you can use this troubleshooter to diagnose and correct the problem.

Status	Meaning
Healthy	The volume is accessible and has no known problems. This is the normal volume status. No user action is required. Both dynamic volumes and basic volumes display the **Healthy** status.
Healthy (At Risk)	The volume is currently accessible, but I/O errors have been detected on the underlying disk. If an I/O error is detected on any part of a disk, all volumes on the disk display the **Healthy (At Risk)** status. A warning icon appears on the volume. Only dynamic volumes display the **Healthy (At Risk)** status. When the volume status is **Healthy (At Risk)**, an underlying disk's status is usually **Online (Errors)**. To return the underlying disk to the **Online** status, reactivate the disk (using the **Reactivate Disk** command). Once the disk is returned to **Online** status, the volume should return to the **Healthy** status.
Initializing	The volume is being initialized. Dynamic volumes display the **Initializing** status. No user action is required. When initialization is complete, the volume's status becomes **Healthy**. Initialization should be completed very quickly.

Status	Meaning
Resynching	The volume's mirrors are being resynchronized so that both mirrors contain identical data. Both dynamic and basic mirrored volumes display the **Resynching** status. No user action is required. When resynchronization is complete, the mirrored volume's status returns to **Healthy**. Resynchronization may take some time, depending on the size of the mirrored volume. Although you can access a mirrored volume while resynchronization is in progress, you should avoid making configuration changes (such as breaking a mirror) during resynchronization.
Regenerating	Data and parity are being regenerated for the RAID-5 volume. Both dynamic and basic RAID-5 volumes display the Regenerating status. No user action is required. When regeneration is complete, the RAID-5 volume's status returns to **Healthy**. You can access a RAID-5 volume while data and parity regeneration is in progress.
Failed Redundancy	The data on the volume is no longer fault tolerant because one of the underlying disks is not online. A warning icon appears on the volume with **Failed Redundancy**. The **Failed Redundancy** status applies only to mirrored or RAID-5 volumes. Both dynamic and basic volumes display the Failed Redundancy status. You can continue to access the volume using the remaining online disks, but if another disk that contains the volume fails, you will lose the volume and its data. To avoid such loss, you should attempt to repair the volume as soon as possible. A **Failed Redundancy** status will also display if a disk was moved and the volume on it spanned more than the single disk. To correct the problem, you must move the entire disk set that contains all the appropriate volumes.

Status	Meaning
Failed Redundancy (At Risk)	The data on the volume is no longer fault tolerant, and I/O errors have been detected on the underlying disk. If an I/O error is detected on any part of a disk, all volumes on the disk display the **(At Risk)** status. A warning icon appears on the volume. Only dynamic mirrored or RAID-5 volumes display the **Failed Redundancy (At Risk)** status. When the volume status is **Failed Redundancy (At Risk)**, the underlying disk's status is usually **Online (Errors)**. To return the underlying disk to the Online status, reactivate the disk (using the **Reactivate Disk** command). Once the disk is returned to the **Online** status, the volume status should change to **Failed Redundancy**.
Failed	The volume cannot be started automatically. An error icon appears on the failed volume. Both dynamic and basic volumes display the **Failed** status.
Formatting	The volume is being formatted using the specifications you chose for formatting.
No Media	No media has been inserted into the CD-ROM or removable drive. The volume status will become **Online** when you insert the appropriate media into the CD-ROM or removable drive. Only CD-ROM or removable disk types display the **No Media** status.

Volume Troubleshooting Procedures

The following sections describe common volume troubleshooting procedures:

- Bring a Dynamic RAID-5 or Mirrored Volume Back Online
- Reactivate a Dynamic Volume
- Repair Basic Volumes
- Repair Dynamic Volumes
- Repair a Dynamic RAID-5 Volume
- Cannot create a RAID-5 volume
- Drive Letters and Drive Mapping

See also the following sections for these and other troubleshooting procedures:

- Common Troubleshooting Procedures
- Problem Situations and Solutions
- Dell PowerVault 660F and 224F Storage Systems Troubleshooting

Common Troubleshooting Procedures

This section describes commands and procedures that can be used in troubleshooting. Topics covered include:

- Cables attached correctly
- System Requirements
- Drivers and Firmware
- Isolate SCSI device problems
- Rescan to Update Information
- Maintain integrity of redundant (mirrored and parity) information
- Reactivate a Disk
- Replacing a Failed Disk
- Repair Basic Volumes
- Reactivate a Dynamic Volume
- Repair Dynamic Volumes
- Repair a Dynamic RAID-5 Volume
- Repair Basic Volumes
- Drive Letters and Drive Mapping
- Recovering from Removing the Wrong Drive
- Resolving Windows Upgrade Problems

Cables attached correctly

Verify that the power-supply cord and adapter cables are attached correctly. If the system is having trouble with read and write operations to a particular array (if the system hangs, for example), then make sure that the SCSI cables attached to the array are secure. If the connection is secure but the problem persists, you may need to replace a cable. See also the “Isolate SCSI device problems” section.

System Requirements

Make sure that the system meets all system requirements as described in the readme.txt file located in the installation directory. In particular, verify that the correct levels of firmware and drivers are installed on the system. For more information on drivers and firmware, see the “Drivers and Firmware” section.

Drivers and Firmware

Array Manager is tested with the supported controller firmware and drivers. The supported controllers and firmware are listed in the readme.txt file. To avoid possible conflicts or inconsistencies between the controller firmware and drivers, it is recommended that you only use the supported versions. The most current versions can be obtained from the Dell support site at http://support.dell.com.

In a SAN environment, all LS modules in an array should have the same firmware version. When upgrading the firmware on an LS module, make sure to upgrade the firmware on the other LS modules at the same time.

It is also recommended to obtain and apply the latest Dell PowerEdge™ Server System BIOS on a periodic basis to benefit from the most recent improvements. Please refer to the Dell PowerEdge system documentation for more information.

Isolate SCSI device problems

If you receive a “timeout” event related to a SCSI device or if you otherwise suspect that one of the SCSI devices is experiencing a hardware failure, then do the following to confirm the problem:

- Verify that the cables are correctly attached.
- If the cables are correctly attached and you are still experiencing the problem, then disconnect the device cables and reboot the system. If the system reboots successfully, then one of the SCSI devices may be defective. Refer to the SCSI device documentation for more information.

Rescan to Update Information

Use **Rescan** to update disk information. This operation may take a few minutes if there are a number of devices attached to the system. You will see a message “Getting hardware configuration. Please wait.” while the rescan is occurring.

If this does not properly update the disk information, you may need to reboot your system.

Maintain integrity of redundant (mirrored and parity) information

The Check Consistency function determines the accuracy of mirrored data and parity information. When necessary, this feature rebuilds the parity information. For more information, see the following sections:

- "Check Consistency" for PERC Subsystem 1 and PERC Subsystem 2 controllers
- "Check Consistency (PV660F)" for the PowerVault 660F storage system

Reactivate a Disk

1. Reboot your machine to update the list of existing disks.
2. Right-click the disk marked Missing or Offline dynamic disk.
3. Use **Rescan** to change the disk status to Online (errors).
4. Right-click the disk marked Missing or Offline dynamic disk. Select Reactivate Disk from the context menu. The disk should be marked Online after the disk is reactivated.
5. For any volumes that are not Healthy, right-click the volume from the context menu and select **Reactivate Volume**.

Caution When reactivating a volume, be aware that the volume's data is restored, even if it is stale, corrupt, or out-of-date. See "Reactivate a Dynamic Volume" for more information on the consequences of reactivating a volume.

Replacing a Failed Disk

You may need to replace a failed disk in the following situations:

- Replacing a Failed Disk that is Part of a Redundant Virtual Disk
- Replacing a Failed Disk that is Part of a Non-redundant Virtual Disk
- Replacing a Failed Drive in a RAID 1 on a CERC SATA1.5/2s

Replacing a Failed Disk that is Part of a Redundant Virtual Disk

If the failed disk is part of a redundant virtual disk, then the disk failure should not result in data loss. You should replace the failed disk immediately, however, as additional disk failures can cause data loss.

If the redundant virtual disk has a hot spare assigned to it, then the data from the failed disk is rebuilt onto the hot spare. After the rebuild, the former hot spare functions as a regular array disk and the virtual disk is left without a hot spare. In this case, you should replace the failed disk and make the replacement disk a hot spare.

Replacing the Disk:

1. Remove the failed disk.
2. Insert a new disk. Make sure that the new disk is the same size or larger as the disk you are replacing. (On some controllers, you may not be able to use the additional disk space if you insert a larger disk. See "Virtual Disk Considerations for PERC Subsystem 1 Controllers" for more information.)

   A rebuild is automatically initiated because the virtual disk is redundant.

Assigning a Hot Spare:

If a hot spare was already assigned to the virtual disk, then data from the failed disk may already be rebuilt onto the hot spare. In this case, you need to assign a new hot spare. See "Configure Dedicated Hot Spare" and "Assign Global Hot Spare" for more information.

Replacing a Failed Disk that is Part of a Non-redundant Virtual Disk

If the failed disk is part of a non-redundant virtual disk (such as RAID 0), then the failure of a single disk will cause the entire virtual disk to fail. To proceed, you need to verify when your last backup was, and if there is any new data that has been written to the disk since that time.

If you have backed up recently and there is no new data on the drives that would be missed, you can restore from backup.

Do the following:

1. Delete the virtual disk which is currently in a failed state.
2. Remove the failed disk.
3. Insert a new disk.
4. Create a new virtual disk.
5. Create one (or more) new volumes.
6. Restore from backup.

Using the Online Command

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers

If you do not have a suitable backup available, and if the failed disk is part of a virtual disk on a PERC Subsystem 1 controller, then you can attempt to retrieve data by right-clicking the failed disk and selecting **Online** from the pop-up menu.

The Online command attempts to force the failed disk back into a Ready state. If you are able to force the disk into a Ready state, you may be able to recover individual files. How much data you can recover depends on the extent of disk damage. File recovery is only possible if a limited portion of the disk is damaged.

There is no guarantee you will be able to recover any data using this method. A forced Online does not fix a failed disk. You should not attempt to write new data to the virtual disk or volume.

After retrieving any viable data from the disk, replace the failed disk as described previously in this section.

Replacing a Failed Drive in a RAID 1 on a CERC SATA1.5/2s

On a CERC SATA1.5/2s controller, a rebuild may not start automatically when you replace a failed array disk that is part of a RAID 1 virtual disk. In this circumstance, use the following procedure to replace the failed array disk and rebuild the redundant data.

1. Disconnect the SATA cable on the failed array disk in the RAID 1 virtual disk.
2. Shut down the system.
3. Replace the failed array disk with a formatted array disk. You can format the array disk using the Disk Utilities in the controller BIOS.
4. Reboot the system. When rebooted, the RAID 1 virtual disk should display a Failed Redundancy state.
5. Expand the controller object in the tree view and select the **Array Disks** object.
6. Execute the **Rebuild** task for the array disk you added.

Bring a Dynamic RAID-5 or Mirrored Volume Back Online

A RAID 5 volume's status can appear as Failed Redundancy and the disk's status is Offline. The disk's name may be Missing, and an error icon **(X)** appears on the missing or offline disk. In this case, do the following.

1. Rescan the disk to make sure the disk, controller, or cable problem is fixed.
2. Try to reactivate the disk by right-clicking on the disk and selecting **Reactivate Disk**.
3. If the volume remains as Failed Redundancy or Failed, right-click the volume, then select **Reactivate Volume**. If all disks on this volume are Online, the volume should be brought back to a healthy state.

Caution When reactivating a volume, be aware that the volume's data is restored, even if it is stale, corrupt, or out-of-date. See "Reactivate a Dynamic Volume" for more information on the consequences of reactivating a volume.

Reactivate a Dynamic Volume

Reactivating a volume attempts to restart all volumes regardless of the volume's state. If data corruption exists, you can reactivate the volume and then run the **chkdsk** utility. However, in the case of a mirrored or RAID-5 volume, reactivating a volume with stale data can cause that data to be used when it is inaccurate.

Reactivating a volume should be done only if you understand that the volume's data, which might be corrupted, will be restored. For example, if one mirror in a mirrored volume fails and data is written to the remaining mirror, the data is now out of sync. Then, if the remaining mirror (the one with accurate data) fails and the first mirror is reactivated, the stale data becomes "real" data.

For this reason, it is important to act on data failures as soon as possible. You should use care when reactivating volumes.

Repair Basic Volumes

Make sure that the underlying physical disk is turned on, plugged in, and attached to the computer. No other user action is possible for basic volumes.

Repair Dynamic Volumes

1. If the disks are not online, use the **Rescan** and then the **Reactivate Disk** commands to return the disk to the Online status. If this succeeds, the volume automatically restarts and returns to the Healthy status. A mirrored volume repairs itself by resynchronizing the data in its mirrors. A RAID-5 volume repairs itself by regenerating its parity and data.

2. If the disk returns to the Online status but the volume does not return to the Healthy status, you can reactivate the volume manually (using the **Reactivate Volume** command).

Caution When reactivating a volume, be aware that the volume's data is restored, even if it is stale, corrupt, or out-of-date. See "Reactivate a Dynamic Volume" for more information on the consequences of reactivating a volume.

3. If the volume is a mirrored or RAID-5 volume with stale data, bringing the underlying disk online will not automatically restart the volume. If the disks that contain non-stale data are disconnected, you should bring those disks online first (to allow the data to become synchronized). Otherwise, restart the mirrored or RAID-5 volume manually (using the **Reactivate Disk** command), and then run Chkdsk.exe. To run Chkdsk.exe, click **Start**, click **Run**, type chkdsk, and then click **OK**.

4. If the disk does not return to the Online status and the volume does not return to the Healthy status, there may be something wrong with the disk. You should replace the failed mirror or RAID-5 disk region. To replace the failed mirror in a mirrored volume, use the **Remove Mirror** command to remove the failed mirror, then use the **Add Mirror** command to create a new mirror on another disk. To replace the failed disk region in a RAID-5 volume, use the **Repair RAID-5 Volume** command.

There are particular considerations regarding dynamic disks and volumes on NetWare, Windows Server 2003, and Linux. See "Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux" for more information.

Repair a Dynamic RAID-5 Volume

1. Right-click volume, then click **Repair RAID-5 volume**.

2. A message appears that indicates that the repair will be attempted if there is another dynamic disk with adequate unallocated space. Click **Yes** to confirm the repair.

3. The volume should be brought back to a healthy state.

You should be able to repair a RAID-5 volume if it is in a state of Failed Redundancy, and if there is unallocated space on another dynamic disk available. To avoid data loss, you should attempt to repair the volume as soon as possible.

Drive Letters and Drive Mapping

A Drive Letter is Unavailable

After deleting a basic disk, the drive letter used by that disk may no longer be available. To correct this problem, reboot the server.

Drive Mapping is Not Working

Drive mapping may not work properly on Windows 2000 systems with PERC 3/DC, PERC 3/DCL, PERC 3/QC, PERC 2/DC, PERC 3/SC, PERC 2/SC, PERC 4/SC, PERC 4/DC, 4e/DC, PERC 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers. After creating a virtual disk on these controllers, the disk may not be visible in the disk folder until the system is rebooted. After rebooting the system, the mapping between the newly created disk and the corresponding Windows 2000 disk may not be displayed in the Array Manager console.

Solution for Windows 2000:

When using a PERC 2/SC or 2/DC controller, upgrade your driver to MRAID 35X.SYS version 2.68 or later.

Recovering from Removing the Wrong Drive

If the drive that you mistakenly removed is part of a redundant virtual disk that also has a hot spare, then the virtual disk rebuilds automatically either immediately or when a write request is made. After the rebuild has completed, the virtual disk will no longer have a hot spare since data has been rebuilt onto the disk previously assigned as a hot spare. In this case, you should assign a new hot spare.

If the drive that you removed is part of a redundant virtual disk that does not have a hot spare, then replace the drive and do a rebuild.

See the following sections for information on rebuilding drives and assigning hot spares:

- Understanding Hot Spares for PERC Subsystem 1 and PERC Subsystem 2 controllers
- "Rebuild" for PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers
- "Rebuild" and "Make Global Hot Spare" for the PowerVault 660F storage system

You can avoid removing the wrong drive by blinking the LED display on the drive that you intend to remove. See the following sections for information on blinking the LED display:

- "Blink" for PERC Subsystem 1 and PERC Subsystem 2 controllers
- "Blink Disk (PV660F)" for the PowerVault 660F storage system

Resolving Windows Upgrade Problems

If you upgrade the Windows operating system on a server, you may find that Array Manager no longer functions after the upgrade. The installation process installs files and makes registry entries on the server that are specific to the operating system. For this reason, changing the operating system can disable Array Manager.

To avoid this problem, you should uninstall Array Manager before upgrading. If you have already upgraded without uninstalling Array Manager, however, you should uninstall Array Manager after the upgrade. See "Uninstalling Array Manager" for more information.

After you have uninstalled Array Manager and completed the upgrade, reinstall Array Manager using the Array Manager install media. You can download Array Manager from the Dell support site at http://support.dell.com.

Problem Situations and Solutions

This section contains additional trouble-shooting problem areas. Topics include:

- Rebuild does not work
- A Rebuild Completes with Errors on PERC Subsystem 1 Controllers
- Receive Event 691 with "BBR Error" or "Error with Sense," or "Medium Error" Message
- Cannot create a virtual disk (option is inactive)
- Cannot create a RAID-5 volume
- Cannot create a mirror
- When expanding the Disks object, error icons appear
- Missing Disk displays error icon
- Read and Write Operations Experience Problems
- Problems after Installing a PERC 2/SC or 2/DC Controller
- I/O Stops on a Channel Redundant Channel

- Error message that the connection to remote computer has terminated
- Error Message: The stripe depth is out of range
- Tree view object for PowerEdge RAID controller cannot be expanded after the software and driver are installed
- An option is inactive
- To bring a disk that is Offline and Missing back online
- To bring a disk that is Offline (not Missing) and is still named Disk # back online
- A disk on a PERC 4/Di, 4e/Si, or 4e/Di controller does not return online after a Prepare to Remove
- A disk is marked as Foreign
- A Disk is Marked as Offline or Foreign after Upgrading to Dynamic with a PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, or CERC ATA100/4ch Controller
- Virtual Disk Initialization Causes a Foreign, Offline, or Missing Disk
- Initialize Command on a Virtual Disk does not Appear to Erase the Volumes on the Disk
- A Disk's Functions become Inactive or It is Marked as Foreign or Offline after Upgrading to Dynamic with a PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s Controller
- The Online Help behaves strangely, or will not come up at all
- When attempting to bring up the Help file, Dr. Watson reports an Access Violation in HH.EXE
- During reboot, a message displays about a "corrupt drive," suggesting that you run autocheck
- When attempting to access a remote computer, you are denied access or get an error message
- Problems Occur after NetWare Upgrades or Configuration Changes
- After creating a virtual disk with a PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, or CERC ATA100/4ch controller, the virtual disk does not appear under the Disks storage object
- "The Connection Failed" message displays when connecting to the NetWare server
- A Disk is Marked as Failed when Rebuilding in a Cluster Configuration
- Erroneous Status and Error Messages after a Windows Hibernation
- Cannot Connect to Remote System from Windows Server 2003

- Event Description Missing from Event Details Message Box
- "Unsafe Removal of Device" message displays after uninstalling Array Manager on Windows 2000

Note If you are using the Dell PowerVault 660F storage system and the PowerVault 224F enclosure, see "Dell PowerVault 660F and 224F Storage Systems Troubleshooting," for additional issues specific to the PowerVault 660F storage system and PowerVault 224F enclosure.

Rebuild does not work

A rebuild will not work in the following situations:

- The virtual disk is non-redundant. For example, a RAID 0 virtual disk cannot be rebuilt because RAID 0 does not provide data redundancy.
- There is no hot spare assigned to the virtual disk. As long as the virtual disk is redundant, you can do the following to rebuild it:
  - Pull out the failed array disk and replace it. A rebuild will automatically start on the new disk.
  - Assign a hot spare to the virtual disk and then perform a rebuild.
- An array group contains both redundant and non-redundant virtual disks. On the PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4/IM, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers, a rebuild is not performed for an array disk that is used by both redundant and non-redundant virtual disks. In order to rebuild the redundant virtual disk, you need to delete the non-redundant virtual disk. Before deleting this disk, however, you can attempt to recover data from the failed array disk by forcing it back online.
- An array disk has been removed, and the system has not yet attempted to write data to the removed disk. In this case, the system will not recognize the removal of an array disk until it attempts a write operation to the disk. If the array disk is part of a redundant virtual disk, then the system will rebuild the disk after attempting a write operation. This situation applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/IM, 4e/DC, 4/Di, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers.

A Rebuild Completes with Errors on PERC Subsystem 1 Controllers

This section applies to PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4/IM, 4e/Si, 4e/Di, and CERC ATA100/4ch controllers

In some situations, a rebuild may complete successfully while also reporting errors. This may occur when a portion of the disk containing redundant (parity) information is damaged. The rebuild process can restore data from the healthy portions of the disk but not from the damaged portion.

When a rebuild is able to restore all data except data from damaged portions of the disk, it will indicate successful completion while also generating event 690. The rebuild may also report sense key errors. In this situation, take the following actions to restore the maximum data possible:

1. Back up the degraded virtual disk onto a fresh (unused) tape.
   - **If the backup is successful.** If the backup completes successfully then the user data on the virtual disk has not been damaged. In this case, you can continue with step 2.
   - **If the backup encounters errors.** If the backup encounters errors then the user data has been damaged and cannot be recovered from the virtual disk. In this case, the only possibility for recovery is to restore from a previous backup of the virtual disk.
2. Perform a Check Consistency on the virtual disk that you have backed up onto tape.
3. Restore the virtual disk from the tape onto healthy array disks.

Rebuilding a Virtual Disk on the CERC SATA1.5/2s Controller

On the CERC SATA1.5/2s controller, you can rebuild a RAID 1 virtual disk that has a global hot spare. The following considerations apply to rebuilding a virtual disk on this controller:

- If an array disk in the RAID 1 virtual disk fails and you replace the array disk with a brand new array disk that does not contain a previous configuration, then you need to assign the new array disk as a global hot spare. Once you have done so, the rebuild begins automatically. See Replacing a Failed Drive in a RAID 1 on a CERC SATA1.5/2s for more information on replacing the failed array disk.

- If an array disk in the RAID 1 virtual disk fails and you replace the array disk with an array disk that contains a previous configuration, then you need to delete the previous configuration. The previous configuration will be displayed as a virtual disk in the Array Manager tree view. After you have deleted the previous configuration, assign the new array disk as a global hot spare. When you have done so, the rebuild begins automatically. See Replacing a Failed Drive in a RAID 1 on a CERC SATA1.5/2s for more information on replacing the failed array disk.
- If one array disk in the RAID 1 virtual disk becomes unavailable because it is removed or the cable is disconnected, and the other array disk in the RAID 1 continues to update its data, then the data on the two array disks will be out of sync. Once the removed or disconnected array disk is restored, the virtual disk will be in Degraded state. In this situation, the **Rebuild** menu option for the virtual disk is enabled. You can right-click the virtual disk and select **Rebuild** in order to update the data on the restored array disk.

Receive Event 691 with "BBR Error" or "Error with Sense," or "Medium Error" Message

Event 691 is generated when a portion of a disk is damaged. This damage is discovered when the controller performs an operation that requires scanning the disk. Examples of operations that may result in event 691 are as follows:

- Consistency check
- Rebuild
- Virtual disk format
- I/O

If you receive event 691 as the result of doing a rebuild or while the virtual disk is in a degraded state, then data cannot be recovered from the damaged disk without restoring from backup. If you receive event 691 under circumstances other than a rebuild, then data recovery may be possible. The following describes each of these situations.

Event 691 Received during a Rebuild or while a Virtual Disk is Degraded

Do the following if you receive event 691 during a rebuild or while the virtual disk is in a degraded state:

1. Replace the damaged drive.
2. Create a new virtual disk and allow the virtual disk to completely resynchronize. To monitor the progress of the resynchronization, select the parent of the virtual disk in the Array Manager tree view and look at the status in the right pane. While the resynchronization is in progress, the status of the virtual disk will be **Resynching**.
3. Restore data to the virtual disk from backup.

Event 691 Received while Performing I/O, Consistency Check, Format, or Other Operation

If you receive event 691 while performing an operation other than a rebuild, you should replace the damaged disk immediately to avoid data loss.

Do the following:

1. Back up the degraded virtual disk to a fresh (unused) tape.
2. Replace the damaged disk.
3. Do a rebuild.

Cannot create a virtual disk (option is inactive)

Check:

- How many virtual disks exist? You can create a maximum of 8 virtual disks on one PERC 2/SC or PERC 2/DC controller and a maximum of 16 virtual disks on a CERC ATA100/4ch controller. You can create a maximum of 24 virtual disks on a PERC 2, PERC 2/Si, PERC 3/Si, PERC 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controller, and 40 virtual disks on a PERC 3/SC, 3/DCL, PERC 3/DC, PERC 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, or 4e/Di controller.
- Is there adequate unallocated space on the disk? You must have adequate available disk space to create a virtual disk.

Cannot create a RAID-5 volume

Check:

- Is there adequate unallocated space on three or more disks? You must have at least three disks to create a RAID-5 volume.

- Are you using a Windows 2000 Professional machine? You cannot create a RAID-5 volume on this machine. This restriction applies to software RAID only (that is, dynamic volumes). You need to use a Windows 2000 Server machine. Another option is that you can implement a hardware RAID-5 virtual disk and then create a volume on the virtual disk.

Cannot create a mirror

Check:

- Is there adequate unallocated space on two different disks? You must have two disks to create a mirrored volume. You also can create a mirror only on a simple or spanned volume.
- Are you using a Windows 2000 Professional machine? You cannot create a RAID-1 volume on this machine. This restriction applies to software RAID only (that is, dynamic volumes). You need to use a Windows 2000 Server machine. Another option is that you can implement a hardware RAID-1 virtual disk and then create a volume on the virtual disk.

When expanding the Disks object, error icons appear

Situation:

Windows is not aware of the status of these disks. Most likely, the virtual disks that were associated with these have been deleted.

Check:

To remove these error status icons from the **Disks** object, the computer must be restarted to allow Windows to find the current information.

Situation:

If the type of disk shows **No Signature**, you need to write a signature to the disk. When creating a new virtual disk, the software must write a signature to the virtual disk that prepares it for use. This signature is not written automatically in case this disk has been merged from another operating system and the configuration information needs to be kept intact.

Check:

For instructions on writing a disk signature, see the section “Write a Disk Signature” in the “Disks” chapter.

Missing Disk displays error icon

The corresponding virtual disk has been removed, or the disk has been rendered inactive because of a problem.

Check:

- If the corresponding virtual disk for this disk has been deleted, select **Remove Disk** and remove the disk from list of disks.
- If the disk has been rendered inactive because of some problem, see To bring a disk that is Offline and Missing back online

Once you have repaired the disk, controller, or cable problem, you need to:

1. **Rescan** to see the disk within Array Manager. If Array Manager finds the disk, this should bring the disk Online. If Array Manager does not find the disk, a reboot may be required.

2. **Reactivate Disk** to bring all the volumes on the disk to the Healthy status.

Read and Write Operations Experience Problems

If the system is hanging, timing out, or experiencing other problems with read and write operations, then there may be a problem with the adapter cables or a SCSI device. For more information, see the "Cables attached correctly" and "Isolate SCSI device problems" sections.

Problems after Installing a PERC 2/SC or 2/DC Controller

If you install a PERC 2/SC or 2/DC controller after you have already installed Array Manager, you may experience problems with Drive mapping, system hangs, and other performance problems. Reinstall Array Manager to resolve these problems.

I/O Stops on a Channel Redundant Channel

If you have implemented channel redundancy on a PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, or 4e/Di controller, a failure of one channel causes I/O to stop on the other channels included in the channel redundant configuration. The resolution to this problem is described in the "Considerations for Implementing Channel Redundancy on PERC 3/DCL, 3/DC, 3/QC, 4/DC, 4e/DC, 4/Di, and 4e/Di controllers." section.

Note Channel redundancy only applies to controllers that have more than one channel and that attach to an external disk enclosure.

Error message that the connection to remote computer has terminated

The full message is: “The connection to the remote computer has terminated. Remote computer will be removed from view.” The remote computer that you were connected to has been disconnected from your console. Most often, there is a problem with the network connection and the transmissions timed out. This can also occur if the remote machine was restarted or the service on the remote machine was stopped.

Check:

Make sure that the remote machine is turned on and is available to the network, and that the service is started. Reconnect to the resource.

Error Message: The stripe depth is out of range

Array Manager displays “The stripe depth is out of range” error message when you attempt to apply a RAID 0 or RAID 5 to more array disks than the controller can support in a single virtual disk. For example, the PERC 4/SC and 4/DC controllers can support up to 32 array disks in a virtual disk when using RAID 0 or RAID 5. Attempting to create a RAID 0 or RAID 5 using more than 32 array disks on these controllers will cause this error message to be displayed.

Tree view object for PowerEdge RAID controller cannot be expanded after the software and driver are installed

The installation detects any drivers that you have installed for PowerEdge RAID controllers. If these drivers (and/or the card itself) are installed after the software is installed, support for the controller will need to be added.

Check:

Close the console. Open the Array Manager Utilities and check the box next to the appropriate controller. This action will restart the service, and the disks should be available the next time you launch the console.

An option is inactive

When an operation is inactive or dimmed in a menu, the task cannot be performed on the object at this time. Certain operations are valid only for certain types of objects. (For example: RAID levels that are not fault tolerant will not allow you to check the consistency of the virtual disk.) If there is a task currently running on that object, wait until it has finished and try again. Otherwise, the operation may not be appropriate at this time.

To bring a disk that is Offline and Missing back online

If this was a virtual disk, then check that the virtual disk still exists. If it no longer exists, use the **Remove Disk** command to remove the disk from the list of disks.

Repair any disk, controller, or cable problems and make sure that the physical disk is turned on, plugged in, and attached to the computer. From the **View** pull-down menu, select **Rescan**. The disk should change from Offline to Online, but the volumes remain Failed. (If they do not change to Online, you may need to reboot.) Right-click the disk and select **Reactivate Disk**. The disk status changes to Healthy. (You can also select each volume one at a time and select **Reactivate Volume**. It is recommended you do a **chkdsk**.

Caution When reactivating a volume, be aware that the volume's data is restored, even if it is stale, corrupt, or out-of-date. See "Reactivate a Dynamic Volume" for more information on the consequences of reactivating a volume.

If the disk status remains Offline and Missing and you determine that the disk has a problem that cannot be repaired, you can remove the disk from the system (using the **Remove Disk** command). However, before you can remove the disk, you must delete all volumes on the disk. You can save any mirrored volumes on the disk by removing the mirror that is on the Missing disk instead of the entire volume. Deleting a volume destroys the data in the volume, so you should remove a disk only if you are absolutely certain that the disk is permanently damaged and unusable.

To bring a disk that is Offline (not Missing) and is still named Disk # back online

Use the **Reactivate Disk** command to bring the disk back online. If the disk status remains Offline, check the cables and disk controller, and make sure that the physical disk is healthy. Correct any problems and try to reactivate the disk again. If the disk reactivation succeeds, any volumes on the disk should automatically return to the Healthy status.

A disk on a PERC 4/Di, 4e/Si, or 4e/Di controller does not return online after a Prepare to Remove

When you do a Prepare to Remove command on an array disk attached to a PERC 4/Di, 4e/Si, or 4e/Di controller, you may find that the disk does not display in the Array Manager tree view even after doing a rescan or a reboot.

In this case, do the following to redisplay the disk in the Array Manager tree view:

1. Manually remove and then replace the array disk.

2. Either do a Rescan or reboot the system.

A disk is marked as Foreign

The disk has been moved to your computer from another Windows computer and has not been set up for use. Only dynamic disks display this status. To add the disk so that it can be used, right-click the disk and select **Merge Foreign Disk**. All existing volumes on the disk will be visible and accessible.

Because a volume can span more than one disk (e.g., a mirrored volume), it is important that you first verify your disk configurations and then move the entire disk set that the volume is on. If only part of the disk set is moved, some of the volumes will show Failed Redundancy or Failed error condition.

Note When adding or moving a disk that already contains data to a new system, you must format the disk using the controller BIOS before the disk can be used by Array Manager. Until you format the disk using the controller BIOS, the disk will have an Offline or Foreign status and may generate other errors. Using the **Merge Foreign Disk** procedure will not correct this situation.

A Disk is Marked as Offline or Foreign after Upgrading to Dynamic with a PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, or CERC ATA100/4ch Controller

If you initialize a virtual disk that has been upgraded to dynamic, the status of the dynamic disk may change to “offline” or “foreign.” You can view a disk’s status by selecting the disk’s **General** tab. When using a PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, or CERC ATA100/4ch controller, you can resolve this problem by reverting the “offline” or “foreign” disk to a basic disk. See Reverting a Dynamic Disk to Basic

Virtual Disk Initialization Causes a Foreign, Offline, or Missing Disk

Because initializing a virtual disk destroys the data on the virtual disk, you may find that after initializing a virtual disk on a Windows system, a disk marked as "foreign" or "missing" is displayed under the Disks folder. In addition, initializing a virtual disk containing a dynamic disk changes the status of the dynamic disk to foreign or offline.

To reuse a Windows disk that is set to foreign or offline, right-click the disk and select **Merge Foreign Disk** or **Revert to Basic Disk** from the pop-up menu.

If the disk is marked as missing, right-click the disk and select **Remove Disk**.

Initialize Command on a Virtual Disk does not Appear to Erase the Volumes on the Disk

Cause of Problem

When running the Initialize Disk command on a system where Array Manager is installed without Disk and Volume management, the operating system may display stale data from before the operation. This situation is particularly likely to occur on Windows Server 2003.

Solution

Reboot the server after a destructive operation like initializing the disk so that the operating system stays synchronized with Array Manager.

A Disk's Functions become Inactive or It is Marked as Foreign or Offline after Upgrading to Dynamic with a PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s Controller

If you format a virtual disk that has been upgraded to dynamic, the disk functions may become inactive or the status of the disk may change to "foreign" or "offline." You can view a disk's status by selecting the disk's **General** tab. When using a PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s controller, you can resolve these problems by doing a global rescan.

The Online Help behaves strangely, or will not come up at all

The Help file uses a technology known as HTML Help, a Microsoft standard. Some software will attempt to update the core files with an older version of HTML Help and make Array Manager's Help file unusable. The required HTML Help update is located on the Array Manager CD-ROM in the Help Update folder. Double-click HHUPD.EXE and follow the instructions.

When attempting to bring up the Help file, Dr. Watson reports an Access Violation in HH.EXE

HH is Microsoft's HTML Help format, which reads precompiled HTML files for Array Manager's Help sections.

Check:

Delete the HH.DAT file in your Windows directory. Deleting this file will remove any customizations that have been made to your HTML help files.

During reboot, a message displays about a "corrupt drive," suggesting that you run autocheck

Let autocheck run, but do not worry about the message. Autocheck will finish and the reboot will be complete. If you have a large system (more than 1 gigabyte), this may take about 10 minutes.

When attempting to access a remote computer, you are denied access or get an error message

There are several situations where this occurs.

You are denied access and do not even get a connection login box

This occurs when you log in to the local computer originally as a local user, local administrator, or domain user and the remote computer is not in your domain or a trusted domain. The Windows security model does not allow you to have access under these circumstances. The workaround is to log in to your local computer with an account that has the same user name and password as an administrator account on the remote computer.

You are denied access after typing the login information in the connection box

Access can be denied here if you do not type in a user name and password that match a local or domain administrator account on the remote computer or if you mistype the login information.

“Connection Failed” message

If the remote computer is not on or there are network problems, you will get the message “Connection Failed.”

For a NetWare system, refer to “The Connection Failed” message displays when connecting to the NetWare server.

You are unable to connect to a Windows 2000 server with Disk Management after a client-only installation

Another situation where you may get an error message is when you have just done a client-only installation of Array Manager and you bring up the Array Manager client and attempt to connect to a remote server that has Windows 2000 Disk Management.

Array Manager assumes that its client will connect first to a remote server running Array Manager before connecting to a system running Windows 2000 Disk Management.

Once you connect to a server with Array Manager, you will then be able to connect successfully to a remote system running Disk Management.

Windows 2000 Disk Management is the disk and volume management program that comes with Windows 2000. Because Array Manager and Disk Management are related programs, Array Manager is able to remotely manage the storage on a Windows 2000 computer with Disk Management.

You are unable to connect to a NetWare server

If you are having problems connecting to a NetWare® server from a local machine, use the ping and nslookup TCP/IP network diagnostic tools to determine whether the NetWare server is accessible from the local machine and whether the system running the NetWare server has a valid DNS name. If the system running the NetWare server does not have a valid DNS name, then you can edit the Hosts file on the local machine with an entry for the system running the NetWare server. The Hosts file is located in the winnt/system32/drivers/etc directory. The entry in the Hosts file should consist of the IP address and the host or server name of the system running the NetWare server.

If you do not connect by using a valid DNS name or an entry in the Hosts file, then you will need to use the IP address.

When you want to connect to a NetWare server, Array Manager expects the server to be identified by one of three types of entries:

- A DNS entry name
- An IP address
- A host name from a Hosts file listing

If you identify the name of the machine by a NetWare server's name that is not one of the three items above, the connection will fail. It is suggested that the name assigned to the NetWare server be the same name as its DNS or Hosts file entry.

Note that the DNS and Hosts file entries do not allow for a computer name that consists of all numbers. In addition, the DNS name does not allow a computer name that starts with a number. If the NetWare server has a numeric name or a name that starts with a number, you can use the IP address to identify that server. You can also put quotation marks around the computer's name for the entry in DNS or the Hosts file (such as "12345").

The Hosts file has to be on the client computer (local machine) that has the Array Manager console.

In addition, connecting to a remote system requires that you have administrator authority on both the local and remote system.

Note Dell does not offer NetWare in Japan.

Problems Occur after NetWare Upgrades or Configuration Changes

Upgrading NetWare or making configuration changes may cause problems for Array Manager. The following sections discuss how to avoid or recover from these problems.

Compatibility Issues with NetWare 6.x QLCSFTE.CDM and DellSEP.CDM Drivers

There is an incompatibility between the QLCSFTE.CDM and DellSEP.CDM drivers in NetWare 6 and the enclosure management provided by Array Manager. In NetWare 6, the hardware discovery logic detects the Qlogic chips within the server and auto installs the QLCSFTE.CDM and DellSEP.CDM drivers. These drivers interfere with the Array Manager enclosure management.

To resolve this problem, the Array Manager installation process comments out the lines for loading the QLCSFTE.CDM and DellSEP.CDM drivers in the STARTUP.NCF file. You should not remove the ";(Commented out by install...)" designation within the STARTUP.NCF file or attempt to manually re-install these drivers. If you do restore the

QLCSFTE.CDM and DellSEP.CDM drivers, then the enclosure management features of Array Manager will not function on NetWare 6. When you uninstall Array Manager, the QLCSFTE.CDM and DellSEP.CDM drivers are restored.

If you make configuration changes on NetWare 6.5 after installing Array Manager, the HDETECT utility may reactivate the QLCSFTE.CDM or the DellSEP.CDM drivers by removing the ";(Commented out by install...)" designation within the STARTUP.NCF file. Once the QLCSFTE.CDM and DellSEP.CDM drivers are reactivated, Array Manager is disabled. The solution to this problem is to reinstall Array Manager. Examples of configuration changes that will cause this problem are installing a new RAID or NIC card or installing NWASPI.CDM.

Note Dell does not offer NetWare in Japan.

After creating a virtual disk with a PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, or CERC ATA100/4ch controller, the virtual disk does not appear under the Disks storage object

If there are no virtual disks configured at boot time on a PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, or CERC ATA100/4ch controller on Windows 2000, the Windows disk driver may not be loaded. The solution is to reboot after creating the first virtual disk or create the first virtual disk in the BIOS (use **Ctrl-m** to invoke the BIOS utility).

"The Connection Failed" message displays when connecting to the NetWare server

If you are trying to connect to the NetWare server with the Array Manager console, you may receive a "The Connection Failed" error message. There can be a variety of reasons for why the connection between the Array Manager console and the NetWare managed system fails. (See also "Connection Failed" message.)

To identify why the connection failed, perform the following steps:

1. Ping the NetWare server from the system running the Array Manager console. If this fails, then you are experiencing network problems.

2. Verify that the correct NDS tree, user ID, and password are correct for the target NetWare server. Also verify that the user ID has administrator rights.

3. Verify that the server name is included in some form of DNS (DNS server, hosts file, and so forth).

4. Verify that the server name does not start with a number. This can cause problems with DNS.

5. Restart the console if an entry for the NetWare server was just added to the hosts file.

Note You may be able to avoid this connection problem by using the NetWare server's IP address instead of the server.

Note Dell does not offer NetWare in Japan.

A Disk is Marked as Failed when Rebuilding in a Cluster Configuration

When a system in a cluster attempts to rebuild a failed disk but the rebuild fails, then another system takes over the rebuild. In this situation, you may notice that the rebuilt disk continues to be marked as failed on both systems even after the second system has rebuilt successfully. To resolve this problem, perform a rescan on both systems after the rebuild completes successfully.

Erroneous Status and Error Messages after a Windows Hibernation

Activating the Windows hibernation feature may cause Array Manager to display erroneous status information and error messages. This problem resolves itself when the Windows operating system recovers from hibernation.

Cannot Connect to Remote System from Windows Server 2003

Certain conditions must be met before you can connect to a remote system from Windows Server 2003. For a description of these conditions, see Remote Connection and Windows Server 2003

Event Description Missing from Event Details Message Box

Situation: The events log lists events, but Description Details message boxes have no text in the Description field.

If you have uninstalled Array Manager or modified, repaired, or removed installation components without clearing the event log first, you might have removed the component that generates the descriptions. See also "Managing the Event Log" in the "Event Monitoring" chapter.

"Unsafe Removal of Device" message displays after uninstalling Array Manager on Windows 2000

After uninstalling Array Manager on Windows 2000, you may see the following error message: "Unsafe Removal of Device" with the listed device as "ROOT\vxio\0000." Click OK, then reboot the server. There is no adverse effect to the system.

System Performance Problems

This section describes problems that may deteriorate system performance.

Unusual CPU Usage

You may notice unusual surges in your system's CPU usage. These surges may be caused by Array Manager's volume capacity monitoring. This function monitors NTFS volumes on the local server for the amount of space used. When the space used on an NTFS volume reaches 90%, a warning event is logged in the Array Manager and Windows event log. When the space used reaches 98%, an error event is logged.

If the surges in CPU usage pose a problem, you can disable volume capacity monitoring.

To disable volume capacity monitoring:

1. Launch the Array Manager Utilities by clicking **Start | Programs | Dell OpenManage Applications | Array Manager** and selecting **Array Manager Utilities**.
2. Deselect the **Volume Capacity Monitoring** check box on the **Windows** tab.
3. Click **Apply** and then **Close** to exit the Array Manager Utilities.

For information related to volume capacity monitoring, see the following:

- Volume Management Events
- Volume Troubleshooting Procedures
- Volume Capacity Monitoring

Dell PowerVault 660F and 224F Storage Systems Troubleshooting

This section presents possible problem situations with accompanying solutions for the Dell PowerVault 660F and 224F storage systems. The problem situations are organized as follows:

- Event Monitoring and Logging
- Fibre Channel RAID Controller Status Events
- Enclosure Status Events
- Drive Status Events
- Error Messages
- Other Problems
- Events Generated by the PowerVault 660F Storage System

The situations in the first three topics are categorized by their event number. A brief discussion of event messages is included at the beginning of this section in the topic Event Monitoring and Logging The fourth topic describes general problems not related to a specific event.

You will also find a full listing of the events associated with the Dell PowerVault 660F Fibre Channel RAID controller at the end of this section in the topic Events Generated by the PowerVault 660F Storage System

Event Monitoring and Logging

Event messages help identify significant incidents such as an array disk failure or an array disk addition. Event monitoring and logging starts when the Array Manager managed system starts up. If the managed system service stops on Windows or the Array Manager Service stops in NetWare, then event monitoring and logging stops. If array disks are S.M.A.R.T. (Self Monitoring Analysis and Reporting Technology) enabled, the RAID controllers check array disks for failure predictions, and if found, pass this information on to the Array Manager console. Array Manager immediately displays an alert icon on the array disk and also raises an alert under the Events tab and in the Windows event log. Windows has three event logs; Array Manager uses the application log.

Note When a controller's I/O is paused, Array Manager does not receive S.M.A.R.T. events.

Fibre Channel RAID Controller Status Events

The following incidents are included in this topic:

- Event 708, Rebuild stopped with error
- Event 840, Controller has been reset
- Event 857, Warm boot failed
- Event 858, Controller entered Conservative Cache mode
- Event 870, Killed partner
- Event 872, Controller boot ROM image needs to be reloaded
- Event 873, Controller is using default nonunique world-wide name
- Event 887, Back End Fibre Dead

Event 708, Rebuild stopped with error

Cause of Problem

Because of some unknown error on the Fibre Channel RAID controller, the rebuild has failed.

Solution

Try rescanning the controller: from the Array Manager tree view, click to expand the **Arrays** storage object, right-click the **PV660F Subsystem** object, and then select **Rescan** from the context menu that comes up. This action will update the controller status within the GUI.

If the controller has been removed and reinserted, check to see that the controller is inserted correctly: the DB9 connector should be located at the top of the module. For details, see the *Dell PowerVault 660F and 224F Storage Systems Service Manual* . Also, check that all cable connections are correctly and firmly connected. Try to rebuild again: right-click the **Array Disk** storage object in the tree view, and then select **Rebuild** from the context menu that comes up.

If controller and connections are correct and the problem continues, contact customer service.

Event 840, Controller has been reset

Cause of Problem

One of the following may have occurred on one of the Fibre Channel RAID controllers in the PowerVault 660F storage system:

- The controller failed.
- The controller was removed from the system.
- The controller has been powered off.

Solution

- If the controller failed, replace the controller according to the *Dell PowerVault 660F and 224F Storage Systems Service Manual*.
- If the controller was removed, reinsert it according to the *Dell PowerVault 660F and 224F Storage Systems Service Manual*.
- If the controller has been powered off, restore power.

Then access the Array Manager console, click to expand the **Arrays** storage object, right-click the **PV660F Subsystem** storage object, and select **Rescan** from the context menu that comes up. This action will update the controller status.

Event 857, Warm boot failed

Cause of Problem

A memory error was detected during the warm boot scan, which may result in possible data loss.

Solution

Power cycle the PV660F subsystem.

If the error persists:

- Remove and reinsert the controller. If the controller fails to come online, contact customer service.
- If needed, follow the procedures in the *Dell PowerVault 660F and 224F Storage Systems Service Manual* to replace the controller module.

Event 858, Controller entered Conservative Cache mode

Cause of Problem

- Enclosure Management is offline.
- Partner controller disabled (held in reset) or failed.
- BBU battery is low.
- Expand Capacity or Add Virtual Disk operation is in progress.
- Out-of-limit or over temperature.
- Non-redundant power.
- Power supply failed.

Solution

If Enclosure Management has been enabled, check to see whether one or more LS modules has failed. LS modules are cards installed in slots at the front of the enclosure. Each LS module has an SES processor that monitors environmental functions and a Loop Redundancy Circuit (LRC) function, which maintains the viability of the Fibre Channel loop. See the *Dell PowerVault 660F and 224F Storage Systems Installation and Troubleshooting Guide* for troubleshooting tips.

If the controller entered Conservative Cache mode because of user's intended action, proceed per user's intent. When finished, right-click the controller and select either **Enable Partner** controller or **Enable BBU** to exit Conservative Cache mode.

If the BBU battery is low, recondition the battery. If the battery needs to be replaced, see the *Dell PowerVault 660F and 224F Storage Systems Service Manual*. Then access the Array Manager console, and in the tree view, click to expand the **Arrays** storage object, right-click the **PV660F Subsystem** storage object to bring up the context menu, and select **Rescan**. This action will update the controller.

If there is an Expand Capacity or Add Virtual Disk operation in progress, wait until this activity has finished. Then access the Array Manager console, and in the tree view, click to expand the **Arrays** storage object, right-click the **PV660F Subsystem** storage object to bring up the context menu, and select **Rescan**. This action will update the controller.

See also "Conservative Cache Mode" in the "The Dell PowerVault 660F Storage System" chapter.

Event 870, Killed partner

Cause of Problem

- The controller is waiting for the **Enable Partner** command before completing the startup process. If the controller option **Auto Restore** is set, the **Enable Partner** command is not needed. If the Auto Restore option is not set, the partner controller will wait for the **Enable Partner** command.
- Couldn't allocate chunk of memory.
- SCSI communication failed.
- Mismatch in the number of array disk channels present.
- Mismatch in the number of host channels present.
- Mismatch in the firmware version.
- Mismatch in the firmware header type.
- Memory read of partner controller failed.
- Mismatch in cache memory size.
- Received **Disable Partner** controller command.
- Negotiation finished, but nexus not entered in time.
- Write-back sync to partner controller, channels 0-5.
- Mismatch in firmware build.
- Device cables are crossed.
- Partner controller removal detected while nexus active.
- Partner controller missing at negotiation time.
- BBU powerfailed before failover finished.
- BBU powerfailed before relinquish finished.
- Lock time-out.
- Lock SCSI failed.
- General ctoc (controller to controller) message failure.
- Failed for some unknown reason.

Note Nexus refers to the state in which both redundant controllers are in communication. In this state, each controller can copy write-back data to its partner controller and can determine whether the other controller is operating.

Solution

Go to the Array Manager console and click to expand the **Arrays** storage object, right-click the **PV660F Subsystem** storage object, and select **Rescan** from the context menu that comes up. This action will update the partner controller. If the situation does not improve, try one of the following (**Rescan** as before when necessary after troubleshooting the partner controller):

- The controller is shipped with **Auto Restore** set. To see whether the partner is enabled, right-click the controller. If **Enable Partner** is on the menu, then the partner controller is disabled: click **Enable Partner** to enable the partner controller. If **Disable Partner** is on the menu, then the partner controller is enabled; use the following troubleshooting solutions to fix the event.
- Unless a firmware or memory mismatch is suspected, reboot the disabled partner controller. If a firmware or memory mismatch is the issue, see below.
- If memory problems exist, replace the memory module if necessary. See the *Dell PowerVault 660F and 224F Storage Systems Service Manual*.
- Check the firmware version of the good controller: in the Array Manager tree view right-click the **Controller** and click the **Properties** command on the context menu that appears. Replace the failed controller; see the *Dell PowerVault 660F and 224F Storage Systems Service Manual*. Go to the Array Manager console and click to expand the **Arrays** storage object, right-click the **PV660F Subsystem** storage object, and select **Rescan** from the context menu that comes up. This action will update the controller status. Check the firmware version of the new controller; if there is a mismatch, update both controllers with the latest firmware. See "Firmware version mismatch" in this chapter for instructions.
- Check the memory size of the good controller: right-click the **Controller** in the tree view and click the **Properties** command on the context menu that appears. Replace the failed controller with a controller of the same memory size; see the *Dell PowerVault 660F and 224F Storage Systems Service Manual*. Go to the Array Manager console and click to expand the **Arrays** storage object, right-click the **PV660F Subsystem** storage object, and select **Rescan** from the context menu that comes up. This action will update the controller status.
- Check that the partner controller is present. If it is not, reinstall the partner controller. See the *Dell PowerVault 660F and 224F Storage Systems Installation and Troubleshooting Guide* for slot module installation instructions.
- If the user gave the **Disable Partner** command, continue per user's intent.

If none of these solutions apply, contact customer service.

Event 872, Controller boot ROM image needs to be reloaded

Cause of Problem

The Media Access Control (MAC) address is corrupted.

Solution

Replace the controller according to the *Dell PowerVault 660F and 224F Storage Systems Service Manual*. Then go to the Array Manager console and click to expand the **Arrays** storage object, right-click the **PV660F Subsystem** storage object, and select **Rescan** from the context menu that comes up. This action will update the controller status.

Event 873, Controller is using default nonunique world-wide name

Cause of Problem

The Media Access Control (MAC) address has been lost or was not set.

Solution

Replace the controller according to the *Dell PowerVault 660F and 224F Storage Systems Service Manual*. Then go to the Array Manager console and click to expand the **Arrays** storage object, right-click the **PV660F Subsystem** storage object, and select **Rescan** from the context menu that comes up. This action will update the controller status.

Event 887, Back End Fibre Dead

Cause of Problem

- Lost access to data on the Fibre Channel. The Fibre Channel cable may have been disconnected.
- LS module has failed.
- I/O module has failed.

Note LS modules are cards installed in slots at the front of the enclosure. Each LS module has an SES processor that monitors environmental functions and a Loop Redundancy Circuit (LRC) function, which maintains the viability of the Fibre Channel loop.

Solution

Check that the Fibre Channel cable is connected to the controller and the switch box. If not, reconnect it as necessary. If the Fibre Channel cable is connected, try replacing the cable. See the *Dell PowerVault 660F and 224F Storage Systems Installation and Troubleshooting Guide* for troubleshooting the LS and I/O modules.

When troubleshooting is complete, go to the Array Manager console and click to expand the **Arrays** storage object, right-click the **PV660F Subsystem** storage object, and select **Rescan** from the context menu that comes up. This action will update the Fibre Channel.

If problem persists, contact customer service.

Enclosure Status Events

The following incidents are included in this topic:

- Event 818 and 820, Fan failure or Fan is not present
- Event 821 and 823, Power supply failure or Power supply is not present
- Event 825, Temperature is above working limit
- Event 828, Enclosure access is Critical
- Event 831, Enclosure Soft Addressing detected
- Event 844, BBU power low
- Event 863, BBU Recondition suggested

Event 818 and 820, Fan failure or Fan is not present

Cause of Problem

- A fan is bad.
- A fan is not present.
- A possible broken cable connection.
- The LS (Loop Redundancy circuit/SCSI enclosure services) module connection may be broken, or the management hardware is bad.

Solution

To locate a fan, right-click the bad fan and click **Properties**. The Enclosure ID field indicates the ID number of the enclosure where this fan is located. Be aware that the enclosure ID number displayed by Array Manager does not match the Enclosure ID set on

the switch module ID indicator on the back of the PowerVault 660F or 224F enclosures. When the switch module ID indicators on all the enclosures are configured properly, the enclosure ID numbers start at 0 and continue sequentially through 7. The enclosure ID numbers displayed by Array Manager, however, start at 1 and continue sequentially through 8. For this reason, the enclosure ID number displayed by Array Manager will be one greater than the number set on the switch module ID indicator on the back of the PowerVault 660F or 224F enclosures. For example, if the switch module ID indicator identifies the enclosure as 0, the Array Manager enclosure ID will be 1.

See the *Dell PowerVault 660F and 224F Storage Systems Installation and Troubleshooting Guide* for information on how to troubleshoot the Advanced Cooling Module (ACM). See the *Dell PowerVault 660F and 224F Storage Systems Service Manual* for part replacement procedures.

After troubleshooting or replacing the ACM, go to the Array Manager console and click to expand the **Arrays** storage object, right-click the **PV660F Subsystem** storage object, and select **Rescan** from the context menu that comes up. This action will update the fan status within the GUI.

Event 821 and 823, Power supply failure or Power supply is not present

Cause of Problem

- The LS (Loop Redundancy circuit / SCSI enclosure services) module connection may be broken.
- The management hardware is bad.
- The power supply is not present.
- The cable connection is broken.
- The power supply is bad.

Solution

See the *Dell PowerVault 660F and 224F Storage Systems Installation and Troubleshooting Guide* for information on how to troubleshoot the power supply. See the *Dell PowerVault 660F and 224F Storage Systems Service Manual* for part replacement procedures.

Note LS modules are cards installed in slots at the front of the enclosure. Each LS module has an SES processor that monitors environmental functions and a Loop Redundancy Circuit (LRC) function, which maintains the viability of the Fibre Channel loop.

After troubleshooting or replacing the power supply, go to the Array Manager console and click to expand the **Arrays** storage object, right-click the **PV660F Subsystem** storage object, and select **Rescan** from the context menu that comes up. This action will update the power supply status within the GUI.

Event 825, Temperature is above working limit

Cause of Problem

- The room temperature is high.
- A fan is bad.

Solution

Check all fans to see whether they are functioning properly. If yes, check that the ambient temperature is within limit. If necessary, adjust the room temperature. If the problem persists, power-cycle the system. If this does not solve the problem, replace the affected Advanced Cooling Module (ACM). See the *Dell PowerVault 660F and 224F Storage Systems Installation and Troubleshooting Guide* for information on how to troubleshoot the ACM. See the *Dell PowerVault 660F and 224F Storage Systems Service Manual* for part replacement procedures.

After fixing the temperature problem, go to the Array Manager console and click to expand the **Arrays** storage object, right-click the **PV660F Subsystem** storage object, and select **Rescan** from the context menu that comes up. This action will update the temperature and / or the fan status within the GUI.

Event 828, Enclosure access is Critical

Cause of Problem

The LS module connection may be broken or the management hardware is bad. LS modules are cards installed in slots at the front of the enclosure. Each LS module has an SES processor that monitors environmental functions and a Loop Redundancy Circuit (LRC) function, which maintains the viability of the Fibre Channel loop

Solution

Check to see that the LS module is enabled.

Follow the *Dell PowerVault 660F and 224F Storage Systems Installation and Troubleshooting Guide* for troubleshooting an LS module. For part replacement, see the *Dell PowerVault 660F and 224F Storage Systems Service Manual*.

After resolving the hardware problem and providing corrective action, go to the Array Manager interface and click to expand the **Arrays** storage object, right-click the **PV660F Subsystem** storage object, and select **Rescan**. This action will update the enclosure status within the GUI.

Event 831, Enclosure Soft Addressing detected

Cause of Problem

The enclosure has duplicate loop IDs (soft addressing).

Solution

Make sure shelf ID switches on all PV660s and PV224s in the subsystem are set to different numbers. See the *Dell PowerVault 660F and 224F Storage Systems Installation and Troubleshooting Guide* on how to set shelf IDs.

Event 844, BBU power low

Cause of Problem

A battery backup unit (BBU) with a low charge was found on the controller.

Solution

If this message occurs without power failure, replace the BBU.

To replace the BBU, see the *Dell PowerVault 660F and 224F Storage Systems Service Manual* for part replacement procedures.

The BBU requires two reconditioning cycles prior to first time use. This reconditioning process will take several hours and cannot be interrupted. Refer to "Recondition" in the chapter "The Dell PowerVault 660F Storage System" for detailed instructions on performing a BBU recondition.

After troubleshooting or replacing the BBU, go to the Array Manager console and click to expand the **Arrays** storage object, right-click the **PV660F Subsystem** storage object, and select **Rescan** from the context menu that comes up. This action will update the BBU status within the GUI.

Event 863, BBU Recondition suggested

Cause of Problem

- The controller detects a new battery.
- The charge of the battery dropped below 85% more than 30 times.
- The battery was temporarily disconnected from the controller from a fan pack pull.
- The controller was removed and replaced.

Solution

This message will appear each time the controller is (re)booted until the BBU is manually reconditioned. In this state, the battery is still protecting the controller cache, but the maximum level of the battery may be uncertain. To re-establish the maximum level, you can perform a BBU recondition at your convenience using the procedure described in the "Recondition" section. Be aware that the reconditioning process takes several hours. The recondition process must complete its full cycle without interruption. Interrupting the recondition will invalidate the recondition, thus the next invocation will start at the beginning of the cycle. In other words, the recondition process cannot be suspended and then restarted in the middle of the process.

The BBU requires an automatic and a manual recondition before the maximum level will be accurate. Until the first manual recondition is completed, the maximum level from the battery specifications is used. The automatic recondition will start if one of the following event occurs:

- When powering up a new PowerVault 660F storage system for the first time
- When reinstalling or replacing a fan pack into the PowerVault 660F storage system
- When reinstalling or replacing a PowerVault 660F storage system

After the automatic recondition is complete, Array Manager will show "BBU Recondition Suggested" when the system is rebooted.

After manually reconditioning the BBU, go to the Array Manager console and click to expand the **Arrays** storage object, right-click the **PV660F Subsystem** storage object, and select **Rescan** from the context menu that comes up. This action updates Array Manager on the status of the BBU.

If the "BBU Recondition Suggested" error message continues to be displayed after you have manually reconditioned the BBU twice, then the battery may be defective and require replacement. In this case, see the *Dell PowerVault 660F* and *224F Storage Systems Service Manual* for part replacement procedures.

Drive Status Events

The following incidents are included in this topic:

- Event 702, Hard disk error found
- Event 703, Hard disk PFA condition found
- Event 707, Rebuild is cancelled
- Event 709, Rebuild stopped with error: New device failed
- Event 710, Rebuild stopped because logical drive failed
- Event 711, A hard disk has failed
- Event 718, SCSI command timeout on hard device
- Event 721, Parity error found
- Event 749, Physical device status changed to Offline
- Event 753, Physical device failed to start
- Event 756, Physical drive missing on startup
- Event 758, Physical drive is switching from one channel to another
- Event 765, Consistency check on logical drive error
- Event 766, Consistency check on logical drive failed
- Event 767, Consistency check failed due to physical device failure
- Event 768, Logical drive has been made offline
- Event 769, Logical drive is Critical
- Event 781, Logical drive initialization failed
- Event 786, Expand Capacity stopped with error
- Event 787, Bad blocks found

Procedure for Replacing a Drive

Many of the incidents in this section can be resolved by replacing one or more drives. The procedure for replacing a drive is as follows:

- Identify the failed drive's location. If necessary, refer to the *Dell PowerVault 660F and 224F Storage Systems Installation and Troubleshooting Guide*.
- Prepare the drive for removal. See the topic "Prepare to Remove" in the chapter The Dell PowerVault 660F Storage System

- Replace the drive; see the removing and replacing hard-disk drives section in the *Dell PowerVault 660F and 224F Storage Systems Service Manual*.
- The rebuild process begins immediately if the Automatic Rebuild Management and Enclosure Management controller options are enabled. See "Advanced Fibre Channel RAID Controller Options" in the chapter The Dell PowerVault 660F Storage System If these controller options are not enabled, issue a **Rebuild** command. See "Rebuild" in the chapter "The Dell PowerVault 660F Storage System" for instructions.
- Rescan the subsystem. Click to expand the **Arrays** storage object, right-click the **PV660F Subsystem** storage object, and select **Rescan** to verify that the replaced drive has been recognized.

Event 702, Hard disk error found

Cause of Problem

- A bad sector was found on the physical media.
- A mechanical failure on the device.
- The Host SCSI device detected illegal instruction.
- Target device generated unknown phase sequence.

Solution

If the virtual disk is offline, try forcing it online with the **Force Online** command. Right-click the disk and select **Force Online** from the context menu that appears. See "Force Online" in the chapter "The Dell PowerVault 660F Storage System" for detailed instructions.

If you cannot force the virtual disk online, remove and replace the affected hard drive. See the topic "Procedure for Replacing a Drive" at the beginning of this section for instructions on replacing a drive.

If you still get a hard disk error after replacing the drive, contact customer service.

Event 703, Hard disk PFA condition found

Cause of Problem

A hard disk predicted a future failure condition. This disk may fail soon.

Solution

Replace and rebuild the drive. See the topic "Procedure for Replacing a Drive" at the beginning of this section for instructions on replacing and rebuilding a drive.

Event 707, Rebuild is cancelled

Cause of Problem

- User cancelled rebuild.
- Higher-priority rebuild started.

Solution

If the rebuild was cancelled, that disk will remain in an unusable state until a successful rebuild has been performed on it.

See "Rebuild" in the chapter "The Dell PowerVault 660F Storage System" for instructions.

Event 709, Rebuild stopped with error: New device failed

Cause of Problem

- A new array disk failed.
- A new array disk may not be compatible with controller hardware or firmware.

Solution

Try to rebuild the drive again. See "Rebuild" in the chapter "The Dell PowerVault 660F Storage System" for instructions.

If the problem persists, replace the disk drive. It may be helpful to look at the other events that were generated in order to identify the malfunctioning drive.

See the topic "Procedure for Replacing a Drive" at the beginning of this section for instructions on replacing a drive.

Event 710, Rebuild stopped because logical drive failed

Cause of Problem

- At least one or more array disks failed in the virtual disk.

- Bad data table overflow.

Solution

It may not be possible to recover from this error; therefore, you may lose your virtual disk.

Try to rebuild the virtual disk. See "Rebuild" in the chapter "The Dell PowerVault 660F Storage System" for instructions.

If the problem persists, contact customer service.

Event 711, A hard disk has failed

Cause of Problem

A hard disk failed because the user either changed the status to Offline or removed the hard disk.

Solution

A drive is usually manually taken offline to replace it. If the drive was physically removed from the enclosure, replace and rebuild the drive (using a drive at least as large as the other disk drives in the virtual disk).

See the topic "Procedure for Replacing a Drive" at the beginning of this section for instructions on replacing and rebuilding a drive.

Event 718, SCSI command timeout on hard device

Cause of Problem

- An array disk has been removed.
- An array disk has failed.

Solution

If the drive has been removed or has failed, replace the drive.

If the time-out cannot be reset on the existing array disk, replace the disk.

See the topic "Procedure for Replacing a Drive" at the beginning of this section for instructions on replacing and rebuilding a drive.

It may be necessary to do a complete reboot after the drive is replaced.

If you need more help, contact customer service.

Event 721, Parity error found

Cause of Problem

- An array disk did not generate proper parity.
- A cable connection is bad.
- Another array disk interfered.

Solution

Check all cables, making sure they are correctly and firmly connected and that none are crossed. See the *Dell PowerVault 660F and 224F Storage Systems Installation and Troubleshooting Guide* for cabling procedures.

Replace the affected disk. See the topic "Procedure for Replacing a Drive" at the beginning of this section for instructions on replacing and rebuilding a drive.

Event 749, Physical device status changed to Offline

Cause of Problem

The array disk encountered too many errors, causing the drive to fail and its status to change to Offline.

Solution

It is not possible to recover this physical drive. Replace the disk drive. See the topic "Procedure for Replacing a Drive" at the beginning of this section for instructions on replacing and rebuilding a drive.

If the replacement drive still does not work, contact customer service.

Event 753, Physical device failed to start

Cause of Problem

Drive failed to spin up during controller bootup.

Solution

Check that the new array disk is seated properly. If not, remove and reinsert the disk.

If the problem persists, see the *Dell PowerVault 660F and 224F Storage Systems Installation and Troubleshooting Guide* on how to troubleshoot the Fibre Channel hard disk drives.

When completed, go to the Array Manager console and click to expand the **Arrays** storage object, right-click the **PV660F Subsystem** storage object, and select **Rescan**. This action will update the array disk status within the GUI.

Event 756, Physical drive missing on startup

Cause of Problem

A previously configured array disk no longer appears in the Array Manager GUI.

Solution

Make sure that all enclosures are powered on.

Remove and reinsert the physical drive. Then go to the Array Manager console and click to expand the **Arrays** storage object, right-click the **PV660F Subsystem** storage object, and select **Rescan** to verify that the replacement drive has been recognized.

If the drive is still missing or not found, try replacing the drive. See the topic "Procedure for Replacing a Drive" at the beginning of this section for instructions on replacing and rebuilding a drive.

If the replacement drive does not work, contact customer service.

Event 758, Physical drive is switching from one channel to another

Cause of Problem

Communication to a drive on a particular channel has failed.

Solution

If this event appears for all existing drives, then a Loop ID problem may be present. See the *Dell PowerVault 660F and 224F Storage Systems Installation and Troubleshooting Guide* on how to troubleshoot the I/O module.

If the drive has failed, replace the drive. See the topic "Procedure for Replacing a Drive" at the beginning of this section for instructions on replacing and rebuilding a drive.

If the replacement drive does not work, contact customer service.

Event 765, Consistency check on logical drive error

Cause of Problem

- Inconsistent data was found.
- Bad sectors were found.
- An array disk reliability problem.

Solution

Try performing a Consistency Check again. If the problem persists, replace the disk drive. It may be helpful to look at the other events that were generated in order to identify the malfunctioning drive. See the topic "Procedure for Replacing a Drive" at the beginning of this section for instructions on replacing and rebuilding a drive.

Event 766, Consistency check on logical drive failed

Cause of Problem

- A virtual disk became Critical (lost one fault tolerant drive).
- A virtual disk failed (lost two fault tolerant drives).

Note A virtual disk that is Critical will be shown as Degraded status within the Array Manager console.

Solution

Try performing a consistency check again. If the problem persists, replace the disk drive(s). It may be helpful to look at the other events that were generated in order to identify the malfunctioning drive(s). See the topic "Procedure for Replacing a Drive" at the beginning of this section for instructions on replacing and rebuilding a drive.

Event 767, Consistency check failed due to physical device failure

Cause of Problem

An array disk failed.

Solution

Replace the affected disk. See the topic "Procedure for Replacing a Drive" at the beginning of this section for instructions on replacing and rebuilding a drive.

Event 768, Logical drive has been made offline

Cause of Problem

If you have a non-fault-tolerant virtual disk, a single array disk failure may have caused the virtual drive to go offline. If you have a fault tolerant virtual disk, multiple array disk failures may have caused the virtual drive to go offline.

Solution

Try making the virtual drive Online.

Verify through the LED lights that power is supplied to the enclosure.

Identify the location of the failed drive(s). If necessary, refer to the *Dell PowerVault 660F and 224F Storage Systems Installation and Troubleshooting Guide*.

Replace the array disk(s) if necessary. See the topic "Procedure for Replacing a Drive" at the beginning of this section for instructions on replacing and rebuilding a drive.

It may not be possible to recover from this error. Contact customer service.

Event 769, Logical drive is Critical

Cause of Problem

One fault tolerant virtual disk has been degraded.

Solution

Replace the array disk. See the topic "Procedure for Replacing a Drive" at the beginning of this section for instructions on replacing and rebuilding a drive.

Event 781, Logical drive initialization failed

Cause of Problem

- One or more array disks failed (fault tolerant disks).
- Controller has been removed or powered off.

Solution

Reinsert the controller and power on the system and/or the controller.

Initialize the virtual disk manually. See "Initialize" in the chapter "The Dell PowerVault 660F Storage System" for instructions.

When completed, go to the Array Manager console and click to expand the **Arrays** storage object, right-click the **PV660F Subsystem** storage object, and select **Rescan** to verify that the virtual disk has been initialized and is recognized.

Replace the array disk. See the topic "Procedure for Replacing a Drive" at the beginning of this section for instructions on replacing and rebuilding a drive.

Event 786, Expand Capacity stopped with error

Cause of Problem

- An error was detected during phase one of the migration process.
- A drive has failed.

Solution

Replace the array disk. See the topic "Procedure for Replacing a Drive" at the beginning of this section for instructions on replacing and rebuilding a drive.

See "Expand Capacity" in the chapter "The Dell PowerVault 660F Storage System" for more information.

Event 787, Bad blocks found

Cause of Problem

A bad sector was found on an array disk during one of the following operations: consistency check, rebuild, or RAID expansion.

Solution

- Run a consistency check.
- Restore data from a backup.

For information on consistency check and rebuild functions, see "Check Consistency (PV660F)" and "Rebuild" in the "The Dell PowerVault 660F Storage System" chapter.

If the problem persists, replace the hard disk. See the topic "Procedure for Replacing a Drive" at the beginning of this section for instructions on replacing and rebuilding a drive.

Error Messages

This section describes error messages for the PowerVault 660F storage subsystem that may be displayed by Array Manager.

Degraded, LS module chip missing or failed

In addition to a missing or failed LS module chip, this error message may also be displayed when the SES firmware of the active LS module is version 1.9.5 while the SES firmware for the failover LS module is version 1.3b2. The SES firmware for both LS modules should be at the same level. You can resolve this error by upgrading the SES firmware so that both LS modules have the same SES firmware version.

Unsafe Removal of Device

When assigning virtual disk 0 (LUN 0) to a Windows server, a Windows "Unsafe Removal of Device" message may display for the PV660F PSEUDO SCSI Disk Device.

You can disregard the Disk Removed message. This is caused by the PV660F Pseudo Disk (used to manage the Fibre Channel array when no virtual disks are assigned) being replaced by the newly assigned virtual disk.

Disk Removed

A Windows Disk Removed dialog may display on the server under the following conditions:

- In a storage area network (SAN) hardware configuration, you do a console rescan when assigning virtual disk 0 of a PowerVault 660F array to the server.
- You create virtual disk 0 in a direct-attached hardware configuration.

You can disregard the Disk Removed message. This is caused by the PV660F Pseudo Disk (used to manage the Fibre Channel array when no virtual disks are assigned) being replaced by the newly assigned virtual disk.

Other Problems

This section contains the following general problem situations:

- Cannot add a hot spare

- Cannot expand a disk group
- After creating a virtual disk, cannot locate it in the Disks folder
- Cannot delete a volume
- Cannot delete a virtual disk
- Firmware version mismatch
- Unable to manage objects below the controller object
- No objects appear under Physical Array and Logical Array
- Unable to reset the PowerVault 660F storage system
- Application Transparent Failover autofails all LUNs on the PV650F

How to Detect, Recondition, and Replace a Failed Battery Backup Unit

The following sections describe how to troubleshoot problems with the PowerVault 660F battery backup unit (BBU).

How to Detect a Failed BBU

The PowerVault 660F generates error messages that indicate the battery is either not installed, not recognized, or too low on power to function properly. Here are a few examples of the types of event messages you may receive:

- BBU is non-functional and should be replaced. The BBU may have been physically removed.
- BBU is not currently installed or detected.
- Battery power is currently insufficient to sustain the write-back cache for the user's configured amount of time. Note that this event can be generated if the battery is new and has not been charged yet and also during a recondition cycle as the battery is discharged and recharged.

How to Recondition a Failed BBU

Certain conditions require that the BBU needs to be reconditioned before it can be used. If the power went off or if you are installing a new controller with a new BBU, then the BBU needs to be reconditioned, sometimes referred to as recharged. However, before you start a recondition process, you need to be aware that this operation cannot be interrupted for any reason. Do not try to initiate a fast charge or another recondition cycle, once you have started a recondition process. If for any reason, there is an interruption in the cycle, the process will need to be completely restarted.

To recondition a battery, right-click on a controller, click on C0 Battery or C1 Battery, and click on Recondition.

How to Replace a Failed BBU

When all attempts to recharge or recondition the BBU have failed, it is time to replace it. However, before you remove a BBU, it is best to put the system into a ready for Shutdown State. Then replace the component and clear the Shutdown State to re-enable caching. The BBU is contained in the Fan CRU (Dell/Euro) and can be replaced separately from the controller without shutting down the system.

Note Replacing a BBU or a controller can be done in a "hot" condition–you can leave the main system running and shut off power to or disable only the component being replaced.

For more information, refer to "Battery Properties" in the "The Dell PowerVault 660F Storage System" chapter.

How to Detect and Replace a Failed Controller

The PowerVault 660F has two controllers. Each controller is encased in its own canister and each controller is on one side of the enclosure. If you receive an event notice that indicates a controller has failed or is not responding, there are certain steps to follow to shutdown and replace the bad controller while maintaining accessibility and functionality to your system. If one controller goes down, the other controller will handle the same functions, (fault-tolerance with the controllers) although performance will be degraded until the bad controller is replaced and brought on line.

Causes of Controller Failure

- Debug dump to the controller or its partner controller
- Mechanical failure on the controller
- Improper seating of the controller
- Faulty cable connections
- A change in the controller's World Wide Name (WWN) can make storage unavailable

Solutions

Before performing any of these procedures, be sure the good controller is functioning at full performance. You may now remove the faulty controller (pull it out of the enclosure).

Note You may test the faulty controller by removing it and inserting it into another system for diagnostics. In any case, be sure that the good controller is working. Replace the bad controller with a new one as soon as possible.

- Check to see if the faulty controller will boot up on its own by pulling it out of the enclosure, check the connectors, and push it back into place, making sure it is seated properly. If it does not boot up, replace it with a new controller.

Note Do not use the power cycle as a preparation for a hot swap. You will lose access to data if the system is currently online through the surviving controller.

- Reseat the faulty controller and check the pin connections. First, power down the faulty controller. Then, remove and reinstall the controller making sure the pin connections are straight and that the cable connections are tight. See the controller hardware documentation installation instructions. Reinsert the canister into the enclosure and perform a rescan to check the hardware setup.

World Wide Name (WWN) Change

Upgrading to PowerVault 660F RAID controller firmware version 7.7 or higher can change the controller's World Wide Name (WWN) if the controller is at a version prior to 7.7. When the WWN changes, you may not be able to see the storage with Array Manager or another management application. Resolving this problem can include updating pathing of the FC HBA, as well as FC switch zoning configuration if zoning is being performed by WWN. See the latest SAN or PowerVault 660F documentation that came with your system or that is available from the Dell support site at http://support.dell.com for information on editing the FC switch zone table and resolving problems associated with a new WWN.

How to Detect and Replace a Failed Disk Drive and Restore a System to a Fault-tolerant Condition

The following sections describe how to troubleshoot problems with a failed drive on the the PowerVault 660F storage system.

Detect and Replace a Failed Disk Drive

If you receive an event notice that indicates a disk drive has failed or is not responding, there are certain steps to follow to help detect which disk drive has failed and how to do a spindown. There are several disk drives installed in the enclosure, so that you can configure them to your system requirements. You can use the Blink Disk (PV660F) command to locate a failed disk.

Restore a System to a Fault-tolerant Configuration

After the faulty disk drive has been replaced, it needs to be restored to full operation. Do this by applying the rebuild and rescan functions. Fault-tolerant is known as mirror technology in RAID configurations, commonly used as RAID 5, requiring three or more disk drives in an array. For more information, see Dell PowerVault 660F Storage System Configuration Overview

Cause of Problem

- Drive has no power
- Drive has faulty cable connections
- A mechanical failure on the device
- Disk drive offline

Other error (event) messages that may indicate a disk drive problem:

- Rebuild stopped with error: New device failed
- A new array disk failed.
- A new array disk may not be compatible with controller hardware or firmware.
- Rebuild stopped because logical drive failed
- At least one or more array disks failed in the virtual disk.
- Bad data table overflow.
- A hard disk has failed
- A hard disk failed because the user either changed the status to Offline or removed the hard disk.

Solutions

Before you remove a physical disk, you must consider the logical drives that are mapped to the disk. If the drive is *not marked dead*, and it is configured, *do not remove it*, or data will be lost. It is imperative that only the known failed disk drive is the one removed. If a second drive fails for any reason, data will be lost.

Logical Drives: A logical drive is a partition you create within an extended partition on a basic disk. A logical drive can be formatted and assigned a drive letter. Only basic disks can contain logical drives. A logical drive cannot span multiple disks. For more information on basic disks, basic volumes, and extended partitions, see the chapters on “Disks” and Volumes

If Automatic Rebuild Management (ARM) is enabled, then when the firmware detects a failed drive, it will look for a spare. If a suitable spare is found, then the failed drive is removed from all logical drives that it was part of, the spare is inserted into those logical drives, and a rebuild is started.

If ARM is not enabled, and the failed drive is removed and replaced, in the SAME SLOT, then the user may issue a manual rebuild start.

In order to use the ARM feature, go to the Advanced Controller Options settings: Right-click on a controller, click on Controller Options, and Click on the Advanced Tab to continue. When the Advanced Tab opens, make your custom changes. For more information, refer to "Advanced Fibre Channel RAID Controller Options" in the "The Dell PowerVault 660F Storage System" chapter.

Use the following procedures related to locating and removing the failed drive:

- Blink Disk (PV660F)
- Prepare to Remove
- Rebuild
- Make Global Hot Spare
- Rescan Subsystem. You may need to perform a rescan to update the disk information. Refer to "General Disk Functions" to learn more about adding a new disk, updating disk information by using rescan, how to use rescan, and other disk features.
- Write a Disk Signature. You will need to write a disk signature. Windows requires that a disk have a signature on it before it can be used.

Cannot add a hot spare

Cause of Problem

The **Assign Global Hot Spare** function will be selectable only when disks are available within the system.

Solution

Add drives to the system. See the topic *Dell PowerVault 660F and 224F Storage Systems Installation and Troubleshooting Guide* for instructions on adding new drives.

After adding the new drives, go to the Array Manager console and click to expand the **Arrays** storage object, right-click the **PV660F Subsystem** storage object, and select **Rescan.**

This action will update the drive status within the GUI. You are now ready to create a hot spare.

Cannot expand a disk group

Cause of Problem

- Disk groups containing one virtual disk are the only disk groups that can be expanded. Disk groups containing two or more virtual disks can not be expanded.
- The full capacity of the disk group must be used before additional disk groups can be created.
- Redundant controllers must be in a failover condition. To do this, send a **Disable Partner** controller command.
- There must be at least one available hot spare whose capacity is greater than or equal to the smallest array disk in the disk group.

Solution

Make sure that the situation you have supports the **Expand Capacity** command. See "Expand Capacity" in the chapter "The Dell PowerVault 660F Storage System" for more information.

After creating a virtual disk, cannot locate it in the Disks folder

Cause of Problem

A virtual disk has not been made visible to the operating system.

Solution

For a direct connect, reboot the system to recognize the virtual disk.

For a SAN connection, two steps must be performed after creating a virtual disk before the virtual disk will appear in the Disks folder.

- Assign the virtual disk to the server by using the Dell OpenManage Storage Consolidation software.
- After the disk is assigned to the server, reboot the server so that the server will recognize the new virtual disk.

Once these two steps have been performed, the virtual disk will appear in the Disks folder of the Array Manager console.

Cannot delete a volume

Cause of Problem

A volume that has been marked as a primary partition cannot be deleted. Primary partitions are protected because they contain a bootable operating system.

Solution

To override this protection feature, delete the virtual disk that the primary partition belongs to. See "Delete Virtual Disk" in the chapter "The Dell PowerVault 660F Storage System" for detailed instructions.

Cannot delete a virtual disk

Cause of Problem

Only the last virtual disk that was created can be deleted.

Solution

Delete the virtual disks in the reverse order that they were created. See the topic "Delete Virtual Disk" in the chapter "The Dell PowerVault 660F Storage System" for detailed instructions.

Firmware version mismatch

Cause of Problem

This problem may occur in the following situations:

- If you receive a firmware mismatch warning after downloading the SES firmware to enclosures attached to the PV660F storage system, then rescan the PV660F storage system to resolve this warning.
- The partner controller has different firmware. This may occur during the replacement of a failed controller.

Solution: Firmware mismatch warning after downloading SES firmware

Rescan the PV660F storage system.

Solution: The partner controller has different firmware

Both controllers must have the same version of firmware to operate in a redundant configuration. If a failed controller is replaced with a controller with a different version of firmware, the replacement controller will not be allowed to start and will be disabled by the existing controller.

Use the following procedure to download a common firmware image:

1. Check the firmware version of the existing controller. See "Fibre Channel RAID Controller Properties" in the chapter The Dell PowerVault 660F Storage System
2. Power off the subsystem. See the *Dell PowerVault 660F and 224F Storage Systems Service Manual*.
3. Remove the existing controller. See the *Dell PowerVault 660F and 224F Storage Systems Service Manual*.
4. Insert the replacement controller. See the *Dell PowerVault 660F and 224F Storage Systems Service Manual*.
5. Power on the subsystem. See the *Dell PowerVault 660F and 224F Storage Systems Service Manual*.
6. Check the firmware version of the replacement controller. See "Fibre Channel RAID Controller Properties" in the chapter The Dell PowerVault 660F Storage System
7. If the replacement controller has an older version of firmware, download the same firmware version as the that on the existing controller. This may not be the latest version of firmware available.
8. If the replacement controller has more recent firmware, power off the subsystem, exchange controllers, and download the newer firmware to the existing controller.
9. Downgrading firmware versions is not recommended. See "Fibre Channel RAID Controller Properties" in the chapter The Dell PowerVault 660F Storage System
10. Insert the second controller. See the *Dell PowerVault 660F and 224F Storage Systems Service Manual*.
11. Issue **Enable Partner** if the second controller is in the Disable Partner mode (held in reset). See "Enable Partner" in the chapter The Dell PowerVault 660F Storage System
12. Once both controllers have the same firmware version, the latest version of firmware can be downloaded to both controllers at the same time. See "Fibre Channel RAID Controller Firmware" in the chapter The Dell PowerVault 660F Storage System

Note If an empty enclosure is available, the firmware can be downloaded to the replacement controller without having to power off the subsystem. However, this will not work if the replacement controller has more recent firmware.

Unable to manage objects below the controller object

Cause of Problem

The **Enclosure Management** Advanced controller option is disabled.

Solution

Change the **Enclosure Management** Advanced controller option to enabled. See the "Advanced Fibre Channel RAID Controller Options" section of the "The Dell PowerVault 660F Storage System" chapter for details.

If the **Enclosure Management** Advanced controller option is already enabled, perform a **Rescan** on the PV660F Subsystem storage object. If this isn't successful, perform a **Reset** on the controller object. See "Reset" in the "The Dell PowerVault 660F Storage System" chapter for details.

No objects appear under Physical Array and Logical Array

It is possible that the Fibre Channel array might not discover some of the PV660F Fibre Channel arrays if any of the controllers are not fully rebooted.

To resolve this problem, right-click **PV660F SubSystem** and select **Rescan**.

Unable to reset the PowerVault 660F storage system

Cause of Problem

Either the controller has activity occurring, or the queue limit is set to zero.

Solution

Verify that no I/O is occurring to the PowerVault 660F array. You can verify this by viewing the activity LEDs on the controllers.

Verify that no controller activity is occurring. Virtual disk initializations, Expand Capacity, or any other controller activity will prohibit resets from occurring.

Verify that the Queue Limit in the controller options is not set to zero. If set to zero, reset it to a higher value (such as 32).

Application Transparent Failover autofails all LUNs on the PV650F

Starting the Array Manager service or doing a rescan when the Qlogic HBA to switch cable failure occurs on one of the servers in a SAN fabric may cause the Application Transparent Failover to autofail all LUNs on the PV650F that are assigned to other servers.

To resolve this problem, disable the PV660F management on the server with HBA failure and enable PV660F management on another server in the SAN until the cable failure is resolved.

Glossary

Use this glossary to find the definitions of terms pertinent to Array Manager.

Active partition

The partition where the computer starts up. The active partition must be a primary partition on a basic disk. If you use one Windows operating system exclusively, the active partition can be the same as the system partition. If you have a dual-boot system with two operating systems (such as Windows 2000 and Windows Server 2003) the active partition must contain the startup files for both operating systems. See also Basic disk, Extended partition, Extended partition, Partition, Primary partition.

Array

The term array refers to a RAID controller and its associated logical and physical disks.

Array disk

A physical disk controlled by a RAID controller. These disks can be placed in an array group, and virtual disks can be created from them.

Array group

An array group provides a mechanism to group disks that are controlled by a particular RAID controller. You can create virtual disks from disks in an array group if you have a SCSI RAID controller. However, with the PowerVault™ 660F controller, you will need to further define the array group into one or more disk groups before creating virtual disks. See also Disk groups.

Background initialization

Background initialization enables a virtual disk to maintain processes such as read and write operations while it is being initialized. Depending on the controller, background initialization is similar to an initialize, format, or check consistency operation. See "Background Initialization on PERC Subsystem 1 Controllers" and "Understanding Virtual Disk Initialization (PV660F)" for more information.

Basic disk

A physical disk that can contain primary partitions, an extended partition, and logical drives. See also Basic volume, Extended partition, Logical drive, and Primary partition.

Basic volume

A logical volume or partition created on a basic disk. You can create basic volumes only on basic disks. See also Basic disk and Volume.

BBU

Abbreviation for battery backup unit.

Breaking a mirror

Breaking a mirror creates two simple volumes with individual drive letters. Each volume contains the data on the mirror at the time the mirror was broken. The data is no longer redundant, but it remains intact. Compare with Removing a mirror.

Boot volume

The volume, formatted for either an NTFS or FAT file system, that contains the Windows operating system and its support files. The boot volume can be the same as the system volume. See also FAT and FAT32, NTFS, System volume, and Volume.

Cache

A fast storage area that keeps a copy of data or instructions for quicker data retrieval. Storage cache usually resides on RAID controllers and boosts performance because the CPU doesn't have to wait for a disk head to spin. Data can be written and read directly from cache.

C-C nexus

Controller-controller nexus. A controller-controller nexus refers to the state in which both redundant controllers are in communication with each other. In this state, each controller can copy write-back data to its partner controller and can determine whether the other controller is operating.

CERC

Cost Effective RAID Controller.

CERC ATA

Cost Effective RAID Controller Advanced Technology Attachment. A RAID controller that uses ATA rather than SCSI technology.

CERC SATA

Cost Effective RAID Controller Serial Advanced Technology Attachment. A RAID controller that uses Serial ATA rather than SCSI technology.

Chunk

See "Strip."

Clustering

Clustering is the process of associating multiple servers so that they function as a single, virtual server. Clustering can provide greater overall system availability, performance, and capacity for the server platforms.

Concatenation

Storing data either on one array disk (simple) or on disk space that spans more than one array disk (spanned). When spanning or concatenating more than one array disk, the entire storage capacity of each disk is combined in a manner that enables the operating system to view the concatenated array disks as a single disk. The storage space on concatenated array disks can be fully consumed by a virtual disk.

Consistency check

A consistency check verifies the accuracy of redundant data in a RAID group. For example, on a virtual disk with parity, checking consistency computes the parity of the data drives and compares the results with the contents of the parity drives.

Dell Utility Partition

This is a factory-installed, bootable partition on the hard-disk drive that provides utilities and diagnostics for your Dell hardware and software. When activated, the partition boots and provides an executable environment for the partition's utilities.

DIMM

Abbreviation for dual in-line memory module, a method for packaging memory.

Direct Attach

A Fibre Channel hardware configuration in which the host bus adapter(s) of a single server are directly connected to Fibre Channel RAID controller(s). This attachment method uses point-to-point Fibre Channel connectivity. In this configuration, the Fibre Channel RAID controller storage is confined to that server and is not shared across multiple servers. See also Redundant controllers.

Disk

A physical data storage device attached to a computer. See also Basic disk, Dynamic disk.

Disk groups

With redundant Fibre Disk RAID controllers that support a large number of physical disks, it is necessary to define one or more disk groups from the array group before creating virtual disks. You then create virtual disks from a disk group rather than from the larger array group. See also Array group.

Disk striping

See "Striping."

Disk mirroring

See "Mirroring."

Drive mapping

In Array Manager, drive mapping refers to the corresponding display of disks under the Array Group and the Disks objects in the tree view. For example, the naming convention for virtual disks under the Array Group object indicates the disk number of the virtual disk as displayed under the Disks object. Likewise, the naming convention for virtual disks under the Disks object indicates the virtual disk name as displayed under the Array Group object.

Dynamic disk

A logical disk that is managed by Array Manager software. Dynamic disks can contain only dynamic volumes (that is, volumes created with Array Manager software). Dynamic disks cannot contain partitions or logical drives, nor can they be accessed by MS-DOS or Windows 95®/98. See also Dynamic volume, Partition.

There are particular considerations regarding dynamic disks and volumes on NetWare, Windows Server 2003, and Linux. See "Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux" for more information.

Dynamic volume

A logical volume that is created with Array Manager software. Dynamic volumes include simple, spanned, striped, mirrored, and RAID-5. You must create dynamic volumes on dynamic disks. See also Dynamic disk, Volume.

Encapsulation

The process of upgrading partitions to dynamic volumes by upgrading the disk(s) they are on to dynamic. See the section "Upgrade Disks with Legacy Volumes to Dynamic on Windows 2000" in the "Volumes" chapter for details.

Extended partition

A portion of a basic disk that can contain logical drives. Use an extended partition if you want to have more than four volumes on your basic disk. Only one of the four partitions allowed per physical disk can be an extended partition, and no primary partition needs to be present to create an extended partition. Extended partitions can be created only on basic disks. See also Basic disk, Logical drive, Partition, Primary partition, Active partition.

Failback

Failback is a feature of RAID in which a RAID controller operating in failover mode detects the presence of a second operating controller, exits failover mode, and resumes its normal mode of operation. For more information, see "Failover and Failback" in the "The Dell PowerVault 660F Storage System" chapter.

Failover

Failover is a feature of RAID in which a pair of RAID controllers that are in communication with one another simultaneously process the same input/output commands to an array of disks. If one of the controllers fails, the surviving controller detects the failure and automatically assumes the function of the failed controller. For more information, see "Failover and Failback" in the "The Dell PowerVault 660F Storage System" chapter.

FAT and FAT32

FAT and FAT32 are file systems that are defined as follows:

FAT (File Allocation Table) — A file system used by MS-DOS, Windows 3.x, Windows 95, and Windows 98. Windows NT and Windows 2000 also can use the FAT file system. The operating system maintains a table to keep track of the status of various segments of disk space used for file storage.

FAT32 (File Allocation Table) — A derivative of the file allocation table (FAT) file system. FAT32 supports smaller cluster sizes than FAT, thus providing more efficient space allocation on FAT32 drives. FAT32 is designed for larger disks than FAT.

Fault tolerance

Ensures data integrity when hardware failures occur.

Fibre Channel

For Fibre Channel terms, see "Fibre Channel Concepts" in the "The Dell PowerVault 660F Storage System" chapter.

Firmware

Software (programs or data) that has been written onto read-only memory (ROM). Firmware can boot and operate a device. Each controller contains firmware which helps provide the controller's functionality.

Free space

Available space used to create logical drives within an extended partition. See also Extended partition, Logical drive, Unallocated space.

FT disk

Refers to a disk that is completely supported by Microsoft's Disk Administrator.

FT volume types

These are legacy systems that Array Manager can identify. These are basic volumes. You can use Disk Administrator in parallel with Array Manager if you need to keep these disks as basic disks.

HBA

Host bus adapter. A PCI adapter card which resides in the server whose only function is to convert data commands from PCI-bus format to storage interconnect format (examples: SCSI, Fiber Channel) and communicate directly with disk drives, tape drives, CD-ROMS, or other storage devices.

Hot plug

Adding or removing a component from the server while the server is running. Some of the PCI controllers enable this functionality. In order to "hot plug" a component, both the component and the server must support this functionality.

Hot spare

A hot spare is an extra, unused disk that is part of a disk subsystem. It is assigned as a backup disk that can take over when a primary disk fails. Hot spares remain in standby mode, ready for activation in case of a disk failure. Hot spares can replace failed drives without interrupting the system or requiring user intervention.

Hotswap

A hotswap is the replacement of a component (such as a hard drive, power supply, PCI card, or bus) while the computer is still running.

Logical drive

A logical drive is a partition you create within an extended partition on a basic disk. A logical drive can be formatted and assigned a drive letter. Only basic disks can contain logical drives. A logical drive cannot span multiple disks. See also Basic disk, Basic volume, Extended partition.

LS modules

LS modules are cards installed in slots at the front of the enclosure. Each LS module has an SES processor that monitors environmental functions and a Loop Redundancy Circuit (LRC) function, which maintains the viability of the Fibre Channel loop.

LUN

Logical unit number. A LUN is a logical representation of physical storage. You determine whether a LUN is a disk drive, a number of disk drives, or a partition on a disk drive, depending on the RAID configuration.

Mirrored volume (RAID-1)

A fault-tolerant volume that duplicates data on two physical disks. It provides data redundancy by using a copy (mirror) of the volume to duplicate the information contained in the volume. The mirror is located on a different disk. If one of the physical disks fails, the data on the failed disk becomes unavailable, but the system continues to operate using the unaffected disk.

A mirrored volume is slower than a RAID-5 volume in read operations but faster in write operations. You can create mirrored volumes only on dynamic disks. You cannot extend mirrored volumes. In Windows NT 4.0 Disk Administrator, a mirrored volume is known as a mirror set. See also Dynamic disk, Dynamic volume, Fault tolerance, RAID, Volume.

Mirroring

Duplicating data on two array disks. Mirroring provides data redundancy by using a copy (mirror) of the RAID group to duplicate the information contained in the RAID group. The mirror is located on a different array disk. If one of the array disks fails, the system can continue to operate using the unaffected disk. Both drives contain the same data at all times. Either drive can act as the operational drive.

Nexus

See C-C nexus.

NLM

Abbreviation for NetWare Loadable Module.

NTFS

An advanced file system designed for use specifically within the Windows NT operating system. It supports file system recovery, extremely large storage media, long file names, and various features for the POSIX subsystem. It also supports object-oriented applications by treating all files as objects with user-defined and system-defined attributes. NTFS is also called Windows NT file system. See also FAT and FAT32.

Online Volume Expansion (OLVE)

The capability of extending an existing simple or spanned dynamic volume while the volume is online. This is done with the **Extend Volume** command. For more information, see "Extending a Dynamic Simple or Spanned Volume" in the "Volumes" chapter.

Parity

Redundant information that is associated with a block of information. In Array Manager software, parity is a calculated value used to reconstruct data after a failure.

RAID-5 volumes stripe data and parity intermittently across a set of disks. Within each stripe, the data on one disk is parity data and the data on the other disks is normal data. RAID-5 volumes therefore require at least three disks to allow for this extra parity information. When a disk fails, Array Manager software uses the parity information in those stripes in conjunction with the data on the good disks to recreate the data on the failed disk. See also Fault tolerance, Volume, Striped volume (RAID-0), RAID 5 volume (stripe set with parity), RAID policy.

Partition

A portion of a physical disk that functions as though it were a physically separate disk. Partitions can be created only on basic disks. Partitions cannot span disks; they must be a contiguous region. See also Basic disk, Extended partition, Extended partition, Partition, Active partition.

PCI

Peripheral Component Interconnect. PCI is the predominant 32-bit or 64-bit bus used in PCs and servers.

PERC

PowerEdge Expandable RAID Controller. A RAID controller based on I20 technology.

PERC /DC

A PERC dual-channel controller.

PERC /Di

A PERC dual-integrated channel controller.

PERC /QC

A PERC quad-channel controller.

PERC /SC

A PERC single-channel controller.

PERC /Si

A PERC single-integrated channel controller.

Physical array

A physical array is a collection of physical disk drives governed by RAID management software. A physical array appears to Array Manager as one or more logical drives.

Primary partition

A volume you create using unallocated space on a basic disk. Microsoft Windows NT and other operating systems can start from a primary partition. You can create up to four primary partitions on a basic disk, or three primary partitions and an extended partition. Primary partitions can be created only on basic disks and cannot be subpartitioned. See also Basic disk, Dynamic volume, Extended partition, Partition.

Providers

Providers run on a managed system along with the server on a Windows system. Providers are plug-ins to the server framework. Each provider is specific to the hardware controller or software component that the provider manages.

RAID

Hardware RAID (Redundant Array of Independent Disks) refers to:

An array of multiple independent hard disk drives that, in varying levels, provide high performance and fault tolerance.

The term that describes a method used to standardize and categorize fault-tolerant disk systems.

RAID levels provide various mixes of performance, reliability, and cost. Array Manager supports the available RAID levels of the particular hardware included in your system.

Also, Array Manager supports the following RAID levels: Level 0 (striping), Level 1 (mirroring), Level 5 (striping with parity), Levels 10 and 0+1 (striping with mirroring). See also Fault tolerance, Striped volume (RAID-0), Mirrored volume (RAID-1), RAID 5 volume (stripe set with parity), RAID policy.

RAID-0 (striping)

See Striped volume (RAID-0).

RAID 0+1

See RAID 10 and RAID 0+1 (striping over mirrored sets).

RAID 1 (mirroring)

See Mirrored volume (RAID-1).

RAID 10 and RAID 0+1 (striping over mirrored sets)

These two RAID levels involve striped data that is mirrored. They provide both speed (from striping) and fault tolerance (from the mirroring). RAID-10 is available with PERC controllers; RAID 0+1 is a feature of the PowerVault 660F controller. See also Fault tolerance, Parity, RAID.

RAID 5 volume (stripe set with parity)

A RAID 5 volume is a fault-tolerant volume with data and parity striped intermittently across three or more physical disks. Parity is a calculated value that is used to reconstruct data after a failure. If a portion of a physical disk fails, you can recreate the data that was on the failed portion from the remaining data and parity. You can create RAID 5 volumes only on dynamic disks. You cannot mirror or extend RAID 5 volumes. In Windows NT 4.0 Disk Administrator, a RAID 5 volume is known as a striped set with parity. See also Fault tolerance, Parity, RAID.

RAID 50

RAID 50 is a concatenation of RAID 5 across more than one three-drive spans. For example, a RAID 5 array that is implemented with three drives and then continues on with three more array drives would be a RAID 50 array. RAID 50 provides data redundancy through parity. See also Parity, RAID, RAID 5 volume (stripe set with parity).

RAID Level 1-Concatenated

RAID Level 1-Concatenated is a RAID 1 array that spans across more than a single pair of array disks. This combines the advantages of concatenation with the redundancy of RAID 1. No striping is involved in this RAID type. See also Concatenation, RAID 1 (mirroring), and Striping.

RAID policy

An automated RAID configuration method that can be used with the Dell PowerVault™ 660F Storage System.

Read-ahead

A controller using read-ahead policy reads sequential sectors of the disk when seeking data.

Rebuild

Recreating data that existed on a drive before the drive failed. You can rebuild data only on virtual disks with data redundancy.

Redundant controllers

Refers to two identical Fibre Channel RAID controllers that share a common set of array disks. The two controllers can be connected to a server in a direct attach or in a SAN configuration. The two controllers communicate with each other to verify that both are functioning properly through a ping/acknowledgment sequence. Failure to acknowledge the ping triggers failover. When one controller fails, the surviving controller takes over the resources of the failed controller. This failover process is transparent to the applications running on the host. See also Direct Attach.

Region

Contiguous area of storage on a disk.

Removing a mirror

Removing a mirror from a volume destroys the data from the selected mirror and leaves the other mirror intact. After you remove a mirror, the space on the disk used by the removed mirrored volume becomes unallocated space. The remaining volume becomes a simple volume on the disk. Compare with Breaking a mirror.

SAF-TE

SCSI Accessed Fault-Tolerant Enclosure. See also Fault tolerance.

SAN

Abbreviation for Storage Area Network.

SCSI

Small Computer System Interface. A processor-independent standard for system-level communications between a computer and intelligent devices including hard-drives, floppy disks, CD drives, printers, scanners, and so on.

SES

Abbreviation for SCSI-3 Enclosure Services. The SES commands are used to obtain the status of enclosure components.

Simple volume

A volume that consists of contiguous space on a single disk. You can extend a simple volume within the same disk or onto additional disks. If you extend a simple volume across multiple disks or across noncontiguous areas on the same disk, it becomes a spanned volume. You can create simple volumes only on dynamic disks. Simple volumes are not fault tolerant but can be mirrored.

When a basic disk with a partition is upgraded, the partition becomes a simple volume. An extended partition on a basic disk also becomes a simple volume when the disk is upgraded to dynamic. See also Dynamic disk, Dynamic volume, Fault tolerance, Spanned volume.

S.M.A.R.T. (Self Monitoring Analysis and Reporting Technology)

This is a method of determining the reliability of an array disk by performing a predictive failure analysis on each disk, as defined by that disk's vendor. Array Manager picks up alerts sent from each RAID controller, if the hard disk supports S.M.A.R.T.

SNMP

Simple NetWork Management Protocol. SNMP, a popular network control and monitoring protocol, is part of the original TCP/IP protocol suite. SNMP provides the format in which vital information about different network devices, such as network servers or routers, can be sent to a management application.

Spanned volume

A volume made up of disk space on more than one physical disk. A spanned volume can also consist of two or more noncontiguous areas on a single disk. You can add more space to a spanned volume by extending it at any time. You can create spanned volumes only on dynamic disks. Spanned volumes are not fault tolerant. In Windows NT 4.0 Disk Administrator, a spanned volume is known as a volume set. See also Dynamic disk, Dynamic volume, Fault tolerance, Simple volume.

Strip

The portion of a stripe that resides on a single array disk. A strip may also be referred to as a "chunk."

Strip size

The amount of array disk space used by a single strip in a stripe. The strip size is the same for all strips in the stripe. Strip size may also be referred to as "stripe depth."

Stripe depth

Synonym for strip size. See “Strip size.”

Striped volume (RAID-0)

A volume that stores data in stripes on two or more physical disks. Data in a striped volume is allocated alternately and evenly (in stripes) to the disks of the striped volume. Striped volumes are not fault tolerant, but you can mirror a striped dynamic volume to ensure redundancy. You can also extend a striped dynamic volume. See also Dynamic disk, Dynamic volume, Fault tolerance, Simple volume, RAID.

Striping

Disk striping writes data across 3 or more array disks. Each “stripe” spans 3 or more array disks but only consumes a portion of each disk. The amount of space consumed by a stripe is the same on each array disk included in the stripe. A virtual disk may use several stripes on the same set of array disks. See also “Strip.”

Subdisk

Refers to a region of contiguous space on a disk. Subdisks are building blocks for dynamic volumes.

Subsystem

Refers to all the objects in a particular array enclosure. It includes physical and logical objects.

System volume

The volume that has the hardware-specific files needed to load Windows NT/2000. See also Volume, Partition.

TCP/IP

Abbreviation for Transmission Control Protocol/Internet Protocol. A system for transferring information over a computer network containing dissimilar systems, such as PC and UNIX.

Unallocated space

Available disk space that is not allocated to any partition, logical drive, or volume. The type of object you can create on unallocated space depends on the disk type (basic or dynamic). For basic disks, you can use unallocated space outside partitions to create primary or extended partitions. You can use free space inside an extended partition to create a logical drive. For dynamic disks, you can use unallocated space to create dynamic volumes. Unlike with basic disks, you do not select the exact disk region used to create the volume.

Virtual disk

A virtual disk is an abstract entity used by Array Manager that allows you to use different hardware RAID layouts on selected array disks. Virtual disks are hardware RAID logical drives. After the virtual disk is created and the system is rebooted, the virtual disk appears as a disk in Windows. A virtual disk is functionally identical to a physical disk from the standpoint of applications.

Volume

A logical / virtual entity that is made up of portions of one or more physical disks. A volume may be formatted and may have a file system and / or drive letter. A volume has a type (dynamic) and a layout (simple, spanned, mirrored, striped, RAID-5, and RAID-10 or RAID-0+1). See also Dynamic volume, Simple volume, Spanned volume, Striped volume (RAID-0), Mirrored volume (RAID-1), RAID 5 volume (stripe set with parity), RAID policy.

There are particular considerations regarding dynamic disks and volumes on NetWare, Windows Server 2003, and Linux. See "Dynamic Disk and Volume Support on NetWare, Windows Server 2003, and Linux" for more information.

Index

A

Access rights 13
Active partition 557
Add Member (PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, CERC ATA100/4ch) 161
Add Spares (PV660F) 248
Add Virtual Disk (PV660F) 259
Adding
  drive letter or path 201
  enclosures (PV660F) 283
  hot spares
    PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, CERC ATA100/4ch 165, 190
    PV660F 248, 272
  mirrors 224
  new disks 186
  virtual disk (PV660F) 259
Advanced Controller Options (PV660F)
  Automatic Rebuild Management 279
  Coalesce Device Queues 279
  Conservative Cache mode 278
  Enclosure Management 279
  Queue Limit 279
  Restrict Reassign to One Block 278
  S.M.A.R.T. Large Host Transfers 278
  True Verification Of Data 278
  Write Through Verify 278
Alarm Disabled (PV660F) 301
Alarm Enabled (PV660F) 301
Alarm, disable (PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, CERC ATA100/4ch) 108, 132
Alarm, enable (PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, CERC ATA100/4ch) 108, 131
Alarm, quiet (PERC 2/SC, 3/SC, 2/DC, 3/DCL, 3/DC, 3/QC, CERC ATA100/4ch) 109
Allocation unit size 178, 221
Apply RAID policy (PV660F) 255
Array and volume migration
  PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, CERC ATA100/4ch 211
Array disk commands
  PV660F 270
Array disk status descriptions 491
Array disk, defined 557
Array group commands
  PV660F 246
Array group, defined 557
Array Manager categories, changing 43
Array Manager components 2
Array Manager console
  configuring 43
  refreshing 188
Array Manager Service 29, 473
Array Manager system discovery 323
Array Manager Utilities 473
  adding and removing PERC RAID controller configuration 474
  command line options, NetWare 474
  connecting to local and remote servers 476
  fibre channel RAID and enclosure management 475
  version checking 478
Arrays, defined 557
Asset Settings (PV660F) 301
At Events Log, Description Missing at Event Details Message Box 522
Auto Restore (PV660F) 242, 281
Automatic Rebuild Management (PV660F) 279

B

Background Initialize
cancel virtual disk (PV660F) 267
Backup and restore events 336
Basic disks
defined 175, 558
functions 176
limitations 177
upgrading to dynamic 180
Basic volumes
defined 199, 558
overview 199
repairing 503
Battery
calibration 311
PERC 2, 3/Di 103
PV660F 289
recondition 312
BBU (PV660F) 310
defined 558
enable 288
Blink (PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, CERC ATA100/4ch) 193
Blink Disk (PV660F) 270
Blink Virtual Disk (PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, CERC ATA100/4ch) 172
Boot volume, defined 558
Breaking a mirror 225
defined 558
Bring a dynamic RAID-5 or mirrored volume back online 503

C

C0 Battery and C1 Battery (PV660F) 289
Cache
defined 558
Cached data 242
Calibration
battery 311
indicators 312
Cancel Background Initialize (virtual disk), PV660F 267
Cancel Check Consistency (PV660F) 268
Cancel Consistency Check 168
Cancel Initialize (virtual disk), PV660F 265
Cancel Rebuild (PV660F) 271
Cancel Recondition battery (PV660F) 290
Capacity expansion
PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, CERC ATA100/4ch 161
PV660F 257
C-C nexus 241, 242, 558
CERC
defined 558
CERC ATA
defined 558
CERC SATA
defined 559
CERC SATA/6ch 75, 85, 113
Change Low Threshold (battery), PV660F 290
Change Policy
PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, CERC ATA100/4ch 164
PV660F 263
Changing the Array Manager category display 43
Check Consistency
PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, CERC ATA100/4ch 167
PV660F 267
Chunk, see "strip" 559
Clearing Event Log entries 337
Clustering
defined 559
Coalesce Device Queues (PV660F) 279
Coalesce, Smart Large Host Transfers (PV660F) 278
Color settings, graphical layout 44
Computer Management console, Windows 2000 34
Computer Management, Windows 2000 34
Concatenation 50, 559
Configuration, view 77
Configuring
Array Manager console 43
dynamic volumes 218
virtual disks
PV660F 244
Conservative Cache mode (PV660F) 238, 242, 278
Consistency check
defined 559

Controller Options (PV660F) 276
Advanced 277
Expert 280
Fibre 282
Revert 277
Controller redundancy 241, 567
Controller-Controller Nexus, see C-C nexus
Controllers
supported 2
Create Disk Group (PV660F) 247
Creating
disk groups (PV660F) 247
dynamic volumes 67, 69, 218
logical drives 179
partitions 68, 177
virtual disks 66
PV660F 252
Customize headers 44

D

Dead disk segments, restore (PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA/6ch) 171
Degraded array disk status 492
Delete Disk Group (PV660F) 261
Delete virtual disk
PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, CERC ATA100/4ch 163
PV660F 263
Dell Utility Partition, defined 214, 559
Diagnostics status for array disks 492
DIMM
defined 559
Direct attach 243, 559
Disable Alarm (PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, CERC ATA100/4ch) 108, 132
Disable BBU (PV660F) 288
Disable Busy On Failback (PV660F) 280
Disable Check Condition For Invalid LUN (PV660F) 281
Disable Partner (PV660F RAID controller) 286
Disk Group Commands (PV660F) 251
Disk groups (PV660F) 247, 560
Disk management 175
Disk signature 188
Disk status changes 100
Disk striping, see "striping" 560
Disk View tab 41
Disks
adding 186
basic 175
creating partitions 68
defined 560
dynamic 176
foreign 516
reactivating 183, 500
status descriptions 489
updating information 187
DM View tab 42
Download (controller and SES firmware), PV660F 284
Download Firmware (PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di) 133
Drive letter or path 201
Drive status changes 100
Drivers and firmware 499
Dynamic disks 181
advantages 181
defined 176, 560
limitations 182
moving 184
reverting to basic 185
Dynamic volumes 216
configuring 218
creating 67, 69, 218
defined 200, 560
extending 221
mirrored 224
mount at an empty folder 203
RAID-5 226
reactivating 503
repairing 504
types 216
volume size 218

E

Edit RAID policy (PV660F) 249
Enable Alarm (PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, CERC ATA100/4ch) 108, 131
Enable BBU (PV660F) 288
Enable Partner (PV660F) 242, 287
Encapsulation 561
Enclosure (PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di)

commands 130
Properties 139
Enclosure (PV660F)
commands 300
Properties 302
Enclosure Management (PV660F) 279, 296
Enclosure, add (PV660F) 283
Event 522
Event Log 334
backing up and restoring 336
clearing entries 337
Event Details dialog box 335
exporting events 337
setting the size 337
viewing entries 334
Event Monitoring 333
Expand Capacity (PV660F) 257
Expert Controller Options (PV660F)
Auto Restore 281
Disable Busy On Failback 280
Disable Check Condition For Invalid LUN 281
Multi-Port Reset 281
No Pause On Controller Not Ready 281
on Queue Full Give Busy 280
Rebuild Rate 281
Reset Propagation 281
Vendor Unique Test Unit Ready Status 281
Export Log File (PERC 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4e/DC, 4/Di, 4e/Si, 4e/Di, CERC ATA100/4ch, PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA1.5/6ch, and CERC SATA1.5/2s) 110
Exporting events 337
Extend a dynamic volume 221
Extended partitions 177, 179

F

Failback 241, 242, 561
Failed array disk status 492
Failed Redundancy (At Risk) message 496
Failed Redundancy status for a dynamic volume 495
Failed Redundancy, troubleshooting 227
Failed status for a dynamic volume 496
Failover 241, 242, 561, 567
Fan Properties (PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di) 142
Fan Properties (PV660F) 303
Fast Charge (battery), PV660F 290
FAT, FAT32 178, 561
Fault tolerance 49, 561
Favorites 39, 322, 330
Fibre Channel
arbitrated loop 243
drive channel 243
fabric topology 243
host channel 243
point-to-point topology 243
Fibre Controller Options (PV660F)
Frame Control 282
Hard Loop IDs 283
Node Name Retention 282
PCI Latency Control 282
Firmware and drivers
defined 562
version 479
Firmware download (controller and SES firmware), PV660F 284
Firmware, download (PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di) 133
FlexRAID virtual sizing (PERC 2/SC, 2/DC) 98
Force Online (virtual disk), PV660F 268
Force Unit Access (PV660F) 278
Foreign disk status 491
Foreign disks 183, 516
Format
dynamic volumes 221
partitions 178
virtual disk (PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA/6ch) 170
Frame Control (PV660F) 282
Free space 562
FT disk 562
FT volume types 562

G

General disk functions 186
General volume functions 201
Global Hotspare, see Hot spare
Graphical layout, color settings 44

H

Hard Loop IDs (PV660F) 283
HBA

defined 562
Headers, customize 44
Healthy (At Risk) message 494
Healthy message 494
History 38, 330
Hot spare 272
defined 48, 562
PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, CERC ATA100/4ch 165, 190
PV660F 248, 272
Hotswap 49
defined 563

I

I/O errors 490, 494
Initialize
array disk (PERC 2, 2/Si, 3/Si, 3/Di) 190
array disk (PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA/6ch) 109
cancel virtual disk (PV660F)l 265
virtual disk
PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, CERC ATA100/4ch 169
PV660F 264
Initializing status
array disks 493
dynamic volume 494
Installation 5
changing installed components 24
components to install 6
considerations for a SAN 13
initial install 14
preliminary steps 5
remote 22
requirements 8
unattended 22
uninstalling NetWare 27
Windows 2000 506
Integrated
mirror 97

L

Link 243
Logical Array Properties (PV660F) 296
Logical drive
creating 179
defined 200, 563
LS modules, defined 563
LUN
defined 563

M

Make Global Hotspare (PV660F) 272
Managing storage remotely 315, 323
Mark primary partition as active 213
Merging a Foreign disk 183
Microsoft Management Console (MMC) snap-in 3, 481
Mirrored dynamic volumes 218, 224
adding 224
breaking 225
reactivating 225
removing 225
repairing 226
troubleshooting 512
Mirroring
defined 53, 563
integrated 97
Missing disk 490, 513, 515
Mount a volume at an empty folder 203
Moving dynamic disks 184
Multi-Port Reset (PV660F) 281
My Network Places 38

N

NetWare command line options in Array Manager Utilities 474
Nexus, see C-C nexus
NLM
defined 564
No Media array disk status 492
No Pause On Controller Not Ready (PV660F) 281
Node Name Retention (PV660F) 282
NTFS 178, 564

O

Offline
array disk status 492
disk status 490
On Q Full Give Busy (PV660F) 280
Online
array disk status 492
disk status 490
Online (Errors) disk status 490
Online command for array disks (PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, CERC ATA100/4ch) 192

Online Volume Expansion (OLVE) 564

P

Parity 48, 564
Partitions
creating 68, 177
defined 564
extended 179, 200
primary 200, 213
Paths, drive 201, 204
PCI
defined 564
PCI Latency Control (PV660F) 282
PERC
defined 564
PERC /DC
defined 565
PERC /Di
defined 565
PERC /QC
defined 565
PERC /SC
defined 565
PERC /Si
defined 565
PERC 2 75, 85, 113
PERC 2, 2/Si, 3/Si, 3/Di controllers
issues 105
maximum virtual disks 154
RAID levels supported 102
PERC 2, 3/Di controllers
battery 103
PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/Di, CERC ATA100/4ch controllers
RAID configuration 151
RAID levels supported 94
PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di controllers
limitations 151
PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, CERC ATA100/4ch controllers
issues 101
PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, CERC ATA1004/ch controllers
maximum virtual disks 154
PERC 2/Si 75, 85, 113
PERC 3/Di 75, 85, 113
PERC 3/Si 75, 85, 113
PERC 4/IM 97
PERC subsystem 1 commands
Properties 82
Save 80
PERC subsystem 2 75, 85, 113
PERC subsystem 2 controllers 75, 85, 113
Physical array 565
Physical Array Properties (PV660F) 295
Power Supply Properties (PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di) 143
Power Supply Properties (PV660F) 304
Prepare to Remove command for array disks (PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di) 192
Primary partitions 177, 200, 213, 565
Properties
array disks
PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, CERC ATA100/4ch 195
PV660F 274
array group
PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, CERC ATA100/4ch 115
PV660F 251
battery (PV660F) 291
channel
PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, CERC ATA100/4ch 118
controller
PV660F 292
enclosure (PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di) 139
enclosure (PV660F 302
fan (PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di) 142
fan (PV660F) 303
logical array (PV660F) 296
PERC subsystem 1 82
physical array (PV660F) 295
power supply (PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di) 143
power supply (PV660F) 304

subsystem (PV660F) 295
temperature (PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di) 144
temperature (PV660F) 306
virtual disks
PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, CERC ATA100/4ch 173
PV660F 268
Providers 565
PV660F
configuration overview 244
Conservative Cache mode 238
description and functions 235
Properties 292
RAID levels supported 237
PV660F and PV224F storage systems troubleshooting 524
controller status incidents 525
drive status incidents 536
enclosure status incidents 531
event monitoring and logging 524
other general problems 545
procedure for replacing a drive 536

Q

Query Max Size button 69, 220
Queue Limit (PV660F) 279
Quick Access window 65
Quiet Alarm (PERC 2/SC, 3/SC, 2/DC, 3/DCL, 3/DC, 3/QC, CERC ATA100/4ch) 109

R

RAID
defined 565
RAID level 0 (striped) 52
RAID level 0+1 237, 566
RAID level 1 (mirrored) 53, 563
RAID level 10 (mirrored striped) 60, 566
RAID level 1-concatenated (concatenated mirror) 60
RAID level 50 (concatenated distributed parity) 56
RAID levels 3 and 5 (striped with parity) 55, 566
RAID Policies 307
custom 309
PV660F 249, 255
standard 308
RAID policies (PV660F) 307
RAID-5 dynamic volumes 218, 226
problems when creating 511
reactivating 226
repairing 227
troubleshooting 511
Reactivating
disks 183, 500
dynamic volume 503
mirrored volumes 225
RAID-5 volumes 226
Ready status for array disks 491
Rebuild an array disk
PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, CERC ATA100/4ch 190
PV660F 271
Rebuild rate
PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/Di, CERC ATA100/4ch 100
PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, CERC ATA100/4ch 109
PV660F 281
Rebuild status for array disks 492
Rebuild, cancel (PV660F) 271
Rebuilding, defined 49
Recondition battery 312
Recondition battery, PV660F 289
Recondition indicators 312
Reconfiguring a virtual disk
PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, CERC ATA100/4ch 159
Recovering status for array disks 492
Redundant controllers
configuration 241
definition 567
requirements 241
Refresh 188
Regenerating status for dynamic RAID-5 volume 495
Region, definition 567
Remote connection 315
caching effect 321
disconnecting 329
discovery 323
NetWare computer 327, 519

specific steps 323
troubleshooting 514, 518
understanding 315
Remote connection sequence
NetWare computer 316
Windows computer 318
Windows computer with Disk Management 320
Remote installation 22
Remote storage management 315, 323
Removable-media devices 180
Remove Orphan / Dead Disk Segments for array disks (PERC 2, 2 / Si, 3 / Si, 3 / Di, CERC SATA / 6ch) 194
Removed status for array disks 492
Removing a mirror 225
Repairing
basic volumes 503
mirrored volumes 226
RAID-5 volumes 227
Rescan channel
PERC 2 / SC, 2 / DC, 3 / SC, 3 / DCL, 3 / DC, 3 / QC, 4 / SC, 4 / DC, 4 / Di, CERC ATA100 / 4ch 118
Rescan computer's storage objects 187, 499
Rescan controller
PERC 2 / SC, 2 / DC, 3 / SC, 3 / DCL, 3 / DC, 3 / QC, 4 / SC, 4 / DC, 4 / Di, 4 / IM, CERC ATA100 / 4ch 107
Rescan Subsystem (PV660F) 295
Reset (PV660F) 287
Reset Configuration
PERC 2 / SC, 2 / DC, 3 / SC, 3 / DCL, 3 / DC, 3 / QC, 4 / SC, 4 / DC, 4 / Di, CERC ATA100 / 4ch 114
PV660F 248
Reset Propagation (PV660F) 281
Resolving Windows 2000 Upgrade Problems 506
Restore Dead Disk Segments (PERC 2, 2 / Si, 3 / Si, 3 / Di, CERC SATA / 6ch) 171
Restoring event entries 336
Restrict Reassign to One Block (PV660F) 278
Resynching status
array disks 492
dynamic volume 495
Revert controller options (PV660F) 277

S

S.M.A.R.T alerts 101
S.M.A.R.T. (Self Monitoring Analysis and Reporting Technology) 568
S.M.A.R.T. Large Host Transfers (PV660F) 278
SAF-TE, definition 567
Save command for PERC subsystem 1 80
Set Tag Data (PERC 2 / SC, 2 / DC, 3 / SC, 3 / DCL, 3 / DC, 3 / QC, 4 / SC, 4 / DC, 4 / Di, CERC ATA100 / 4ch) 132
Set Thresholds for Temperature (PERC 2 / SC, 2 / DC, 3 / SC, 3 / DCL, 3 / DC, 3 / QC, 4 / SC, 4 / DC, 4 / Di) 135
Set Thresholds for Temperature (PV660F) 305
Show / Hide Array Disks
PERC 2 / SC, 2 / DC, 3 / SC, 3 / DCL, 3 / DC, 3 / QC, 4 / SC, 4 / DC, 4 / Di, CERC ATA100 / 4ch 115
Silent installation 22
Simple volumes 200, 217, 568
SNMP support, configuring 472
SNMP traps 335
Source 338
Spanned volume 217, 568
Split Mirror (PERC 2, 2 / Si, 3 / Si, 3 / Di, CERC SATA / 6ch) 171
Strip size, defined 568
Strip, defined 568
Stripe depth, see "Strip size" 569
Striped volume 217, 569
Striping, defined 569
Subdisk 569
Subsystem 569
Subsystem commands (PV660F) 294
Subsystem, rescan (PV660F) 295
Supported controllers 2
System partition, defined 214
System volume 569

T

Temperature Probes (PERC 2 / SC, 2 / DC, 3 / SC, 3 / DCL, 3 / DC, 3 / QC, 4 / SC, 4 / DC, 4 / Di) 129
Temperature Probes (PV660F) 299
Temperature Properties (PERC 2 / SC, 2 / DC, 3 / SC, 3 / DCL, 3 / DC, 3 / QC, 4 / SC, 4 / DC, 4 / Di) 144

Temperature Properties (PV660F) 306
Temperature settings for enclosure (PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di) 135
Temperature settings for enclosure (PV660F) 305
Transform Type (PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, CERC ATA100/4ch) 162
Troubleshooting 489
array disk status descriptions 491
cannot connect to remote computer 514, 518
common problems 506
disk status descriptions 489
Failed Redundancy 227
I/O errors 490
procedures 498
PV660F and 224F Storage Systems 524
PV660F and PV224F Storage Systems 524
reactivating missing disks 183
reactivating RAID-5 volumes 226
viewing Event Log entries 334
volume status descriptions 494
True Verification Of Data (PV660F) 278

U

Unallocated space 569
Unassign Global Hotspare (PV660F) 273
Unattended (Silent) Uninstall 28
Unattended installation 22
Unblink (PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, CERC ATA100/4ch) 194
Unblink Virtual Disk (PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, CERC ATA100/4ch) 172
Understanding Hot spares
PERC 3/QC, 3/DC, 3/DCL, 4/Di 88
Uninstalling Array Manager on Windows 27
Uninstalling NetWare 27
Unknown status for array disks 491
Unmirror (PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA/6ch) 172
Unmount a volume 204
Unreadable disk status 490
Unrecognized disk status 490
Unsigned disk type 188
Updating disk information 187, 499
Upgrading a basic disk to dynamic 71, 180

V

Vendor Unique Test Unit Ready Status (PV660F) 281
Viewing
configuration settings 77
drive paths 204
events 334
SNMP traps 335
Virtual disks
Change Policy
PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, CERC ATA100/4ch 164
PV660F 263
Check Consistency
PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, CERC ATA100/4ch 167
PV660F 267
command listing
PV660F 262
configuring
PV660F 244
Create Disk Group (PV660F) 247
creating 66
PERC 4/IM 571
PV660F 252
definition 570
deleting
PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, CERC ATA100/4ch 163
PV660F 263
Force Online (PV660F) 268
Format (PERC 2, 2/Si, 3/Si, 3/Di, CERC SATA/6ch) 170
initializing
PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, CERC ATA100/4ch 169
PV660F 264
maximum number
PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, CERC ATA100/4ch 153
migrating RAID levels (PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC,

4/SC, 4/DC, 4/Di, CERC ATA100/4ch) 162
Properties
PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, CERC ATA100/4ch 173
PV660F 268
reconfiguring
PERC 2/SC, 2/DC, 3/SC, 3/DCL, 3/DC, 3/QC, 4/SC, 4/DC, 4/Di, CERC ATA100/4ch 159
troubleshooting 511
Virtual sizing (PERC 2/SC, 2/DC) 98
Volume status descriptions 494
Volumes
defined 199, 570
dynamic 216
general functions 201
properties 205
simple 200

W

Windows 2000 Professional machine 220
Write Cache Coherency 242
Write disk signature 188
Write Through Verify (PV660F) 278