

**Dell EMC Server Administrator Storage Management
9.0.1
User's Guide**

Notes, cautions, and warnings

-  **NOTE:** A NOTE indicates important information that helps you make better use of your product.
-  **CAUTION:** A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.
-  **WARNING:** A WARNING indicates a potential for property damage, personal injury, or death.

Copyright

Copyright © 2017 Dell Inc. or its subsidiaries. All rights reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

2017 - 06

Rev. A00

Contents

1 Overview.....	12
What Is New In This Release.....	12
Before Installing Storage Management.....	12
Version Requirements For Controller Firmware And Drivers.....	12
Supported Controllers.....	13
Supported Enclosures.....	13
Support For Disk And Volume Management.....	14
2 Getting Started.....	15
Launching Storage Management.....	15
On Systems Running Microsoft Windows.....	15
On A System Running Linux And Any Remote System.....	15
User Privileges.....	16
Using The Graphical User Interface.....	16
Storage Object.....	16
Health.....	16
Information/Configuration.....	16
Using The Storage Management Command-Line Interface.....	16
Displaying The Online Help.....	16
Common Storage Tasks.....	17
3 Understanding RAID concepts.....	18
What is RAID?.....	18
Hardware and software RAID.....	18
RAID concepts.....	18
RAID Levels.....	19
Organizing Data Storage For Availability And Performance.....	19
Choosing RAID Levels And Concatenation.....	19
Concatenation.....	20
RAID level 0 (striping).....	21
RAID level 1 (mirroring).....	21
RAID level 5 (striping with distributed parity).....	22
RAID level 6 (striping with additional distributed parity).....	22
RAID level 50 (striping over RAID 5 sets).....	23
RAID level 60 (striping over RAID 6 sets).....	24
RAID level 10 (striped-mirrors).....	25
RAID Level 1-Concatenated (Concatenated Mirror).....	26
Comparing RAID Level And Concatenation Performance.....	27
No-RAID.....	28
4 Quick Access To Storage Status And Tasks.....	29
Storage Health.....	29

Hot Spare Protection Policy.....	29
Storage Component Severity.....	30
Storage Properties And Current Activity.....	30
Alerts Or Events.....	30
Monitoring Disk Reliability On RAID Controllers.....	31
Using Alarms To Detect Failures.....	31
Using Enclosure Temperature Probes.....	31
Time Delay In Displaying Configuration Changes.....	31
5 PCI Express Solid-State Device Support.....	32
What Is PCIe SSD?.....	32
PCIe SSD Features.....	32
PCIe SSD Subsystem Properties.....	32
Physical Device Properties.....	33
Physical Device Tasks.....	35
Blinking And Unblinking A PCIe SSD.....	35
Enabling Full Initialization On A Micron PCIe SSD.....	35
Preparing To Remove A PCIe SSD.....	36
Exporting The Log.....	36
Performing A Cryptographic Erase On An NVMe PCIe SSD.....	37
PCIe SSD In Slot Card Properties.....	37
PCIe SSD In Slot Card Tasks.....	39
Exporting The Log For A PCIe SSD In Slot Card.....	39
Performing A Cryptographic Erase On An NVMe PCIe SSD In Slot Card.....	40
PCIe SSD Subsystem Health.....	41
Backplanes.....	41
Backplane Firmware Version.....	41
6 Storage Information And Global Tasks.....	42
Storage Properties.....	42
Global Tasks.....	42
Setting the Remaining Rated Write Endurance Threshold.....	42
Setting the Available Spare Threshold.....	43
Storage Controller Properties.....	44
Storage Components.....	45
7 Controllers.....	46
What Is A Controller?.....	46
RAID Controller Technology: SATA And SAS.....	46
SAS RAID Controllers.....	46
RAID Controller Features.....	46
Controller — Supported RAID Levels.....	47
Controller — Supported Stripe Sizes.....	47
RAID Controller Read, Write, Cache, And Disk Cache Policy.....	47
Read Policy.....	47
Write Policy.....	48

Cache Policy.....	48
Disk Cache Policy.....	48
Background Initialization On PERC Controllers.....	49
Non-RAID Controller Description.....	49
Non-RAID SCSI Controllers.....	50
Non-RAID SAS Controllers.....	50
Firmware Or Driver Versions.....	50
Firmware/Driver Properties.....	50
Controller Health.....	51
Controller Components.....	51
Controllers Properties And Tasks.....	51
Controller Tasks.....	55
Creating A Virtual Disk.....	56
Enabling The Controller Alarm.....	56
Disabling The Controller Alarm.....	56
Turning Off The Controller Alarm.....	56
Testing The Controller Alarm.....	57
Setting The Rebuild Rate.....	57
Resetting The Controller Configuration.....	57
Exporting The Controller Log File.....	58
Foreign Configuration Operations.....	59
Importing Foreign Configurations.....	61
Importing Or Recovering Foreign Configurations.....	61
Clearing Foreign Configuration.....	62
Physical Disks In Foreign Virtual Disks.....	62
Setting Background Initialization Rate.....	65
Setting The Check Consistency Rate.....	65
Setting The Reconstruct Rate.....	66
Setting The Redundant Path Configuration.....	66
Setting The Patrol Read Mode.....	68
Starting And Stopping Patrol Read.....	69
Changing The Controller Properties.....	69
Managing The Physical Disk Power.....	70
Managing The Preserved Cache.....	72
Encryption Key.....	73
Converting To Non-RAID Disks.....	75
Converting To RAID Capable Disks.....	75
Changing The Controller Mode.....	75
Auto Configure RAID0 Operation.....	76
System Configuration Lock Down Mode.....	76
Global Tasks Supported When System Configuration Lock Down Mode Is On.....	77
Controller Tasks Supported When System Configuration Lock Down Mode Is On.....	78
Controller Reports Supported When System Configuration Lock Down Mode Is On.....	78
Physical Disk Tasks Supported When System Configuration Lock Down Mode Is On.....	79
Virtual Disk Tasks Supported When System Configuration Lock Down Mode Is On.....	79

Viewing Available Reports	79
Available Reports.....	79
Viewing Patrol Read Report.....	79
Viewing Check Consistency Report.....	80
Viewing Slot Occupancy Report.....	80
Viewing Physical Disk Firmware Version Report.....	80
8 Support For PERC 9 and PERC 10 Hardware Controllers.....	83
Support For RAID Level 10 Virtual Disk Creation On PERC 9 and PERC 10 Hardware Controllers.....	83
RAID Level 10 Virtual Disk Creation With Uneven Span.....	83
Support For Advanced Format 4KB Sector Hard-Disk Drives.....	84
Hot Spare Considerations — 4KB Sector Hard-Disk Drives.....	85
Reconfiguration Considerations — 4KB Sector Hard-Disk Drives.....	85
9 Support for Marvell RAID Controllers.....	86
10 Enclosures And Backplanes.....	87
Backplanes.....	87
Flexible Backplane Zoning.....	87
Enclosures.....	87
Enclosure Physical Disks.....	88
Enclosure Fans.....	88
Enclosure Power Supplies.....	89
Enclosure Temperature Probes.....	89
Enclosure Management Modules (EMMs).....	91
Enclosure And Backplane Health.....	92
Enclosure And Backplane Properties And Tasks.....	92
Changing The Mode On 220S And 221S Enclosures.....	97
Enclosure Management.....	97
Identifying An Open Connector On The Enclosure.....	97
Enclosure Components.....	98
11 Connectors.....	99
Channel Redundancy.....	99
Creating A Channel-Redundant Virtual Disk.....	99
Creating A Physical Disk For Channel Redundant Virtual Disks On PERC Controllers.....	99
Connector Health.....	100
Controller Information.....	100
Connector Components.....	100
Connector Properties And Tasks.....	100
Logical Connector Properties And Tasks.....	101
Path Health.....	101
Clearing The Connectors Redundant Path View.....	102
Connector Components.....	102
12 Tape Drive.....	103

Tape Drive Properties.....	103
13 RAID Controller Batteries.....	104
Battery Properties And Tasks.....	104
Battery Tasks.....	105
Battery — Available Tasks.....	105
Starting A Learn Cycle.....	105
Battery Transparent Learn Cycle.....	106
Initiating The Battery Delay Learn Cycle.....	106
To Locate Delay Learn Cycle In Storage Management.....	106
14 Physical Disks Or Physical Devices.....	107
Guidelines To Replace A Physical Disk Or Physical Device.....	107
Adding A New Disk To The System.....	107
For SAS Controllers.....	107
Replacing A Physical Disk Receiving SMART Alerts.....	108
Disk Is Part Of A Redundant Virtual Disk.....	108
Disk Is Not Part Of A Redundant Virtual Disk.....	108
Other Disk Procedures.....	108
Physical Disk Or Physical Device Properties.....	108
Physical Disk Or Physical Device Tasks.....	112
Physical Disk Tasks.....	112
Blinking And Unblinking A Physical Disk	113
Removing Dead Segments.....	113
Preparing To Remove.....	113
Rebuilding Data.....	113
Canceling A Rebuild.....	114
Assigning And Unassigning Global Hot Spare.....	114
Setting The Physical Disk Online Or Offline.....	114
Performing A Clear Physical Disk And Cancel Clear	115
Enabling Revertible Hot Spare.....	115
Performing Cryptographic Erase.....	116
Convert To RAID Capable Disk.....	116
Convert To Non-RAID Disk.....	117
15 Virtual Disks.....	118
Considerations Before Creating Virtual Disks.....	118
Virtual Disk Considerations For Controllers.....	118
Virtual Disk Considerations On Systems Running Linux.....	119
Number Of Physical Disks Per Virtual Disk.....	120
Number Of Virtual Disks Per Controller.....	120
Calculation For Maximum Virtual Disk Size.....	120
Channel-Redundant Virtual Disks.....	120
Creating Virtual Disks.....	121
Reconfiguring Or Migrating Virtual Disks.....	121
Starting And Target RAID Levels For Virtual Disk Reconfiguration And Capacity Expansion.....	121

Maintaining The Integrity Of Redundant Virtual Disks.....	122
Rebuilding Redundant Information.....	123
Managing Virtual Disk Bad Block Management.....	123
Recommendations For Clearing Bad Blocks.....	124
Virtual Disk Properties And Tasks.....	124
Virtual Disk Properties.....	125
Virtual Disk Tasks.....	126
Virtual Disk — Available Tasks.....	127
Reconfiguring A Virtual Disk.....	127
Format, Initialize, Slow, And Fast Initialize.....	127
Canceling Background Initialization.....	127
Restoring Dead Segments.....	128
Deleting Data On the Virtual Disk.....	128
Performing A Check Consistency.....	128
Canceling A Check Consistency.....	128
Pausing A Check Consistency.....	128
Resuming A Check Consistency.....	128
Blinking And Unblinking A Virtual Disk.....	128
Renaming A Virtual Disk.....	129
Canceling A Rebuild.....	129
Changing The Virtual Disk Policy.....	129
Replacing A Member Disk.....	129
Clearing Virtual Disk Bad Blocks.....	129
Encrypting A Virtual Disk.....	129
Create Virtual Disk Express Wizard.....	130
Create Virtual Disk Express Wizard (Step 2).....	131
Create Virtual Disk Advanced Wizard.....	131
Create Virtual Disk Advanced Wizard (Step 2).....	133
Create Virtual Disk Advanced Wizard (Step 3).....	134
Span Edit.....	135
Virtual Disk Task: Reconfigure (Step 1 of 3).....	135
To Reconfigure A Virtual Disk: Step 1 of 3.....	136
Virtual Disk Task: Reconfigure (Step 2 of 3).....	136
To Reconfigure A Virtual Disk Expand Virtual Disk Capacity: Step 2 of 3.....	137
Virtual Disk Task: Reconfigure (Step 3 of 3).....	137
Slow And Fast Initialize.....	138
Considerations For Slow Initialize.....	138
Formatting Or Initializing A Disk.....	138
To Locate Virtual Disks Task In Storage Management.....	139
Deleting A Virtual Disk.....	139
To Delete A Virtual Disk.....	139
To Locate Delete In Storage Management.....	139
Renaming A Virtual Disk.....	140
To Rename A Virtual Disk.....	140
To Locate Rename In Storage Management.....	140

Changing The Policy Of A Virtual Disk.....	140
Changing The Read, Write, Or Disk Cache Policy Of A Virtual Disk.....	140
To Locate Change Policy In Storage Management	140
Split Mirror.....	141
Splitting A Mirror.....	141
To Locate Split Mirror In Storage Management	141
Unmirror.....	141
To Unmirror.....	141
To Locate Unmirror In Storage Management.....	141
Assigning And Unassigning Dedicated Hot Spare.....	142
Assigning A Dedicated Hot Spare.....	142
Unassigning A Dedicated Hot Spare.....	142
To Locate Assign Or Unassign Dedicated Hot Spare In Storage Management.....	142
Virtual Disk Task: Replace Member Disk (Step 1 of 2).....	142
Replacing A Member Disk: (Step 1 of 2).....	143
To Locate Replace Member Disk In Storage Management.....	143
Virtual Disk Task: Replace Member Disk (Step 2 of 2).....	143
16 Moving Physical And Virtual Disks From One System To Another.....	144
Required Conditions.....	144
SAS Controllers.....	144
SAS Controller	144
Migrating SAS Virtual Disks To Another System.....	144
17 Protecting Your Virtual Disk With A Hot Spare.....	145
Understanding Hot Spares.....	145
Setting Hot Spare Protection Policy.....	145
Dedicated Hot Spare Protection Policy.....	145
Global Hot Spare Protection Policy.....	146
Considerations For Hot Spare Protection Policy.....	146
Considerations For Enclosure Affinity.....	146
18 CacheCade Using Solid-State Drives.....	147
Managing The CacheCade.....	147
CacheCade Properties.....	148
Creating A CacheCade.....	148
Resizing The CacheCade.....	148
Renaming The CacheCade.....	149
Blinking And Unblinking The CacheCade.....	149
Deleting The CacheCade.....	149
19 Troubleshooting.....	150
Common Troubleshooting Procedures.....	150
Cables Attached Correctly.....	150
System Requirements.....	150
Drivers And Firmware.....	150

Isolate Hardware Problems.....	151
Replacing A Failed Disk.....	151
Using The Physical Disk Online Command On Select Controllers.....	152
Recovering From Removing The Wrong Physical Disk.....	152
Resolving Microsoft Windows Upgrade Problems.....	152
Virtual Disk Troubleshooting.....	152
A Rebuild Does Not Work.....	153
A Rebuild Completes With Errors.....	153
Cannot Create A Virtual Disk.....	153
A Virtual Disk Of Minimum Size Is Not Visible To Windows Disk Management.....	154
Virtual Disk Errors On Systems Running Linux.....	154
Problems Associated With Using The Same Physical Disks For Both Redundant And Non-Redundant Virtual Disks	154
Specific Problem Situations And Solutions.....	154
Physical Disk Is Offline Or Displays An Error Status.....	155
Receive A Bad Block Alert With Replacement, Sense, Or Medium Error.....	155
Alerts 2146 Through 2150 Received During A Rebuild Or While A Virtual Disk Is Degraded.....	155
Alerts 2146 Through 2150 Received While Performing I/O, Consistency Check, Format, Or Other Operation.....	155
Read And Write Operations Experience Problems.....	156
A Task Menu Option Is Not Displayed.....	156
A Corrupt Disk Or Drive Message Suggests Running Autocheck During A Reboot.....	156
Erroneous Status And Error Messages After A Windows Hibernation.....	156
Storage Management May Delay Before Updating Temperature Probe Status.....	156
Storage Management May Delay Displaying Storage Devices After Reboot.....	156
You Are Unable To Log Into A Remote System.....	156
Cannot Connect To Remote System Running Microsoft Windows Server 2003.....	156
Reconfiguring A Virtual Disk Displays Error In Mozilla Browser.....	157
Physical Disks Are Displayed Under The Connector Object Instead Of The Enclosure Object.....	157
PCIe SSD Troubleshooting.....	157
Peripheral Component Interconnect Express (PCIe) Solid-State Drive (SSD) Is Not Visible In The Operating System.....	157
PCIe SSD Is Not Visible In Disk Management In The Operating System.....	157

20 Frequently Asked Questions..... 158

Why Is A Rebuild Not Working?.....	158
How To Avoid Removing The Wrong Disk.....	158
How Can I Safely Remove Or Replace A Physical Disk.....	158
How Do I Recover From Removing The Wrong Physical Disk.....	159
How Do I Identify The Firmware Version That Is Installed.....	159
Which Controllers Do I Have?.....	159
How Do I Turn Off An Alarm.....	159
Which RAID Level Is Best For Me?.....	159

21 Supported Features..... 160

Supported Features On PERC Hardware Controllers.....	160
Controller Tasks Supported On PERC Hardware Controllers.....	161
Battery Tasks Supported On PERC Hardware Controllers.....	162

Connector Tasks Supported By PERC Hardware Controllers.....	162
Physical Disk Tasks Supported By PERC Hardware Controllers.....	163
Virtual Disk Tasks Supported By PERC Hardware Controllers.....	163
Virtual Disk Specifications For PERC Hardware Controllers.....	164
RAID Levels Supported By PERC Hardware Controllers.....	166
Read, Write, Cache And Disk Cache Policy Supported By PERC Hardware Controllers.....	167
Enclosure Support On PERC Hardware Controllers.....	167
Supported Features On The PERC S140 Controller.....	167
Controller Tasks Supported On the PERC S140 Controller.....	168
Physical Disk Tasks Supported By The PERC S140 Controller.....	168
PERC S140 support on Non-volatile memory express (NVMe) devices.....	168
Virtual Disk Specifications For PERC S140 Controller.....	168
RAID Levels Supported By The PERC S140 Controller.....	169
Supported Features On The Non-RAID Controllers.....	169
Controller Tasks Supported On The Non-RAID Controllers.....	170
Battery Tasks Supported By Non-RAID Controllers.....	170
Connector Tasks Supported On The Non-RAID Controllers.....	170
Physical Disk Tasks Supported On The Non-RAID Controllers.....	171
Virtual Disk Tasks Supported By The Non-RAID Controllers.....	171
Enclosure Support On The Non-RAID Controllers.....	172
Enclosure And Backplane Features.....	172
Enclosure And Backplane Tasks.....	172
Maximum Supported Configuration For SAS Controllers.....	173

22 Determining The Health Status For Storage Components..... 174

Health Status Rollup: Battery Is Charging Or Dead.....	174
Health Status Rollup: Physical Disks In A Virtual Disk Are Failed Or Removed.....	174
Health Status Rollup: Physical Disks In A Virtual Disk Are Unsupported, Partially, Or Permanently Degraded.....	175
Health Status Rollup: All Physical Disks In A Virtual Disk Are In Foreign State.....	175
Health Status Rollup: Some Physical Disks In A Virtual Disk Are In Foreign State.....	175
Health Status Rollup: Virtual Disk Is Degraded; Physical Disks Are Failed Or Rebuilding.....	175
Health Status Rollup: Virtual Disk Is Failed.....	176
Health Status Rollup: Unsupported Firmware Version.....	176
Health Status Rollup: Enclosure Power Supply Failed Or Power Connection Removed.....	176
Health Status Rollup: One Enclosure Fan Is Failed.....	176
Health Status Rollup: One Enclosure EMM Is Failed.....	177
Health Status Rollup: One Enclosure Temperature Probe Is Failed.....	177
Health Status Rollup: Lost Both Power Connections To The Enclosure.....	177
Health Status Rollup: One Or More Physical Disks Are Failed.....	177
Health Status Rollup: Physical Disk Is Rebuilding.....	178

Overview

Server Administrator Storage Management provides enhanced features for configuring the locally attached RAID and non-RAID disk storage on a system. Storage Management enables you to perform controller and enclosure functions for all supported RAID and non-RAID controllers and enclosures from a single graphical user interface (GUI) or command-line interface (CLI). The GUI is wizard-driven and includes features for novice and advanced users. The CLI is fully featured and scriptable. Using Storage Management, you can protect your data by configuring data-redundancy, assigning hot spares, or rebuilding failed physical disks. All users of Storage Management should be familiar with their storage environment and Storage Management.

Storage Management supports SATA and SAS, but not Fibre Channel.

For information on Storage Management alerts, see the *Server Administrator Messages Reference Guide* at dell.com/openmanagemanuals.

What Is New In This Release

This release of Storage Management provides the following new features:

- Support for the following operating systems:
 - Support for Red Hat Enterprise Linux 6.9
 - Support for ESXi 6.0 U3
 - Support for ESXi 6.5
- Support for the following browsers:
 - Internet Explorer - 9, 10, 11
 - Google Chrome - 58
 - Safari - 9.1
 - Mozilla Firefox 52, 53
- Support to set available spare threshold for NVMe PCIe SSDs — [Setting The Available Spare Threshold](#).
- ‘Setting the remaining rated write endurance threshold’ feature is supported on SATA SSDs connected to the PERC S140 Controller.
- System Configuration Lock Down mode is supported — If the system configuration mode is set, you will not be allowed to initiate any of the configuration operation from either OMSS GUI or CLI. For more information, see [System Configuration Lock Down Mode](#).
- Support for RAID configuration on non-volatile memory express (NVMe) devices — [PERC S140 support for NVMe devices](#).
- Support for Marvell RAID controllers — [Support for Marvell RAID Controllers](#).

 **NOTE:** For the list of supported operating systems and servers, go to dell.com/openmanagemanuals, browse to OpenManage Software, then access the required version of the *OpenManage Software Support Matrix* document.


Before Installing Storage Management

The following sections describe considerations for installing Storage Management.

Version Requirements For Controller Firmware And Drivers

For Storage Management to function properly, the controllers must have the minimum required version of the firmware and drivers installed. The firmware and drivers listed in the *Server Administrator Release Notes* refer to the minimum supported version for

these controllers. Later versions of the firmware and drivers are also supported. For the most recent driver and firmware requirements, contact your service provider.

 **NOTE: To download the latest storport driver, see the Microsoft Knowledge Base article KB943545 at support.microsoft.com.**

If you install Storage Management without the minimum required firmware and drivers, Storage Management may not be able to display the controllers or perform other functions. Storage Management generates alerts 2131 and 2132 when it detects unsupported firmware or drivers on a controller.

For information on alert messages, see the *Server Administrator Messages Reference Guide*.

Supported Controllers

 **NOTE: The firmware and drivers listed in the *Server Administrator Release Notes* refer to the minimum supported version for these controllers. Later versions of the firmware and drivers are also supported. For the most recent driver and firmware requirements, contact your service provider.**

This release of Storage Management supports the following controllers.

Supported RAID Controllers

Storage Management supports the following RAID controllers. For information on the technology used by the supported RAID controllers, see [RAID Controller Technology: SATA and SAS](#).


- PERC S140
- PERC H840 Adapter
- PERC H330 Adapter, PERC H330 Mini Monolithic
- PERC H730 Adapter, PERC H730 Mini Monolithic
- PERC H730P Adapter, PERC H730P Mini Monolithic
- PERC H740P Adapter, PERC H740P Mini Monolithic

 **NOTE: The order of the controllers displayed on Storage Management may differ with the order of the controllers displayed in the Human Interface (HII) and PERC Option ROM. The order of the controllers does not cause any limitation.**

Supported Non-RAID Controllers


Storage Management supports the following non-RAID controllers:

- SAS 12 Gbps HBA
- HBA 330 Mini
- HBA 330 Adapter

 **NOTE: In Storage Management, the Broadcom SAS HBA controller page display properties such as: ID, Status, Name, Slot ID, State, Driver Version, Storport Driver Version, Number of Connectors, and Controller Tasks.**

 **NOTE: For HBA 330 controller, the backplane is displayed under a single logical connector, and all the physical disks are displayed under the same backplane.**

 **NOTE: For SAS 12 Gbps HBA and HBA 330 controller, the Available Reports that can be viewed are [View Slot Occupancy Report](#) and [View Physical Disk Firmware Version Report](#).**

 **NOTE: At times, the connector ID for SAS 12 Gbps HBA may not be displayed correctly on Storage Management. This occurs because of a firmware limitation in the SAS 12 Gbps HBA. However, this connector ID discrepancy does not cause any functional limitation.**

Supported Enclosures

This release of Storage Management supports the MD1400 and MD1420 storage systems enclosures.

Support For Disk And Volume Management

Storage Management does not provide disk and volume management. To implement disk and volume management, you must use the native disk and volume management utilities provided by your operating system.

Getting Started

Server Administrator Storage Management is designed for system administrators who implement hardware RAID solutions and understand corporate and small business storage environments.

Storage Management enables you to configure the storage components attached to your system. These components include RAID and non-RAID controllers and the channels, ports, enclosures, and disks attached to them. Using Storage Management, you can configure and manage the controller functions without accessing 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 like rebuilds, troubleshooting, setting thresholds. Most functions can be configured and managed while the system remains online and continues to process requests.

Storage Management reports the status of storage components. When the status for a component changes, Storage Management updates the display for that component and sends an alert to the **Alert Log**.

In addition to status changes, Storage Management generates alerts for user actions such as creating or deleting a virtual disk and for many other events. Most alerts also generate SNMP traps.

Other than monitoring and reporting status, Storage Management does not automatically initiate actions independent of user input. Storage Management actions are user-initiated using wizards and drop-down menus. Storage Management does, however, report the actions taken by the controllers, which include generating alerts, initiating tasks, such as a rebuild, and making state changes.

 **NOTE: Storage Management reports the change in state of disks and other storage components as viewed by the controller.**

Launching Storage Management

Storage Management is installed as a Server Administrator service. All Storage Management features are accessible by selecting the **Storage** object in the Server Administrator tree view. For more information on starting Server Administrator, see the *Server Administrator User's Guide*

On Systems Running Microsoft Windows

To start a Server Administrator session on a local system running Microsoft Windows operating system, click the **Server Administrator** icon on your desktop and log in using an account with Administrator privileges.

 **NOTE: Administrative privileges are required for configuration purposes.**

On A System Running Linux And Any Remote System

To start a Server Administrator session on a Linux or any remote system, click the **Server Administrator** icon on your desktop and log in using an account with Administrator privileges.

Or, open a web browser and type one of the following in the address field and press <Enter>:

`https://<localhost>:1311`

where <localhost> is the assigned name for the managed system and 1311 is the default port.

or

`https://<IP address>:1311`

where <IP address> is the IP address for the managed system and 1311 is the default port.

 **NOTE: Type https:// (not http://) in the address field to receive a valid response in your browser.**

User Privileges

Server Administrator provides security through the User, Power User, and Administrator user groups. Each user group is assigned a different level of access to the Server Administrator features.

The Administrator privileges are required to access all Storage Management features. Administrator privilege allows you to execute the drop-down menu tasks, launch wizards, and use the **omconfig storage** command-line interface commands. Without Administrator privileges, you cannot manage and configure the storage component.

User and Power User privileges allow you to view storage status, but not manage or configure storage. With User and Power User privileges, you can use the **omreport storage** command and not the **omconfig storage** command.

For more information on user groups and other Server Administrator security features, see the *Server Administrator User's Guide*.

Using The Graphical User Interface

The following sections describe how to access the Storage Management features using the Server Administrator graphical user interface (GUI).

Storage Object

The Server Administrator tree view displays a **Storage** object. The Storage Management features are accessible by selecting the **Storage** object or expanding the **Storage** object and selecting a lower-level object.

Related links

[Displaying The Online Help](#)

Health

On the **Properties** page, click **Health** to view the status information for the storage components.

Related links

[Storage Health](#)

Information/Configuration

On the **Properties** page, click **Information/Configuration** to view the property information for a storage object. The **Information/Configuration** subtabs also have options for executing storage tasks or launching wizards.


Using The Storage Management Command-Line Interface

Storage Management has a fully featured command-line interface (CLI). For more information on CLI, see the *Server Administrator Command Line Interface User's Guide*.

Displaying The Online Help

Storage Management provides an extensive online Help. This Help is available from the Server Administrator graphical user interface when the **Storage** or lower-level tree view object is selected.

The online Help is available as:

- Context-sensitive Help — Each Storage Management page has a  icon. Click this icon to display the context-sensitive online Help that describes the contents of the displayed page.
- Table of Contents — The table of contents is available in the page that displays the information when you access the context-sensitive Help.

Related links

[Storage Object](#)

Common Storage Tasks

This section provides information on commonly performed storage tasks:

- Create and configure virtual disks (RAID configuration). For more information, see:
 - [Create Virtual Disk Express Wizard](#) — This topic provides information on using the **Express Wizard** to create a virtual disk. Using the Express Wizard is the quickest method for creating a virtual disk. The Express Wizard is appropriate for novice users.
 - [Create Virtual Disk Advanced Wizard](#) — This topic provides information on using the **Advanced Wizard** to create a virtual disk. The Advanced Wizard requires a good knowledge of RAID levels and hardware and is appropriate for advanced users.
 - [Virtual Disks](#) — This topic provides detailed information regarding the virtual disk management. This information includes controller-specific considerations that affect virtual disk creation and management.
- Assign a hot spare to the virtual disk — When a virtual disk uses a RAID level, you can assign a hot spare (backup physical disk) to rebuild data if a physical disk in the virtual disk fails.
 - [Protecting Your Virtual Disk With A Hot Spare](#) — This topic provides information on hot spares and controller-specific information.
- Perform a Check Consistency — The [Maintaining The Integrity Of Redundant Virtual Disks](#) task verifies the accuracy of the redundant data on a virtual disk.
- Reconfigure a Virtual Disk — To expand the capacity of a virtual disk you can add physical disks to the virtual disk. You can also change the RAID levels. For more information, see [Virtual Disk Task: Reconfigure \(Step 1 of 3\)](#).

Understanding RAID concepts


Storage Management uses the Redundant Array of Independent Disks (RAID) technology to provide Storage Management capability. Understanding Storage Management requires an understanding of RAID concepts, as well as some familiarity with how the RAID controllers and operating system view disk space on your system.

Related links

- [What is RAID?](#)
- [Organizing Data Storage For Availability And Performance](#)
- [Choosing RAID Levels And Concatenation](#)
- [Comparing RAID Level And Concatenation Performance](#)

What is RAID?

RAID is a technology for managing the storage of data on the physical disks that reside or are attached to the system. A key aspect of RAID is the ability to span physical disks so that the combined storage capacity of multiple physical disks can be treated as a single, extended disk space. Another key aspect of RAID is the ability to maintain redundant data which can be used to restore data 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. The 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 RAB defines the RAID levels, commercial implementation of RAID levels by different vendors may vary from the actual RAID specifications. An implementation of 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 physical disks. When using software RAID provided by the operating system, the operating system implements the RAID levels. For this reason, using software RAID by itself can slow the system performance. You can, however, use software RAID along with hardware RAID volumes to provide better 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.

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 physical disk to another physical disk. Mirroring provides data redundancy by maintaining two copies of the same data on different physical 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 always. Either side of the mirror can act as the operational side. A mirrored RAID disk group is comparable in performance to a RAID 5 disk group in read operations but faster in write operations.
- **Striping** — Disk striping writes data across all physical disks in a virtual disk. Each stripe consists of consecutive virtual disk data addresses that are mapped in fixed-size units to each physical disk in the virtual disk using a sequential pattern. For example, if the virtual disk includes five physical disks, the stripe writes data to physical disks one through five without repeating any of the physical disks. The amount of space consumed by a stripe is the same on each physical disk. The portion of a stripe that resides

on a physical disk is a stripe element. Striping by itself does not provide data redundancy. Striping in combination with parity does provide data redundancy.

- **Stripe size** — The total disk space consumed by a stripe not including a parity disk. For example, consider a stripe that contains 64KB of disk space and has 16KB of data residing on each disk in the stripe. In this case, the stripe size is 64KB and the stripe element size is 16KB.
- **Stripe element** — A stripe element is the portion of a stripe that resides on a single physical disk.
- **Stripe element size** — The amount of disk space consumed by a stripe element. For example, consider a stripe that contains 64KB of disk space and has 16KB of data residing on each disk in the stripe. In this case, the stripe element size is 16KB and the stripe size is 64KB.
- **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.
- **Span** — A span is a RAID technique used to combine storage space from groups of physical disks into a RAID 10, 50, or 60 virtual disk.

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 also entail an increase or decrease in the I/O (read and write) performance of a system.

Maintaining redundant data requires the use of additional physical disks. The possibility of a disk failure increases with an increase in the number of disks. Since 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 the ability of a system 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 option cannot be avoided in situations where redundant data is a requirement, it does have implications on the workload of the system support staff within your organization.
- **Volume** — Volume refers to a single disk non-RAID virtual disk. You can create volumes using external utilities like the O-ROM <Ctrl> <r>. Storage Management does not support the creation of volumes. However, you can view volumes and use drives from these volumes for creation of new virtual disks or Online Capacity Expansion (OCE) of existing virtual disks, provided free space is available. Storage Management allows Rename and Delete operations on such volumes.

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 topics provide specific information on how each RAID level or concatenation store data as well as their performance and protection characteristics:

- [Concatenation](#)
- [RAID Level 0 \(Striping\)](#)
- [RAID Level 1 \(Mirroring\)](#)

- [RAID Level 5 \(Striping With Distributed Parity\)](#)
- [RAID Level 6 \(Striping With Additional Distributed Parity\)](#)
- [RAID Level 50 \(Striping Over RAID 5 Sets\)](#)
- [RAID Level 60 \(Striping Over RAID 6 Sets\)](#)
- [RAID Level 10 \(Striping Over Mirror Sets\)](#)
- [RAID Level 1-Concatenated \(Concatenated Mirror\)](#)
- [Comparing RAID Level And Concatenation Performance](#)
- [No-RAID](#)

Related links

[Starting And Target RAID Levels For Virtual Disk Reconfiguration And Capacity Expansion](#)

Concatenation


In Storage Management, concatenation refers to storing data on either one physical disk or on disk space that spans multiple physical disks. When spanning more than one disk, concatenation enables the operating system to view multiple physical 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 physical disk.

Data that spans more than one physical disk can be considered a spanned volume. Multiple concatenated disks can also be defined as a virtual disk that comprises more than one physical disk.

A dynamic volume that spans to separate areas of the same disk is also considered concatenated.

When a physical 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. In addition, a concatenated volume can easily be expanded by adding an additional physical disk.



- 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 redundant data is stored. When a disk fails, the large virtual disk fails.
- No performance gain.
- No redundancy.

RAID level 0 (striping)

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




RAID 0 characteristics:

- Groups n disks as one large virtual disk with a capacity of (smallest disk size) $\times n$ disks.
- Data is stored to the disks alternately.
- No redundant data is stored. When a disk fails, the large virtual disk fails with no means of rebuilding the data.
- Better read and write performance.

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 physical disks. If a physical disk fails, data can be rebuilt using the data from the other side of the mirror.




RAID 1 characteristics:

- Groups $n + n$ disks as one virtual disk with the capacity of n disks. The controllers currently supported by Storage Management 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 both the disks.
- When a disk fails, the virtual disk still works. The data is read from the mirror of the failed disk.
- 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.

RAID level 5 (striping with distributed parity)

RAID 5 provides data redundancy by using data striping in combination with parity information. Rather than dedicating a physical disk to parity, the parity information is striped across all physical disks in the disk group.




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.

RAID level 6 (striping with additional distributed parity)

RAID 6 provides data redundancy by using data striping in combination with parity information. Similar to RAID 5, the parity is distributed within each stripe. RAID 6, however, uses an additional physical disk to maintain parity, such that each stripe in the disk group maintains two disk blocks with parity information. The additional parity provides data protection in the event of two disk failures. In the following image, the two sets of parity information are identified as **P** and **Q**.




RAID 6 characteristics:

- Groups n disks as one large virtual disk with a capacity of $(n-2)$ disks.
- Redundant information (parity) is alternately stored on all disks.
- The virtual disk remains functional with up to two disk failures. The data is reconstructed from the surviving disks.
- Better read performance, but slower write performance.
- Increased redundancy for protection of data.
- Two disks per span are required for parity. RAID 6 is more expensive in terms of disk space.

RAID level 50 (striping over RAID 5 sets)

RAID 50 is striping over more than one span of physical disks. For example, a RAID 5 disk group that is implemented with three physical disks and then continues on with a disk group of three more physical disks would be a RAID 50.

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 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 as much parity information as standard RAID 5.
- Data is striped across all spans. RAID 50 is more expensive in terms of disk space.

RAID level 60 (striping over RAID 6 sets)

RAID 60 is striping over more than one span of physical disks that are configured as a RAID 6. For example, a RAID 6 disk group that is implemented with four physical disks and then continues on with a disk group of four more physical disks would be a RAID 60.




RAID 60 characteristics:

- Groups $n*s$ disks as one large virtual disk with a capacity of $s*(n-2)$ 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 6 span.
- Better read performance, but slower write performance.
- Increased redundancy provides greater data protection than a RAID 50.
- Requires proportionally as much parity information as RAID 6.
- Two disks per span are required for parity. RAID 60 is more expensive in terms of disk space.

RAID level 10 (striped-mirrors)

The RAB considers RAID level 10 to be an implementation of RAID level 1. RAID 10 combines mirrored physical disks (RAID 1) with data striping (RAID 0). With RAID 10, data is striped across multiple physical disks. The striped disk group is then mirrored onto another set of physical disks. RAID 10 can be considered a *mirror of stripes*.




RAID 10 characteristics:

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

RAID Level 1-Concatenated (Concatenated Mirror)

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

NOTE: You cannot create a RAID 1-concatenated virtual disk or reconfigure to RAID 1-concatenated with Storage Management. You can only monitor a RAID 1-concatenated virtual disk with Storage Management.



Comparing RAID Level And Concatenation 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. Evaluate your specific environment requirements before choosing a RAID level.

NOTE: The following table does not show all supported RAID levels in Storage Management. For information on all supported RAID levels in Storage Management, see [Choosing RAID Levels And Concatenation](#).

Table 1. RAID Level and Concatenation Performance Comparison

RAID Level	Data Availability	Read Performance	Write Performance	Rebuild Performance	Minimum Disks Required	Suggested Uses
Concatenation	No gain	No gain	No gain	N/A	1 or 2 depending on the controller	More cost efficient than redundant RAID levels. Use for noncritical data.
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, and critical information.
RAID 5	Good	Sequential reads: good. Transactional reads: Very good	Fair, unless using writeback cache	Fair	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	Good	Very Good	Fair	Fair	N + 2 (N = at least 4)	Medium sized transactional or

RAID Level	Data Availability	Read Performance	Write Performance	Rebuild Performance	Minimum Disks Required	Suggested Uses
						data intensive uses.
RAID 6	Excellent	Sequential reads: good. Transactional reads: Very good	Fair, unless using writeback cache	Poor	$N + 2$ ($N =$ at least two disks)	Critical information. Databases and other read intensive transactional uses.
RAID 60	Excellent	Very Good	Fair	Poor	$X \times (N + 2)$ ($N =$ at least 2)	Critical information. Medium sized transactional or data intensive uses.
N = Number of physical disks X = Number of RAID sets						

No-RAID

In Storage Management, a virtual disk of unknown metadata is considered a No-RAID volume. Storage Management does not support this type of virtual disks. These must either be deleted or the physical disk must be removed. Storage Management allows **Delete** and **Rename** operation on No-RAID volumes.

Quick Access To Storage Status And Tasks

This section describes various methods to determine the status or health of the storage components on your system and how to quickly launch the available controller tasks.

Related links

- [Storage Health](#)
- [Hot Spare Protection Policy](#)
- [Storage Component Severity](#)
- [Storage Properties And Current Activity](#)
- [Alerts Or Events](#)
- [Monitoring Disk Reliability On RAID Controllers](#)
- [Using Alarms To Detect Failures](#)
- [Using Enclosure Temperature Probes](#)
- [Time Delay In Displaying Configuration Changes](#)
- [Setting the Remaining Rated Write Endurance Threshold](#)

Storage Health

The **Storage Dashboard** displays the combined status for each controller and lower-level storage components. For example, if the health of the storage system has been compromised due to a degraded enclosure, both the enclosure **Health** and the controller severity on the **Storage Dashboard** display a yellow exclamation mark to indicate a Warning severity. If a controller on the **Storage Dashboard** displays a Warning or Critical status, perform the following actions to investigate the cause:

- Click **Check Alert Log** to display the **Alerts Log**. Review the Alert Log for alerts relating to the status of the controller and its lower-level components. The **Check Alert Log** link is only displayed when the controller displays a Warning or Critical status.
- Select the controller and investigate the status of the lower-level components. For more information, see [Storage Component Severity](#).
- Click the virtual disk that is in degraded state to display the **Physical Disk Properties** page.

 **NOTE: The virtual disk link is displayed only if the physical disks that are part of the virtual disk, are in a Warning or Critical state.**

For more information on how the status of lower-level components is *rolled up* into the status displayed for the controller, see [Determining The Health Status For Storage Components](#).

Related links

- [Health](#)

Hot Spare Protection Policy

The **Set Hot Spare Protection Policy** task allows you to set or modify the number of hot spares to be assigned to the virtual disks. After you set the number of assigned hot spares, any deviation from the protection policy threshold triggers an alert based on the severity level you set.

Related links




- [Setting Hot Spare Protection Policy](#)
- [Dedicated Hot Spare Protection Policy](#)
- [Global Hot Spare Protection Policy](#)

Storage Component Severity

Component status is indicated by the severity. A component with a Warning or Critical/Failure status requires immediate attention to avoid data loss, if possible. The status of the component indicates the combined status of the components and its lower-level objects.

It may be useful to review the **Alert Log** for events indicating why a component has a Warning or Critical status.

Table 2. Component Severity

Severity	Component Status
	Normal/OK — The component is working as expected.
	Warning/Non-critical — A probe or other monitoring device has detected a reading for the component that is above or below the acceptable level. The component may still be functioning, but it could fail. The component may also be functioning in an impaired state. Data loss is possible.
	Critical/Failure/Error/Fatal — The component has either failed or failure is imminent. The component requires immediate attention and may need to be replaced. Data loss may have occurred.

Related links

[Determining The Health Status For Storage Components](#)

Storage Properties And Current Activity

The **Information/Configuration** page displays information regarding the storage component. These properties include details such as the number of connectors (channels or ports) on a controller or the Enclosure Management Modules (EMM) firmware version.

The **State** and **Progress** properties display the current activity of a component. For example, an offline physical disk displays the Offline status while the Progress property displays how close to completion an operation (such as a rebuild) is.

The following sections describe the properties for each component:

- [Storage Information And Global Tasks](#)
- [Battery Properties And Tasks](#)
- [Connector Properties And Tasks](#)
- [Enclosure And Backplane Properties And Tasks](#)
- [Physical Disk Or Physical Device Properties](#)
- [Physical Disk Or Physical Device Tasks](#)
- [EMM Properties](#)
- [Fan Properties](#)
- [Power Supply Properties](#)
- [Temperature Probe Properties And Tasks](#)
- [Virtual Disk Properties And Tasks](#)

Alerts Or Events

Storage activity generates alerts or events that are displayed in the Alert Log. Some alerts indicate normal activity and are displayed for informational purposes only. Other alerts indicate abnormal activity which must be addressed immediately. For more information about alerts and their corrective actions, see the *Server Administrator Messages Reference Guide*.

Monitoring Disk Reliability On RAID Controllers

Storage Management supports Self Monitoring Analysis and Reporting Technology (SMART) on physical disks that are SMART-enabled.

SMART performs predictive failure analysis on each disk and sends alerts if a disk failure is predicted. The RAID controllers check physical disks for failure predictions and, if found, pass this information to Storage Management. Storage Management immediately displays an alert icon on the disk. Storage Management also sends an alert to the Alert Log and the Microsoft Windows application log.

 **NOTE: You may not receive SMART alerts when the I/O of a controller is paused.**

Using Alarms To Detect Failures

Certain storage components have alarms which when enabled, alert you when the component fails.

Related links

[Enabling The Enclosure Alarm](#)

[Enabling The Controller Alarm](#)

Using Enclosure Temperature Probes

Physical disk enclosures have temperature probes that warn you when the enclosure has exceeded an acceptable temperature range.

Related links

[Setting The Temperature Probe Values](#)

Time Delay In Displaying Configuration Changes

When you change the storage configuration, Storage Management quickly generates SNMP traps in response to the configuration changes. The Storage Management, Management Information Base (MIB) is also updated to reflect storage configuration changes. However, it may take up to five minutes to update the MIB with the most recent storage configuration. For this reason, there is a time delay of up to five minutes between the receipt of an SNMP trap and the ability to identify the configuration changes by querying the Storage Management MIB. This time delay is notable when creating a new virtual disk or performing an unmirror or split mirror on a RAID 1-concatenated virtual disk.

PCI Express Solid-State Device Support

This section provides an overview of the Storage Management device management support for Peripheral Component Interconnect Express (PCIe) Solid-State Drive (SSD) and its associated devices like the backplane and extender card.

In Storage Management, PCIe SSD appears under Storage in the tree view. Storage Management reports the PCIe SSD devices and its various properties.

 **NOTE: Storage Management does not support RAID management or configuration on PCIe SSD sub systems.**

What Is PCIe SSD?

Peripheral Component Interconnect Express (PCIe) solid-state device (SSD) is a high-performance storage device designed for solutions requiring low latency, high Input Output Operations per Second (IOPS), and enterprise class storage reliability and serviceability. The PCIe SSD is designed based on Single Level Cell (SLC) NAND flash technology with a high-speed PCIe 2.0 compliant interface. The high-speed PCIe 2.0 compliant interface helps improve performance for I/O bound solutions.

PCIe SSD Features

Following are the key features of PCIe SSD:


- Hot plug capability
- High-performance device
- Support for 2.5-inch HDD Form Factor





PCIe SSD Subsystem Properties

The PCIe SSD subsystem comprises of the following components:

- Backplane
- PCIe Solid-State Device

Table 3. PCIe SSD Subsystem Properties

Properties	Description
ID	Displays the subsystem ID assigned to the PCIe SSD subsystem by Storage Management. Storage Management numbers the controllers and PCIe SSD subsystems attached to the system starting with zero. This number is the same as the PCIe SSD subsystem ID number reported by the <code>omreport</code> command. For information on Command Line Interface, see the <i>Server Administrator Command Line Interface User's Guide</i> .  NOTE: In CLI commands, the PCIe SSD subsystem ID is displayed as the controller ID.
Status	These icons represent the severity or health of the PCIe SSD Subsystem.

Properties	Description
	 — Normal/OK  — Warning/Non-critical  — Critical/Failure/Error
Name	Displays the name of the subsystem.
Slot ID	Displays the slot to which the PCIe SSD subsystem is attached.  NOTE: If Not Applicable is displayed, you can identify the Slot ID by selecting the System → Main → Systems Chassis → Slots object in the tree view and clicking the Information tab. The Slot ID property on this tab may display the correct information.
State	Displays the status of the subsystem. Possible values are: <ul style="list-style-type: none"> • Ready — The subsystem is functioning normally. • Degraded — The subsystem has encountered a failure and is operating in a degraded state. • Failed — The subsystem has encountered a failure and is no longer functioning.
Available Reports	Allows you to view the Slot Occupancy report. For more information, see Available Reports .



Physical Device Properties

You can view information about PCIe SSDs and perform PCIe SSD tasks on the **Physical Device Properties** page. To view the complete PCIe SSD properties, click **Full View** on the top of the page. For more information, see [Physical Device Tasks](#).

The following table lists the physical device properties for PCIe SSD.

Table 4. Physical Device Properties

Properties	Description
Name	Displays the name of the PCIe SSD. The name comprises the bay ID and the slot in which the PCIe SSD is installed.
State	Displays the health state of the PCIe SSD.
Bus Protocol	Displays the technology that the PCIe SSD is using.
Device Protocol	Displays the device protocol of the physical device, such as Non-Volatile Memory Express (NVMe).
Media	Displays the media type of the physical disk.
Device Life Status	Displays the life status of the PCIe SSD. The device life status is determined by the following attributes: Percent Lifetime Used — This attribute is determined by the elapsed time since the start of use (up to three years) or percentage of total bytes written (TBW). Write Protect Progress — This attribute is determined by the reduction in number of available spare sectors. If the available spare sectors are less than 10 percent of the original pool, the drive enters read-only mode. Possible values for the device life status are:

Properties	Description
	<p>Drive Health Good — The drive is used within the TBW specification. The drive health is good as sufficient spare blocks are available. The drive health status is good if the values for percent lifetime used and write protect progress is less than 100 percent.</p> <p>Approaching Warranty Coverage Expiry — The drive is reaching the specified TBW, indicating that it is close to the end of warranty coverage. However, the drive will be functional as the number of spare blocks available are still above the threshold for entering the read-only mode. The drive approaches the warranty coverage expiry if the value for percent lifetime used is greater than or equal to 90 percent and that for write protect progress is less than the threshold value, which is 90 percent.</p> <p>Warranty Coverage Expired — The drive has reached the TBW threshold and met the life expectancy specification. The drive is functional as the number of spare blocks available are still above the threshold for entering the read-only mode. But, the specified data retention period (amount of time that data can be read from the drive after TBW is reached) drops, if TBW specification is exceeded and the warranty for the drive expires. The warranty coverage for the drive expires if the value for percent lifetime used is equal to 100 percent and that for write protect progress is less than 100 percent.</p> <p>Approaching Read Only — The drive is running out of spare sectors and is reaching the read-only mode. However, the health status of the drive is good and data retention is unaffected. The drive is stated to be approaching the read-only mode if the value for percent lifetime used is less than 100 percent and that for write protect progress is greater than or equal to 90 percent.</p> <p>Read Only — The drive is in read-only mode. Users must save open files, if any, to another device and replace or remove the device. If this scenario occurs within three years of the device installation, this failure is covered under warranty. The drive is in read-only mode if the value for percent lifetime used is less than 100 percent and that for write protect progress is equal to 90 percent.</p>
Driver Version	<p>Displays the version of the driver that is installed on the PCIe SSD sub-system.</p> <p> NOTE: Storage Management displays Not Applicable for subsystem for which the driver version cannot be obtained.</p>
Remaining Rated Write Endurance	<p>Displays information on the SSD renewal/replacement based on the amount of write workloads. This field indicates the total remaining programs or erase-cycles available on the SSD, based on the cumulative specification of the total NAND (Negated AND or NOT AND) Flash chips in the SSD.</p> <p> NOTE: This option is applicable to Micron PCIe SSDs, Non-Volatile Memory Express (NVMe) PCIe SSDs, and SAS/SATA SSDs.</p>
Firmware Revision	Displays the firmware version of the physical device.
Model Number	Displays the Piece Part Identification (PPID) of the PCIe SSD.
Capacity	Displays the capacity of the device.
Vendor ID	Displays the hardware vendor of the device.

Properties	Description
Product ID	Displays the product ID of the device.
Serial No.	Displays the serial number of the device.
PCIe Negotiated Link Speed	Displays the current negotiated transfer speed of the physical device in GT/s.
PCIe Maximum Link Speed	Displays the capable transfer speed of the physical device in GT/s.
PCIe Negotiated Link Width	Displays the current negotiated link width of the physical device.
PCIe Maximum Link Width	Displays the capable link width of the physical device.
Sub Vendor	Displays the vendor name of the device.
Available Spare	Displays the threshold value set for all PCIe SSDs. The default value is 10% .

Physical Device Tasks

The physical device tasks for PCIe SSD are as follows:

- [Blinking And Unblinking](#)
- [Performing A Full Initialization On A PCIe SSD](#)
- [Preparing To Remove A PCIe SSD](#)
- [Exporting The Log](#)
- [Performing A Cryptographic Erase on An NVMe PCIe SSD](#)

To perform a physical device task:

1. In the **Server Administrator** window, under the **System** tree, expand the **Storage** tree object to display the storage component objects.
2. Expand the **PCIe SSD Subsystem** object
3. Expand the **Enclosure (Backplane)** object.
4. Select the **Physical Devices** object.
5. Select a task from the **Available Tasks** drop-down menu.
6. Click **Execute**.

Blinking And Unblinking A PCIe SSD

The **Blink** task allows you to find a device within a system by blinking one of the LEDs on the device. You can use this task to locate a failed device. Select **Unblink** to cancel the **Blink** task or to stop the LED on a physical device that is blinking indefinitely.

Enabling Full Initialization On A Micron PCIe SSD

Does my controller support this feature? See [Supported Features](#).

 **CAUTION: Full Initialization permanently erases all data present on the disk.**

 **CAUTION: On VMware ESXi hosts, before performing Full Initialization on the Micron PCIe SSD it is important to first delete any data stores on it. Failure to do so can result in system instability.**

To erase an encrypted physical device, select the **Full Initialization** task. This task is available for:

- Unconfigured SED drive
- Foreign configured encrypted drives

- Unconfigured and foreign SED drive even when an Encryption Key is not present in the controller

Related Tasks

- [Performing A Full Initialization On A Micron PCIe SSD](#)

Performing A Full Initialization On A Micron PCIe SSD

Performing a Full Initialization on a Micron PCIe SSD overwrites all blocks and results in permanent loss of all data on the Micron PCIe SSD. During Full Initialization, the host is unable to access the Micron PCIe SSD.

 **NOTE: If the system reboots or experiences a power loss during Full Initialization, the operation is canceled. You must reboot the system and restart the process.**

Related Task

- [To Locate Full Initialization In Storage Management](#)


To Locate Full Initialization In Storage Management

To locate this task in storage management:


1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Expand a controller object.
3. Expand the enclosure or backplane object.
4. Select the **Physical Devices** object.
5. Select **Full Initialization** from the Tasks drop-down menu of the physical device you want to clear.
6. Click **Execute**.

Preparing To Remove A PCIe SSD

PCIe SSDs support orderly hot swap allowing you to add or remove a device without halting or rebooting the system in which the devices are installed.

 **CAUTION: The identify LED pattern (blink operation) is the same LED pattern as safe to remove. When you initiate a Prepare to Remove operation, ensure that the PCIe SSD is no longer accessible by the system before physically removing the PCIe SSD.**

 **CAUTION: To prevent data loss, it is mandatory that you use the Prepare to Remove task before physically removing a device.**

 **NOTE: Orderly hot swap is only supported when the PCIe SSDs are installed in a supported system running a supported operating system. To ensure that you have the correct configuration for your PCIe SSD, see the system-specific Owner's Manual.**

 **NOTE: The Prepare to Remove task for PCIe SSDs is supported on systems running the VMware vSphere (ESXi) 6.0 operating system. However, this task is not supported on prior versions of VMware vSphere (ESXi) operating system.**

Select the **Prepare to Remove** task to safely remove a PCIe SSD from the system. This task causes the status LEDs on the device to blink. You can safely remove the device from the system under the following conditions after you initiate the **Prepare to Remove** task:

- The PCIe SSD is blinking the **safe to remove** LED pattern.
- The PCIe SSD is no longer accessible by the system.

Exporting The Log

The export log contains debug information for the PCIe SSD and can be useful when troubleshooting. You can export the reliability log from the **Physical Device Available Tasks** drop-down list.

Related links

- [To Locate Export Log In Storage Management For NVMe PCIe SSDs](#)

Performing A Cryptographic Erase On An NVMe PCIe SSD

 **NOTE:** The virtual disk with RAID 10 configuration cannot be created for the selected number of physical disks. Please refer [OpenManage CLI guide](#) for details.


Does my controller support this feature? See [Supported Features](#).

 **NOTE:** After you hot plug an NVMe PCIe SSD, the NVMe PCIe SSD may take several seconds to be displayed on Storage Management.

 **CAUTION:** Cryptographic Erase permanently erases all data present on the disk.

Performing a Cryptographic Erase on an NVMe PCIe SSD overwrites all blocks and results in permanent loss of all data on the NVMe PCIe SSD. During Cryptographic Erase, the host is unable to access the NVMe PCIe SSD.

 **NOTE:** If the system reboots or experiences a power loss during Cryptographic Erase, the operation is canceled. You must reboot the system and restart the process.

 **NOTE:** On systems running the VMware vSphere (ESXi) operating system, Storage Management only performs the Cryptographic Erase task and does not report the status of the task. If the Cryptographic Erase task does not function as expected for some NVMe devices, the status is not reported or displayed on Server Administrator.

Related Task

- [To Locate Cryptographic Erase In Storage Management](#)

Related links

- [To Locate Cryptographic Erase In Storage Management](#)
- [Performing A Cryptographic Erase On An NVMe PCIe SSD](#)
- [Performing A Cryptographic Erase On An NVMe PCIe SSD](#)
- [To Locate Cryptographic Erase In Storage Management](#)

To Locate Cryptographic Erase In Storage Management

 **NOTE:** Cryptographic Erase is not supported for NVMe devices connected to PERC S140 controller.

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Expand a controller object.
3. Expand the enclosure or backplane object.
4. Select the **Physical Devices** object.
5. Select **Cryptographic Erase** from the **Tasks** drop-down menu of the physical device you want to clear.
6. Click **Execute**.

Related links







- [To Locate Cryptographic Erase In Storage Management](#)
- [Performing A Cryptographic Erase On An NVMe PCIe SSD](#)
- [Performing A Cryptographic Erase On An NVMe PCIe SSD](#)
- [To Locate Cryptographic Erase In Storage Management](#)

PCIe SSD In Slot Card Properties

The PCIe SSD in slot card is similar to a physical disk in Storage Management. This PCIe SSD card is directly attached to the PCIe SSD subsystem, unlike a PCIe extender card which is attached to an enclosure or backplane. You can view information about the PCIe SSDs and run PCIe SSD tasks on the PCIe SSD in slot card **Properties** page. To view the complete PCIe SSD in slot card properties, click **Full View** on the top of the page. For more information, see [PCIe SSD in Slot Card Tasks](#)

The following table lists the PCIe SSD in slot card properties.

Table 5. PCIe SSD in Slot Card Properties

Properties	Description
ID	Displays the ID assigned to the PCIe SSD in slot card by Storage Management.
Status	<p>These icons represent the severity or health of the PCIe SSD in slot card.</p> <ul style="list-style-type: none">  — Normal/OK  — Warning/Non-critical  — Critical/Failure/Error  — Unknown
Name	Displays the name of the PCIe SSD in slot card as PCIe SSD in Slot <X> , where <X> represents the front chassis slot number.
State	<p>Displays the status of the PCIe SSD in slot card. Possible values are:</p> <p>Ready — The PCIe SSD in slot card is functioning normally.</p> <p>Degraded — The PCIe SSD in slot card has encountered a failure and is operating in a degraded state.</p> <p>Failed — The PCIe SSD in slot card has encountered a failure and is no longer functioning.</p>
Device Name	Displays the name of the PCIe SSD in slot card. The name comprises the bay ID and the slot in which the PCIe SSD in slot card is installed.
Bus Protocol	Displays the technology that the PCIe SSD is using.
Device Protocol	Displays the device protocol of the PCIe SSD in slot card, such as Non-Volatile Memory Express (NVMe).
Media	Displays the media type of the device.
Driver Version	<p>Displays the version of the driver that is installed on the PCIe SSD subsystem.</p> <p> NOTE: Storage Management displays Not Applicable for subsystem for which the driver version cannot be obtained.</p>
Remaining Rated Write Endurance	<p>Displays information on the SSD renewal/replacement based on the amount of write workloads. This field indicates the total remaining programs or erase-cycles available on the SSD, based on the cumulative specification of the total NAND (Negated AND or NOT AND) Flash chips in the SSD.</p> <p> NOTE: This option is applicable to Micron PCIe SSDs, Non-Volatile Memory Express (NVMe) PCIe SSDs, and SAS/SATA SSDs.</p>
Firmware Revision	Displays the firmware version of the device.
Model Number	Displays the Piece Part Identification (PPID) of the PCIe SSD.
Capacity	Displays the capacity of the device.
Vendor ID	Displays the hardware vendor of the device.
Product ID	Displays the product ID of the device.

Properties	Description
Serial No.	Displays the serial number of the device.
PCIe Negotiated Link Speed	Displays the current negotiated transfer speed of the device in GT/s.
PCIe Maximum Link Speed	Displays the capable transfer speed of the device in GT/s.
PCIe Negotiated Link Width	Displays the current negotiated link width of the device.
PCIe Maximum Link Width	Displays the capable link width of the device.
Form Factor	Displays the form factor of the device. Possible values are: Add-in Card — For HHHL devices 2.5 inches — For physical devices other than a HHHL device.
Sub Vendor	Displays the vendor name of the device.
Available Spare	Displays the threshold value set for all PCIe SSDs. The default value is 10% .

PCIe SSD In Slot Card Tasks

You can perform the following tasks on the PCIe SSD in slot card:

- [Exporting The Log For PCIe SSD In Slot Card](#)
- [Performing A Cryptographic Erase On A PCIe SSD In Slot Card](#)

To perform a PCIe SSD in slot card task:

1. In the **Server Administrator** window, under the **System** tree, expand the **Storage** tree object to display the storage component objects.
2. Expand the **PCIe SSD Subsystem** object.
3. Select the **PCIe SSD in Slot** object.
4. Select a task from the **Available Tasks** drop-down menu.
5. Click **Execute**.

Exporting The Log For A PCIe SSD In Slot Card

The export log contains debug information for the PCIe SSD and can be useful when troubleshooting. You can export the reliability log for the PCIe SSD in Slot Card from the **Available Tasks** drop-down list.

Related links

[To Locate Export Log In Storage Management For NVMe PCIe SSDs](#)

To Locate Export Log In Storage Management For NVMe PCIe SSDs


To locate this task in storage management:

 **NOTE: This operation is supported only on NVMe PCIe SSDs.**

1. Expand the **Storage** tree object to view the controller objects.
2. Expand **PCIe SSD Subsystems**.
3. Select **Physical Disks**.
4. Select **Export Log** from the **Available Tasks** drop-down menu.
5. Click **Execute**.
The **Export Log** page is displayed.
6. In the **Export Log** page, the following attributes are displayed:

- **Host Name**

- **Path**
- **File Name** — In this text box, you can provide a custom file name for the log file. The export log file is saved with a **.log** file extension and the file extension cannot be overwritten by any other file extension entered by the user. The default file name is **NVME_<device name>_<monthdayhourminutesecond>.log**.

 **NOTE: You cannot use MS-DOS reserved words and special characters in the custom file name. For more information on MS-DOS reserved words, see <https://support.microsoft.com/en-us/kb/71843>. For more information on valid and invalid special characters, see <https://support.microsoft.com/en-us/kb/177506>.**

7. Click **Export Log** to export the file.

Related links

[What Is PCIe SSD?](#)

Performing A Cryptographic Erase On An NVMe PCIe SSD In Slot Card

Does my controller support this feature? See [Supported Features](#).

 **NOTE: After you hot plug an NVMe PCIe SSD, the NVMe PCIe SSD may take several seconds to be displayed on Storage Management.**

 **CAUTION: Cryptographic Erase permanently erases all data present on the disk.**

Performing a Cryptographic Erase on an NVMe PCIe SSD overwrites all blocks and results in permanent loss of all data on the NVMe PCIe SSD. During Cryptographic Erase, the host is unable to access the NVMe PCIe SSD.

The following error messages will be displayed based on the actions performed while performing Cryptographic erase


 **NOTE: The error messages will not be displayed in CLI mode.**

- If drive is mountedWhen you perform the 'Cryptographic Erase' operation, any data on the disk will be lost permanently. This device is currently mounted(/mnt/temp) and may still be in use. We recommend un-assigning them before continuing. Are you sure you want to continue?

 **NOTE: The mount path/drive letters is displayed in () brackets, in this case it is (/mnt/temp).**

- If any Raw IO in progressWhen you perform the 'Cryptographic Erase' operation, any data on the disk will be lost permanently.This device may still be in use. We recommend that you stop all applications accessing the device before continuing. Are you sure you want to continue?

 **NOTE: If the system reboots or experiences a power loss during a Cryptographic Erase, the operation is canceled. You must reboot the system and restart the operation.**

 **NOTE: On systems running the VMware vSphere (ESXi) operating system, Storage Management performs only the Cryptographic Erase task and does not report the status of the task. If the Cryptographic Erase task does not function as expected for some NVMe devices, the status is not reported or displayed on Server Administrator.**

Related Task

- [To Locate Cryptographic Erase In Storage Management For A PCIe SSD In Slot Card](#)

To Locate Cryptographic Erase In Storage Management For A PCIe SSD In Slot Card

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Expand a controller object.
3. Select the **PCIe SSD in Slot** object.
4. Select **Cryptographic Erase** from the **Tasks** drop-down menu of the physical device you want to clear.
5. Click **Execute**.

PCIe SSD Subsystem Health

Indicates the roll-up health status of physical devices. The individual health status of the physical devices appears at the respective level.

Related links

[Backplanes](#)

[Backplane Firmware Version](#)

Backplanes

PCIe SSDs are attached to the PCIe SSD backplane of the system. The number of supported PCIe SSDs depend on the system.

 **NOTE: PCIe SSDs must be used with PCIe SSD backplanes. Do not plug in SAS/SATA devices to a PCIe SSD backplane or vice versa.**

Related links

[PCIe SSD Subsystem Health](#)

Backplane Firmware Version

The backplane firmware version information is available in the **Information/Configuration** page of the PCIe SSD sub system.

 **NOTE: The firmware version is the only backplane property supported for PCIe SSD.**

Related links

[PCIe SSD Subsystem Health](#)

Storage Information And Global Tasks

Use the **Storage** information and **Global Tasks** window to view high-level information about the storage components or devices on your system. These windows also allow you to launch global tasks that affect all controllers attached to the system.

Related links

[Storage Properties](#)




[Global Tasks](#)

[Storage Controller Properties](#)

Storage Properties

The Storage tree-view object has the following properties.

Table 6. Storage Properties

Property	Definition
Status	<p>These icons represent the severity or health of the storage component. For more information, see Storage Component Severity.</p> <p> — Normal/OK</p> <p> — Warning/Non-critical</p> <p> — Critical/Failure/Error</p>
Smart Thermal Shutdown	Displays whether thermal shutdown is enabled or disabled.

Global Tasks

To perform a global task, select a task from the **Global Tasks** drop-down menu and click **Execute**.

The available tasks in the Global Tasks drop-down box are:

- [Setting Hot Spare Protection Policy](#)
- [Setting The Remaining Rated Write Endurance Threshold](#)
- [Setting The Available Spare Threshold](#)

Setting the Remaining Rated Write Endurance Threshold

The **Set Remaining Rated Write Endurance Threshold** is supported on PCIe and SAS/SATA SSDs. The **Set Remaining Rated Write Endurance Threshold** option provides added feature functionality in managing the threshold level for only PCIe SSDs, or SAS/SATA SSDs, or both based on your system configuration.

To set the Remaining Rated Write Endurance Threshold:

1. In the **Server Administrator** window, under the **System** tree, select the **Storage** tree object.
The storage **Properties** window is displayed.
2. Click the **Information/Configuration** subtab to display more information.
3. Under **Global Tasks**, select **Set Remaining Rated Write Endurance Threshold** from the drop-down menu.


 **NOTE:** You can also click the **Set Remaining Rated Writer Endurance Threshold** link available on the **Storage Dashboard** to access this operation.

4. Click **Execute**.

The **Set Threshold for Remaining Rated Write Endurance** window is displayed with the following options:

- **Set Threshold for PCIe SSD (1–100)** — Displays the default threshold value for all PCIe SSDs
- **Set Threshold for SAS/SATA SSD (1–100)** — Displays the default threshold value for all SAS/SATA SSDs

5. Type the threshold values for one or both the available options and click **Apply Changes**.

 **NOTE:** If you click **Apply Changes** without entering a new threshold value for either of the available options an error message is displayed. Enter a new threshold value and then click **Apply Changes** to successfully complete the task.

6. After you click **Apply Changes**, Storage Management updates the system with the new threshold values.

After the task is completed, you will be redirected to the **Information/Configuration** window.

 **NOTE:** The newly entered threshold values are retained even after you upgrade **Server Administrator**.


 **NOTE:** When the Remaining Rated Write Endurance Threshold drops below the configured threshold, an Alert is logged and the SNMP trap is received for the same. For more information on SNMP traps, see the *Dell EMC OpenManage SNMP Reference Guide* at dell.com/openmanagemanuals.

Table 7. Remaining Rated Write Endurance Threshold Polling Interval

Device Type	Threshold Polling Interval
NVMe PCIe SSDs	Five minutes
Non-NVMe PCIe SSDs	Twenty-four hours
SAS/SATA SSDs	Seven days

Setting the Available Spare Threshold

The **Set Available Spare Threshold** is supported on PCIe SSDs. The **Set Available Spare Threshold** option provides added feature functionality to configure the threshold limit for the Available Spare NVMe PCIe SSDs.

 **NOTE:** This operation is not supported on SSDs connected to a PERC and HBA.

 **NOTE:** It is supported on PCIe SSDs connected to both PERC S140 controller and PCIe SSD sub systems.

To set the Available Spare Threshold:

1. In the **Server Administrator** window, under the **System** tree, select the **Storage** tree object.
The storage **Properties** window is displayed.
2. Click the **Health** subtab to display more information.
3. Click the **Set Available Spare Threshold** link available on the Storage Dashboard to access this operation.

 **NOTE:** The default threshold is set at 10%.

 **NOTE:** You can select this option from the **Information/Configuration** tab > **Global Tasks** also.

The **Set Available Spare Threshold** window is displayed.

4. In **Set Threshold for PCIe SSD (1-99)**, enter the appropriate threshold value.

 **NOTE:** If available spare threshold value is set an alert message is generated.


 **NOTE:** The threshold value cannot be set to 100%.

5. Click **Apply Changes**.

 **NOTE:** If you click **Apply Changes** without entering a new threshold value for either of the available options an error message is displayed. Enter a new threshold value and then click **Apply Changes** to successfully complete the task.

6. After you click **Apply Changes**, Storage Management updates the system with the new threshold values. After the task is completed, you will be redirected to the **Information/Configuration** window.

 **NOTE:** The newly entered threshold values are retained even after you upgrade Server Administrator.


 **NOTE:** When the Available Spare Threshold drops below the configured threshold, an Alert is logged and based on the alert type selection the SNMP trap is received for the same. For more information on SNMP traps, see the *Dell EMC OpenManage SNMP Reference Guide* at dell.com/openmanagemanuals. The threshold polling interval is scheduled to repeat every five minutes.

Storage Controller Properties

The information displayed for each controller can vary depending on the controller characteristics.

 **NOTE:** The firmware and drivers listed in the *Server Administrator Release Notes*, see the minimum supported version for these controllers. Later versions of the firmware and drivers are also supported. For the most recent driver and firmware requirements, contact your service provider.

Table 8. Controller Properties

Property	Definition
Status	Displays the status of the controller.
ID	Displays the controller ID as reported by the <code>omreport</code> CLI command.
Name	Displays the name of the controller. For more detailed information on a controller, click its name on the controller name.
Slot ID	<p>Displays the slot to which the controller is attached. Storage Management displays Slot Not Applicable for controllers for which it is unable to display the slot ID and Embedded for embedded controllers.</p> <p> NOTE: If Slot Not Available is displayed, you can identify the slot ID by selecting the System → Main → System Chassis → Slots object in the tree view and selecting the Information tab. The Slot ID property on this tab may display the correct information.</p>
State	<p>Displays the state of the controller. Possible values are:</p> <ul style="list-style-type: none"> • 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.
Firmware Version	Displays the version of the firmware available on the controller.
Minimum Required Firmware Version	Displays the minimum firmware version that is required by Storage Management. This property is displayed only if the controller firmware does not meet the minimum requirement.
Driver Version	Displays the version of the drivers installed on the controller.
Minimum Required Driver Version	Displays the minimum driver version that is required by Storage Management. This property is displayed only if the controller driver does not meet the minimum requirement.
Number of Connector	Displays the number of connectors available on the controller. Each connector can be attached to physical disks or an

Property	Definition
	enclosure. Depending on the controller type, the connector can be either a SCSI channel or a SAS port.
Rebuild Rate	The rebuild rate is the percentage of the resources available on the system dedicated to rebuild a failed disk when a rebuild is necessary. For more information on rebuild rate, see Setting The Rebuild Rate .
Alarm State	Displays whether the alarm on the controller is enabled or disabled.
Cluster Mode	Indicates whether the controller is part of a cluster configuration.

Storage Components

For information on attached controllers, see [Controllers](#).

Controllers

This chapter provides information about the supported controllers and controller features in Storage Management.

Related links

[RAID Controller Technology: SATA And SAS](#)

[Which Controllers Do I Have?](#)

[Non-RAID Controller Description](#)

What Is A Controller?

Most operating systems do not read and write data directly from the disks, but instead send read and write instructions to a controller. The controller is the hardware in your system that interacts directly with the disks to write and retrieve data. A controller has connectors (channels or ports) which are attached to one or more physical disks or an enclosure containing physical disks. RAID controllers can span the boundaries of the disks to create an extended amount of storage space — or a virtual disk — using the capacity of more than one disk.

Controllers also perform other tasks, such as initiating rebuilds, initializing disks, and more. To complete their tasks, controllers require special software known as firmware and drivers. In order to function properly, the controller must have the minimum required version of the firmware and drivers installed.

Storage Management supports different types of controllers. If your system has a supported controller, the controller is displayed by expanding the **Storage** object in the System tree view in the GUI. You can select the controller to display tabs for executing controller tasks and viewing controller properties.

Different controllers have different characteristics in the way they read and write data and execute tasks. It is helpful to understand these features to most efficiently manage the storage. The following sections describe the supported controllers and their features.

RAID Controller Technology: SATA And SAS

Storage Management supports RAID controllers using SATA and SAS technology. This section indicates which technology the supported RAID controllers use. For more information on these controllers, see [Supported Features](#) and the controller hardware documentation.

SAS RAID Controllers

The following RAID controllers use Serial Attached SCSI (SAS) technology.

- PERC S140 controllers
- PERC H330 Adapter, PERC H330 Mini Monolithic
- PERC H730 Adapter, PERC H730 Mini Monolithic
- PERC H730P Adapter, PERC H730P Mini Monolithic
- PERC H740P Adapter, PERC H740P Mini Monolithic
- PERC H840 Adapter

RAID Controller Features

Different controllers have different features. If you have more than one controller attached to your system, you may notice that the tasks displayed on the **Information/Configuration** page of the controller are different for each controller.

Controllers may also have differences in their read, write, and cache policies as well as how they handle hot spares. You should be aware of these differences when creating virtual disks and assigning hot spares.

The following describes some of the RAID controller features and provides links to a more detailed explanation. For information on the features supported by the controllers, see [Supported Features](#).

- Hot spares — On RAID controllers, a hot spare is a backup for a disk that fails. See [Protecting Your Virtual Disk With A Hot Spare](#).
- Rebuilding data — You can rebuild data from a failed physical disk if the disk is a member of a redundant virtual disk. See [Rebuilding Redundant Information](#).
- Virtual disk expansion — Virtual disk expansion enables you to expand the capacity of a virtual disk while it remains online by adding additional disks to the virtual disk. This feature is also known as online capacity expansion (OLCE). See [Virtual Disk Tasks](#).
- RAID migration — After creating a virtual disk, you can change the RAID level. See [Reconfiguring Or Migrating Virtual Disks](#).
- Moving physical and virtual disks to another controller — This feature enables you to move the physical and virtual disks from one system to another. See [Moving Physical And Virtual Disks From One System To Another](#).
- Read, write, and cache policies — The manner in which a controller reads and writes data can vary. The read, write, and cache policies have implications for data encryption and system performance. See [RAID Controller Read, Write, Cache, And Disk Cache Policy](#).
- Check consistency — A check consistency determines the integrity of the redundant data on a virtual disk. When necessary, this feature rebuilds the redundant information. See [Maintaining The Integrity Of Redundant Virtual Disks](#).
- Patrol Read — Patrol read identifies disk errors to avoid disk failures, data loss, or corruption. For more information, see [Setting The Patrol Read Mode](#).
- Disk migration or foreign configurations — Some controllers enable you to move physical disks that contain one or more virtual disks to another controller. The receiving controller is able to recognize and import the foreign configuration (virtual disks). For more information, see [Foreign Configuration Operations](#).

Controller — Supported RAID Levels

RAID controllers may support different RAID levels. For information on supported RAID levels for a controller, see [Supported Features](#).

Controller — Supported Stripe Sizes

When creating a virtual disk, you must specify the stripe size for the virtual disk. Different controllers have different limitations on the stripe sizes they can support. For information on the supported stripe sizes for a controller, see the virtual disk specifications section for the controller in [Supported Features](#).

RAID Controller Read, Write, Cache, And Disk Cache Policy

When creating a virtual disk, you can specify the read, write, and cache policies for the virtual disk. The following subsection describes these policies.

Related links

[Changing The Virtual Disk Policy](#)

Read Policy

Does my controller support this feature? See [Supported Features](#).

The read policies indicate whether the controller should read sequential sectors of the virtual disk when seeking data.


- **Read Ahead** — The controller reads sequential sectors of the virtual disk when seeking data. **Read ahead** policy may improve system performance if the data is written to the sequential sectors of the virtual disk.
- **No Read Ahead** — Selecting no read ahead policy indicates that the controller should not use read ahead policy.
- **Read Cache Enabled** — The controller reads the cache information to verify 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 be retrieved more quickly from the cache than from the disk.
- **Read Cache Disabled** — The controller retrieves data directly from the disk and not from the cache.

Write Policy

Does my controller support this feature? See [Supported Features](#).

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

- **Write Back** — 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 retrieve data quickly from the cache then from the disk. However, data loss may occur in the event of a system failure which prevents that data from being written on a disk. Other applications may also experience problems when actions assume that the data is available on the disk.

 **NOTE: Storage Management does not allow you to select the Write Back policy for controllers that do not have a battery. This restriction protects a controller without a battery from the loss of data that may occur in the event of a power failure. On some controllers, the Write Back policy may be available in the controller BIOS even though it is not available in Storage Management.**

- **Force Write Back** — The write cache is enabled regardless of whether the controller has a battery. If the controller does not have a battery and force write back caching is used, data loss may occur in the event of a power failure.


 **NOTE: For PERC 10 Controllers 'Force Write Back' policy is same as the 'Write Back' policy.**

- **Write Back Enabled** — The controller firmware disables the write cache if it does not detect the presence of a charged battery over a specified period. For example, on some controllers, the write cache is disabled if the firmware does not detect a charged battery within 72 hours.

- **Write Through** — 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 is the default write policy setting when cluster mode is enabled.**

- **Write Cache Enabled Protected** — The controller writes data to the write cache before writing data to the physical disk. Because it takes less time to write data to the write cache than it does to a disk, enabling write cache can improve system performance. After 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 physical disk. The **Write Cache Enabled Protected** policy is available only 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 physical disk even when there is power failure.

 **NOTE: Storage Management does not allow you to select the Write Cache Enabled Protected policy for controllers that do not have a battery. This restriction protects a controller without a battery from the data loss that may occur in the event of a power failure. When using the Create Virtual Disk Advanced Wizard on a controller without a battery, the wizard either displays Write Cache Disabled as the only available option or the wizard does not display any option for write policy.**

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

Cache Policy

Does my controller support this feature? See [Supported Features](#).

The Direct I/O and Cache I/O policy apply to readings on a specific virtual disk. These settings do not affect the read-ahead policy. The cache policies are:

- **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 the default setting.

 **NOTE: Cache policy is not supported on controllers that do not have a battery.**

Disk Cache Policy

Does my controller support this feature? See [Supported Features](#).

 **NOTE: The Disk Cache Policy feature is not supported on solid-state disks (SSDs) and on SAS 12 Gbps Adapter.**

The Disk Cache Policy feature is supported on Non-RAID SAS/SATA physical disks (HDDs) connected to PERC hardware controllers and on RAID SAS/SATA physical disks connected to Software RAID S130 controller. This feature is not dependent on the controller mode and can function as usual when the controller is running in **RAID** or **HBA** mode.

Set the physical disk caching policy of all members of a virtual disk by enabling the **Disk Cache Policy**. When this feature is enabled, the physical disk writes data to the physical disk cache before writing it to the physical disk. Because it is faster to write data to the cache than to a disk, enabling this feature improves system performance.

The following are the disk cache policy options for Non-RAID disks connected to a PERC hardware controller:

- **Enabled** — **Disk Cache Policy** is enabled.
- **Disabled** — **Disk Cache Policy** is disabled.
- **Unchanged** — The disk uses the default write cache mode.

 **NOTE: The Unchanged option is applicable only for PERC 9 hardware controllers with the latest firmware version.**

 **NOTE: If you want to revert the Disk Cache Policy settings of the physical disks to the factory default settings, then you must restart the server.**

The following are the disk cache policy options for RAID disks connected to Software RAID S130 controller:

- **Default** — The disk uses the default write cache mode.
- **Enabled** — **Disk Cache Policy** is enabled.
- **Disabled** — **Disk Cache Policy** is disabled.
- **Inconsistent** — The disk cache policy is not uniform across all the physical disks in the virtual disk. For example, if a virtual disk contains three physical disks; the disk cache policy of the first physical disk is **Enabled**, the second physical disk is **Disabled**, and the third physical disk is **Default**. Then the disk cache policy state for the virtual disk is displayed as **Inconsistent**.

 **NOTE: The default setting for Disk Cache Policy for virtual disks based on SATA drives is Enabled, and for virtual disks based on SAS drives is Disabled.**

Related Task

- [Changing The Virtual Disk Policy](#)

Background Initialization On PERC Controllers

On PERC 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 improve write performance. For example, after the background initialization of a RAID 5 virtual disk completes, the parity information is initialized. After the background initialization of a RAID 1 virtual disk completes, the physical disks are mirrored.

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

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

Related links

- [Canceling Background Initialization](#)
- [Setting Background Initialization Rate](#)

Non-RAID Controller Description

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

 **NOTE:** Supported features may vary from controller to controller.

Non-RAID SCSI Controllers


The Broadcom PCIe U320 non-RAID controllers use Small Computer System Interface (SCSI) technology.


Non-RAID SAS Controllers

The following non-RAID controllers use Serial Attached SCSI (SAS) technology:


- SAS 12 Gbps HBA
- HBA 330 Mini
- HBA 330 Adapter

 **NOTE:** Storage Management only enumerates and provides the PCI slot information for Broadcom SAS 9207-8e and SAS 9300-8e controllers. As Storage Management does not monitor these non-RAID controllers, devices connected to these controllers are not enumerated, alerts are not logged, and SNMP traps are not sent for these controllers.

 **NOTE:** In Storage Management, the Broadcom SAS HBA controller page display properties such as: ID, Status, Name, Slot ID, State, Driver Version, Storport Driver Version, Number of Connectors, and Controller Tasks.

 **NOTE:** When a backplane without an expander (passive backplane) is connected to an HBA 330 controller, the enumeration of the physical disks is split between Connector 0 and Connector 1. For example, if the passive backplane is populated with a maximum of eight physical disks, the first four physical disks are listed under Connector 0 and the remaining four physical disks are listed under Connector 1. You will be able to view all the eight physical disks only if both bays are connected. If one bay is connected, only the physical disks associated with that bay are displayed.

 **NOTE:** For SAS 12 Gbps HBA and HBA 330 controller, the Available Reports that can be viewed are [View Slot Occupancy Report](#) and [View Physical Disk Firmware Version Report](#).

 **NOTE:** At times, the connector ID for SAS 12 Gbps HBA may not be displayed correctly on Storage Management. This occurs because of a firmware limitation in the SAS 12 Gbps HBA. However, this connector ID discrepancy does not cause any functional limitation.

Firmware Or Driver Versions

Use the firmware or driver versions window to view information about the controller firmware and drivers. For more information on firmware and drivers, see [Before Installing Storage Management](#).

Related links

[Firmware/Driver Properties](#)




Firmware/Driver Properties

The firmware and driver properties can vary depending on the model of the controller. The firmware and driver properties are listed in the table below.

 **NOTE:** The firmware and drivers listed in the *Server Administrator Release Notes* refer to the minimum supported version for these controllers. Later versions of the firmware and drivers are also supported. For most recent driver and firmware requirements, contact your service provider.

Table 9. Firmware/Driver Properties

Property	Definition
Firmware Version	Displays the version of the firmware that is installed on the controller.

Property	Definition
	 NOTE: Storage Management displays Not Applicable on some controllers for which the firmware version cannot be obtained.
Minimum Required Firmware Version	Displays the minimum firmware version that is required by Storage Management. This property is displayed only if the controller firmware does not meet the minimum requirement.
Driver Version	Displays the version of the driver that is installed on the controller.  NOTE: Storage Management displays Not Applicable on some controllers for which the driver version cannot be obtained.
Minimum Required Driver Version	Displays the minimum driver version that is required by Storage Management. This property is displayed only if the controller driver does not meet the minimum requirement.
Storport Driver Version	Displays the version of the storport driver that is installed on the system.
Minimum Required Storport Driver Version	Displays the minimum storport driver version required by Storage Management. This property is displayed if the operating system storport driver does not meet the minimum requirement. This option is applicable only on systems running Microsoft Windows operating system.  NOTE: To download the latest storport driver, see the Microsoft Knowledge Base article KB943545 at support.microsoft.com.

Controller Health

The controller **Health** page displays the status of the controller and the components attached to the controller.

Related links

- [Storage Component Severity](#)
- [Controller Information](#)
- [Controller Components](#)

Controller Components

For information on attached components, see:

- [RAID Controller Batteries](#)
- [Firmware Or Driver Versions](#)
- [Connectors](#)



NOTE: If you have connected the enclosure in Redundant path mode, the connectors are represented as Logical Connector.

- [Virtual Disks](#)







Controllers Properties And Tasks



Use the controllers properties and tasks window to view information about the controller and perform controller tasks.



NOTE: The firmware and drivers listed in the *Server Administrator Release Notes*, see the minimum supported version for these controllers. Later versions of the firmware and drivers are also supported. For the most recent driver and firmware requirements, contact your service provider. The controller properties can vary depending on the model of the controller.

Table 10. Controller Properties

Property	Definition
ID	<p>Displays the ID assigned to the controller by Storage Management. Storage Management numbers the controllers attached to the system starting with zero. This number is the same as the controller ID number reported by the <code>omreport</code> CLI Command. For information on Command Line Interface, see the <i>Server Administrator Command Line Interface User's Guide</i>.</p>
Status	<p>These icons represent the severity or health of the storage component.</p> <p>For more information, see Storage Component Severity.</p> <ul style="list-style-type: none"> •  — Normal/OK •  — Warning/Non-critical •  — Critical/Failure/Error
Name	<p>Displays the name of the controller.</p> <p> NOTE: For PowerEdge FD332 storage sled, the name is displayed as:</p> <ul style="list-style-type: none"> • Single controller — PERC FD33xS (Integrated RAID Controller <X> in Front Chassis Slot <Y>, where X displays the controller number and Y displays the front chassis slot number. • Dual controller — PERC FD33xD (Integrated RAID Controller <X> in Front Chassis Slot <Y>, where X displays the controller number and Y displays the front chassis slot number.
Slot ID	<p>Displays the slot to which the controller is attached. For PowerEdge FD332 storage sled, the name is displayed as: PCI Slot <Y> <X>, where Y displays the front chassis slot number and X displays the controller number.</p> <p> NOTE: You can also identify the slot ID by selecting the System → Main → System Chassis → Slots object and clicking the Information tab.</p>
State	<p>Displays the status of the controller. Possible values are:</p> <ul style="list-style-type: none"> • Ready — The controller is functioning normally. • Degraded — The controller has encountered a failure and is operating in a degraded state. • Failed — The controller has encountered a failure and is no longer functioning.
Firmware Version	<p>Displays the version of the firmware that is installed on the controller.</p> <p> NOTE: Storage Management displays Not Applicable on some controllers for which the firmware version cannot be obtained.</p>

Property	Definition
Minimum Required Firmware Version	Displays the minimum firmware version that is required by Storage Management. This property is displayed only if the controller firmware does not meet the minimum requirement.
Driver Version	Displays the version of the driver that is installed on the controller.  NOTE: Storage Management displays Not Applicable on some controllers for which the driver version cannot be obtained.
Minimum Required Driver Version	Displays the minimum driver version that is required by Storage Management. This property is displayed only if the controller driver does not meet the minimum requirement.
Number of Connectors	Displays the number of connectors on the controller. Each connector can be attached to physical disks or an enclosure. Depending on the controller type, the connector can be either a SCSI channel or a SAS port.
Rebuild Rate	Rebuild rate is the percentage of resources available on a system dedicated to rebuild a failed disk when necessary. For more information on rebuild rate, see Setting The Rebuild Rate .  NOTE: The value for the Revertible Hot Spare operation is the same as the value set for the Rebuild Rate property.
BGI Rate	The background initialization (BGI) rate is the percentage of resources available on a system dedicated to perform the background initialization of a virtual disk after it is created. For more information on BGI rate, see Setting Background Initialization Rate .
Check Consistency Rate	The check consistency rate is the percentage of resources available on a system dedicated to perform check consistency on a redundant virtual disk. For more information, see Performing A Check Consistency .
Reconstruct Rate	The reconstruct rate is the percentage of resources available on a system dedicated to reconstruct a disk group after adding a physical disk or changing the RAID level of a virtual disk residing on the disk group. For more information on reconstruct rate, see Setting The Reconstruct Rate .
Abort Check Consistency on Error	Enables you to stop the Check Consistency operation on error. This property is available only on controllers that have controller firmware version 6.1 and later.
Allow Revertible Hot Spare and Replace Member	Enables the automatic copying of data from a physical disk to a hot spare (in case of predictive failure) or from a hot spare to a physical disk (in case of replacement of a degraded disk). For more information, see Enabling Revertible Hot Spare .
Load Balance	Provides the ability to automatically use both controller ports or connectors connected to the same enclosure to route I/O requests. This property is available only on SAS controllers that have controller firmware version 6.1 and later.

Property	Definition
Auto Replace Member on Predictive Failure	Enables the automatic copying of data from a physical disk to a hot spare in case of predictive failure. Use this property in conjunction with the Allow Revertible Hot Spare and Replace Member property.
Redundant path view	Indicates whether Storage Management has detected a redundant path configuration. Storage Management detects a redundant path configuration when both controller ports are connected to the same enclosure that is in a unified mode. For more information, see Setting The Redundant Path Configuration .
Encryption Capable	Indicates whether the controller has the capability to support encryption. Possible values are Yes and No .
Encryption Key Present	Indicates whether the controller has an encryption key established. Possible values are Yes and No .
Encryption Mode	Indicates whether the controller is using Local Key Management (LKM) or None . For more information, see Managing The Encryption Key .
T10 Protection Information Capability	Indicates whether the controller supports data integrity. Possible values are Yes and No .
Cache Memory Size	Displays the size of the cache memory on the controller.
Patrol Read Mode	Displays the Patrol Read Mode setting for the controller. Possible values are: <ul style="list-style-type: none"> • Auto — Patrol read runs continuously on the system. When one iteration of the patrol read is complete, the next patrol read is scheduled to start within an interval specified by the controller. You do not have the option of manually starting or stopping the Patrol Read in this mode. • Manual — Allows you to manually start or stop the patrol read process. • Disabled — Indicates that the patrol read process is disabled. For more information about patrol read, see Setting The Patrol Read Mode and Starting And Stopping Patrol Read .
Patrol Read State	Displays the current state of the patrol read process. Possible values are: <ul style="list-style-type: none"> • Ready — The patrol read process is enabled and runs when next scheduled or when manually initiated. • Active — The patrol read process is running. • Stopped — The patrol read has been stopped. For more information about patrol read, see Setting The Patrol Read Mode .
Patrol Read Rate	Displays the percentage of system resources dedicated for running the Patrol Read operation. Patrol Read Rate changes the amount of system resources assigned for the patrol read task. The Patrol Read Rate can be configured between 0% and 100%, where:

Property	Definition
	<ul style="list-style-type: none"> 0% — indicates the lowest priority for controllers and has the least impact on the system performance. 100% — indicates the highest priority for controllers and has a greater impact on the system performance.
Patrol Read Iterations	Displays the number of Patrol Read Iterations . For more information about patrol read, see Setting The Patrol Read Mode .
Cluster Mode	Indicates whether the controller is part of a cluster configuration.
Persistent Hot Spare	Displays if the hot spare is persistent. The possible values are: <ul style="list-style-type: none"> Enabled — The slot corresponding to the hot spare drive is persistent. Any drive in the slot functions as a hot spare if the drive is qualified to be a hot spare. If the drive contains foreign data, it is overwritten. Disabled — The slot corresponding to the hot spare drive is not persistent. If the drive is removed from the slot and any drive is inserted, the slot stops functioning as a hot spare. You must manually assign the drive as a hot spare again.
Controller Tasks	Enables you to configure and manage the controller.
Available Reports	Enables you to view patrol read report, check consistency report, slot occupancy report, and physical disk firmware version report. For more information on available reports, see Available Reports .
Current Controller Mode	Displays the mode of the hardware controller that is selected. The possible values are RAID or HBA . To change the controller mode, see Changing The Controller Mode .
Front Chassis Slot	Displays the front chassis slot number. This property is applicable only for PowerEdge FD332 storage sled. For more information on PowerEdge FD332, see the <i>Dell PowerEdge FD332 Owner's Manual</i> at dell.com/poweredgemanuals .


Related links

- [How Do I Identify The Firmware Version That Is Installed](#)
- [Controller Tasks](#)
- [Available Reports](#)

Controller Tasks

To execute a controller task:

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Select a controller object.
3. Click **Information/Configuration** on the controller **Properties** page.
4. Select a task from the **Available Tasks** drop-down menu.
5. Click **Execute**.

 **NOTE: Different controllers support different features. The tasks displayed on the Available Tasks drop-down menu vary depending on the controller selected. The No Task Available option is displayed when no tasks can be performed due to controller or system configuration limitations.**

Controller Tasks

The following are the list of available tasks on a controller:

- [Creating A Virtual Disk](#)
- [Enabling The Controller Alarm](#)
- [Disabling The Controller Alarm](#)
- [Turning Off The Controller Alarm](#)
- [Testing The Controller Alarm](#)
- [Setting The Rebuild Rate](#)
- [Resetting Configuration](#)
- [Exporting The Controller Log File](#)
- [Foreign Configuration Operations](#)
- [Importing Foreign Configurations](#)
- [Importing Or Recovering Foreign Configurations](#)
- [Clearing Foreign Configuration](#)
- [Setting Background Initialization Rate](#)
- [Setting Check Consistency Rate](#)
- [Setting The Reconstruct Rate](#)
- [Setting The Patrol Read Mode](#)
- [Starting And Stopping Patrol Read](#)
- [Managing The Preserved Cache](#)
- [Changing The Controller Properties](#)
- [Managing The Physical Disk Power](#)
- [Managing The Encryption Key](#)
- [Converting To RAID Capable Disks](#)
- [Converting To Non-RAID Disks](#)
- [Changing The Controller Mode](#)
- [Auto Configure RAID0 Operation](#)

Creating A Virtual Disk

Does my controller support this feature? See [Supported Features](#).

To launch the **Create Virtual Disk Express Wizard**, select the **Create Virtual Disk** task.

Related links

[Create Virtual Disk Express Wizard](#)

Enabling The Controller Alarm

Does my controller support this feature? See [Supported Features](#).

To enable the alarm on the controller, select the **Enable Alarm** task. When enabled, the alarm sounds in the event of a device failure.

Related links

[Using Alarms To Detect Failures](#)

Disabling The Controller Alarm

Does my controller support this feature? [Supported Features](#).

Select the **Disable Alarm** task to disable the alarm on the controller. When disabled, the alarm does not sound in the event of a device failure.

Turning Off The Controller Alarm

Does my controller support this feature? See [Supported Features](#).

Select the **Quiet Alarm** task to turn off the controller alarm. However, the controller alarm remains enabled for future device failure.

Testing The Controller Alarm

Does my controller support this feature? See [Supported Features](#).

Select the **Test Alarm** task to test whether the controller alarm is functional. The alarm sounds for about 2 seconds.

Setting The Rebuild Rate

Does my controller support this feature? See [Supported Features](#).

The **Set Rebuild Rate** task allows you to change the rebuild rate for a controller.

During a rebuild, the complete contents of a physical disk are reconstructed. The rebuild rate, configurable between 0% and 100%, represents the percentage of the system resources dedicated for rebuilding failed physical disks. At 0%, the rebuild has the lowest priority for the controller, takes the maximum time to complete, and is the setting with the least impact to system performance. A rebuild rate of 0% does not mean that the rebuild is stopped or paused.

At 100%, the rebuild is the highest priority for the controller, minimizes the rebuild time, and is the setting with the most impact to system performance.

On the PERC controllers, the controller firmware also uses the rebuild rate setting to control the system resource allocation for the following tasks:

- [Performing A Check Consistency](#)
- Background Initialization. See [Cancelling Background Initialization](#)
- Full Initialization. A BIOS setting determines whether a full or fast initialization occurs. See [Format, Initialize, Slow And Fast Initialize](#)
- Reconfigure. See [Virtual Disk Task: Reconfigure \(Step 1 of 3\)](#)

Related links

[Rebuilding Redundant Information](#)

[To Change The Rebuild Rate](#)

[To locate Set Rebuild Rate In Storage Management](#)

To Change The Rebuild Rate

To change the rebuild rate:

1. Type a numerical value in the **Set New Rebuild Rate (0-100)** field. The value must be within the 0–100 range.
2. Click **Apply Changes**.
To exit and cancel your changes, click **Return to Previous Page**.

To locate Set Rebuild Rate In Storage Management

1. Expand the **Storage** tree object to display the controller objects.
2. Select a controller object.
3. Click **Information/Configuration**.
4. Select **Set Rebuild Rate** from the **Available Tasks** drop-down menu.
5. Click **Execute**.




You can also locate this task from the **Change Controller Properties** drop-down menu.

Resetting The Controller Configuration

Does my controller support this feature? See [Supported Features](#).

The **Reset Configuration** task allows you to erase all the information on the controller so that you can perform a fresh configuration. This operation destroys all data and virtual disks on the controller and unassigns any hot spares.

You must completely reconfigure the storage after performing this reset operation.

-  **CAUTION:** Resetting a configuration permanently destroys all data on all virtual disks attached to the controller. If the system or boot partition resides on these virtual disks, it is destroyed.
-  **NOTE:** Resetting the controller configuration does not remove a foreign configuration. To remove a foreign configuration, select **Clear Foreign Configuration** task.
-  **NOTE:** This task is not supported on PERC hardware controllers running in HBA mode.

Related links

- [Reset The Controller Configuration](#)
- [To Locate Reset Configuration In Storage Management](#)
- [Changing The Controller Properties](#)

Reset The Controller Configuration

To reset the controller configuration:

1. Review the virtual disks that are destroyed by resetting the controller configuration. Make backups as necessary. Click **Blink** at the bottom of the page to blink the physical disks included in the virtual disks.
2. Click **Reset Configuration** when you are ready to erase all information on the controller.
To exit without resetting the controller configuration, click **Go Back to Previous Page**.

To Locate Reset Configuration In Storage Management

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Select a controller object.
3. Click **Information/Configuration**.
4. Select **Reset Configuration** from the **Available Tasks** drop-down menu.
5. Click **Execute**.

You can also locate this task from the **Change Controller Properties** drop-down menu.

Exporting The Controller Log File

Does my controller support this feature? See [Supported Features](#).

The **Export Log** task exports the controller log to a text file. This log provides detailed information on the controller activities and can be useful for troubleshooting.

On a system running Microsoft Windows, the log file is exported to the **windows** or **winnt** directory. On a system running Linux, the log file is exported to the **/var/log** directory.

Depending on the controller, the log file name is **afa_<mmdd>.log** or **lsi_<mmdd>.log**, where **<mmdd>** is the month and date.

 **NOTE:** In the VMware ESXi environment, only one log file is created (**lsiexport.log**). If the file exists, exporting the log file overwrites the existing log file.

 **NOTE:** Controllers without cache cannot store logs and export log files.

Related links

- [What Is PCIe SSD?](#)
- [Export The Controller Log File](#)
- [To Locate Export Log In Storage Management](#)
- [To Locate Export Log In Storage Management For NVMe PCIe SSDs](#)
- [Changing The Controller Properties](#)

Export The Controller Log File

1. Click **Export Log File** when ready.
2. To exit without exporting the controller log file, click **Return to Previous Page**.

To Locate Export Log In Storage Management

To locate this task in storage management:

1. Expand the **Storage** tree object to display the controller objects.
2. Select a controller object.
3. Click **Information/Configuration**.
4. Select **Export Log** from the **Available Tasks** drop-down menu.
5. Click **Execute**.

You can also locate this task from the **Change Controller Properties** drop-down menu.

Foreign Configuration Operations

Does my controller support this feature? See [Supported Features](#).

The **Foreign Configuration Operations** task provides a preview of the foreign configurations that you can import.

 **NOTE: This task is not supported on PERC hardware controllers running in HBA mode.**

 **NOTE: Foreign Configuration Operations task is available only on PERC 6 and SAS controllers with firmware versions 6.1 and later.**

A foreign configuration is data residing on physical disks that have been moved from one controller to another. Virtual disks residing on physical disks that have been moved are considered to be a foreign configuration.

 **NOTE: It is not recommended to remove an external enclosure cable while the operating system is running on the system. Removing the cable could result in a foreign configuration when the connection is re-established.**

The **Foreign Configuration Operations** task is displayed only when a controller detects a foreign configuration. Select this option and click **Execute** to display the **Foreign Configuration Preview** page.

The **Foreign Configuration Preview** page provides a preview of the foreign disks and enables you to perform operations such as, importing, recovering, or clearing the foreign disks. You can also import or clear a locked foreign configuration.

If any foreign configurations locked using **Local Key manager (LKM)** are detected, the associated **Encryption Key Identifier** is displayed prompting you to provide the corresponding passphrase to unlock the drives.

To avoid unlocking foreign configurations and to proceed to preview, import, or clear a foreign configuration that has not been locked, click **Skip** or **Continue**.

If you do not want to import or clear the foreign configurations, or if there is loss of the associated passphrase of the corresponding **Encryption Key Identifier**, perform the **Cryptographic Erase** task for the physical disks.

 **CAUTION: Performing the Cryptographic Erase task erases all data on the physical disk.**

Some conditions, such as an unsupported RAID level or an incomplete disk group, can prevent the import or recovery of foreign virtual disks.

Related links

[Foreign Configuration Properties](#)

[To Locate Foreign Configuration Operations Task In Storage Management](#)

[Importing Foreign Configurations](#)

[Importing Or Recovering Foreign Configurations](#)

[Clearing Foreign Configuration](#)

[Managing The Encryption Key](#)





[Performing Cryptographic Erase](#)

[Performing Cryptographic Erase](#)

Foreign Configuration Properties

The following table describes the properties that are displayed for the Foreign Disks and Global Hot Spares.

Table 11. Foreign Configuration Properties

Property	Definition
<p>Status</p>	<p>These icons represent the severity or health of the storage component.</p> <ul style="list-style-type: none"> •  — Normal/OK •  — Warning/Non-critical •  — Critical/Failure/Error <p>For more information, see Storage Component Severity.</p>
<p>Name</p>	<p>Displays the name of the foreign configuration and is available as a link. This link enables you to access the physical disks that constitute the foreign disk.</p>
<p>State</p>	<p>Displays the current state of the foreign configuration. Possible values are:</p> <ul style="list-style-type: none"> • Ready — The foreign disk can be imported and functions normally after import. • Degraded — The foreign disk is in degraded state and rebuilds after import. • Failed — The foreign disk has encountered a failure and is no longer functioning. You cannot import the foreign configuration. <p>The foreign configuration may be in degraded or failed state due to any of the following reasons:</p> <ul style="list-style-type: none"> • Missing physical disk — One of the physical disk(s) in the potential virtual disk is missing or not available. • Missing span — One or more spans of a hybrid virtual disk is missing. • Stale physical disks — One or more physical disks in the configuration may contain out-of-date data relating to other disks of that virtual disk. Hence, the data integrity of the imported virtual disk is not intact. • Unsupported configuration of the virtual disk — The virtual disk has an unsupported RAID level. • Import and Export — The virtual disks available for import exceed the number of virtual disks available for export. • Incompatible physical disks — Configuration on the physical disks is not recognized by the RAID firmware. • Orphan drive — A physical disk in the foreign configuration has configuration information that matches another physical disk that is already a part of an array (either a foreign or a native array). <p> NOTE: For other applicable physical disk tasks and properties, see Physical Disk Or Physical Device Properties and Physical Disk Or Physical Device Tasks.</p>
<p>Layout</p>	<p>Displays the RAID level of the foreign configuration.</p>
<p>Remarks</p>	<p>Provides information about the foreign virtual disk. If the virtual disk cannot be imported, the reason for failure is displayed.</p> <ul style="list-style-type: none"> • Exceeded maximum — The number of virtual disks selected for import has exceeded the maximum number of supported disks. • Missing physical disk or Missing span — One or more physical disk(s) or span(s) in the virtual disk to be imported is missing. • Unsupported — The selected RAID level is not supported on this controller. • Orphan drive — The physical disk has been replaced and is no longer a part of the RAID volume. The configuration should be cleared. • Stale physical disk — The physical disk to be imported in the virtual disk has outdated data. • Partially foreign — The virtual disk is part of an already existing configuration. Some physical disks in this virtual disk are foreign.

Property	Definition
Dedicated Hot Spare	Displays whether the foreign disk is a dedicated hot spare.

Based on the properties information, you can decide whether you want to import, recover, or clear the foreign configuration.

To Locate Foreign Configuration Operations Task In Storage Management

For SAS controllers with firmware version 6.1:

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Select a controller object.
3. Click **Information/Configuration** on the controller **Properties** page.
4. Select **Foreign Configuration Operations** from the **Controller Tasks** drop-down menu.
5. Click **Execute**.

Importing Foreign Configurations

Some controllers enable you to import a foreign configuration so that virtual disks are not lost after moving the physical disks.

You can import a foreign configuration only if it contains a virtual disk that is either in a **Ready** or **Degraded** state. In other words, all of the virtual disk data must be present, but if the virtual disk is using a redundant RAID level, the additional redundant data is not required.

For example, if the foreign configuration contains only one side of a mirror in a RAID 1 virtual disk, then the virtual disk is in a **Degraded** state and can be imported. On the other hand, if the foreign configuration contains only one physical disk that was originally configured as a RAID 5 using three physical disks, then the RAID 5 virtual disk is in a **Failed** state and cannot be imported.

In addition to virtual disks, a foreign configuration may consist of a physical disk that was assigned as a hot spare on one controller and then moved to another controller. The **Import Foreign Configuration** task imports the new physical disk as a hot spare. If the physical disk was set as a dedicated hot spare on the previous controller, but the virtual disk to which the hot spare was assigned is no longer present in the foreign configuration, then the physical disk is imported as a global hot spare.

The **Import Foreign Configuration** task is only displayed when the controller has detected a foreign configuration. You can also identify whether a physical disk contains a foreign configuration (virtual disk or hot spare) by checking the physical disk state. If the physical disk state is **Foreign**, then the physical disk contains all or some portion of a virtual disk or has a hot spare assignment.

If you have an incomplete foreign configuration which cannot be imported, you can use the [Clearing Foreign Configuration](#) option to erase the foreign data on the physical disks.

 **NOTE: The task of importing foreign configuration imports all virtual disks residing on physical disks that have been added to the controller. If more than one foreign virtual disk is present, all the configurations are imported.**

Related links

[Foreign Configuration Properties](#)

Importing Or Recovering Foreign Configurations

The recover operation attempts to restore degraded, failed, or missing virtual disks to a healthy state. A virtual disk may be in a degraded, failed, or missing state after losing communication with the controller due to a power loss, faulty cable connection, or other failure. A rebuild or background initialization may automatically initiate after the recover operation completes.

The virtual disk data may be inconsistent after recovery. Verify the virtual disk data after the **Import/Recover Foreign Configuration** task completes.

In some cases, the virtual disk data is incomplete and it is not possible to successfully recover the virtual disk.

To import or recover a foreign configuration:

Click **Import/Recover** to import or recover all virtual disks residing on physical disks attached to the controller.

To exit without importing or recovering the foreign configuration, click **Cancel**.

To Locate Import Or Recover Foreign Configuration In Storage Management

For SAS controllers with firmware versions 6.1 and later:


1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Select a controller object.
3. Click **Information/Configuration** on the controller **Properties** page.
4. Select **Foreign Configuration Operations** from the **Controller Tasks** drop-down menu.
5. Click **Execute**.
6. In the **Foreign Configuration Preview** page, click **Import/Recover**.

For SAS controllers with firmware version 6.0 and earlier, select **Import/Recover Foreign Configuration** from the **Controller** tasks.

Clearing Foreign Configuration

Does my controller support this feature? See [Supported Features](#).

After moving a physical disk from one controller to another, you may find that the physical disk contains all or some portion of a virtual disk (foreign configuration). You can identify whether a previously used physical disk contains a foreign configuration (virtual disk) by checking the physical disk state. If the physical disk state is **Foreign**, then the physical disk contains all or some portion of a virtual disk. Use the **Clear Foreign Configuration** task to clear or erase the virtual disk information from the newly attached physical disks.

 **NOTE: The Clear Foreign Configuration task permanently destroys all data residing on the physical disks that are added to the controller. If more than one foreign virtual disk is present, all the configurations are erased. You may prefer to import the virtual disk rather than destroy the data.**

To clear a foreign configuration:

Click **Clear Foreign Configuration** to clear or erase all virtual disks residing on physical disks is added to the controller.

To exit without clearing the foreign configuration, click **Cancel**.

Related links

[To Locate Clear Foreign Configuration In Storage Management](#)

[Importing Foreign Configurations](#)

[Importing Or Recovering Foreign Configurations](#)

To Locate Clear Foreign Configuration In Storage Management

For SAS controllers with firmware versions 6.1 and later:

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Select a controller object.
3. Click **Information/Configuration** on the controller **Properties** page.
4. Select **Foreign Configuration Operations** from the **Controller Tasks** drop-down menu.
5. Click **Execute**.
6. In the **Foreign Configuration Preview** page, click **Clear**.




For SAS controllers with firmware version 6.0 and earlier, select **Clear Foreign Configuration** from the **Controller Tasks** drop-down menu.

Physical Disks In Foreign Virtual Disks

The **Physical Disks in Foreign Virtual Disks** page displays the physical disks and the dedicated hot spare, if any, included in the foreign configuration.

The following table describes the properties for physical disks in the foreign configuration.

Table 12. Physical Disk Properties

Property	Definition
Status	<p>These icons represent the severity or health of the storage component.</p> <ul style="list-style-type: none"> •  — Normal/OK •  — Warning / Non-critical •  — Critical/Failure/Error <p>For more information, see Storage Component Severity.</p>
Name	Displays the name of the physical disk. The name is comprised of the connector number followed by the disk number.
State	Displays the current state of the physical disk.
After Import State	<p>Displays the after-import state of the physical disk. The physical disk can be imported in any of the following states:</p> <ul style="list-style-type: none"> • Online — The physical disk is part of the imported virtual disk and functions normally. • Offline — The physical disk is offline after import to the virtual disk. • Foreign — The virtual disk containing the physical disk cannot be imported and the physical disk remains in foreign state. • Rebuild — After import of virtual disk, the physical disk rebuilds. • Replacing — A Replace Member Disk task is performed on the physical disk. For more information on replacing a member disk, see Replacing A Member Disk and Enabling Reversible Hot Spare.
Capacity	Displays the capacity of the disk.
Failure Predicted	<p>Displays whether the physical disk has received a Self-Monitoring Analysis and Reporting Technology (SMART) alert and is therefore predicted to fail. For more information on SMART predictive failure analysis, see Monitoring Disk Reliability On RAID Controllers. For information on replacing the physical disk, see Replacing A Physical Disk Receiving SMART Alerts.</p> <p>You may also want to review the alert log to see whether the physical disk has generated alerts pertaining to a SMART predictive failure. These alerts can assist you in identifying the cause of the SMART alert. The following alerts may be generated in response to a SMART alert:</p> <ul style="list-style-type: none"> • 2094 • 2106 • 2107 • 2108 • 2109 • 2110 • 2111 <p>For information on alert messages, see the <i>Server Administrator Messages Reference Guide</i>.</p>
Progress	Displays the progress of an operation being performed on the physical disk.
Bus Protocol	<p>Displays the technology that the physical disk is using. Possible values are:</p> <ul style="list-style-type: none"> • SAS — Serial Attached SCSI • SATA — Serial Advanced Technology Attachment
Device Protocol	Displays the device protocol of the physical device, such as Non-Volatile Memory Express (NVMe).

Property	Definition
Certified	Indicates that the drive has firmware which has been tested and fully qualified by the service provider. Drives that are not certified by the service provider may work but, are not supported and recommended for use in servers.
Media	Displays the media type of the physical disk. The possible values are: <ul style="list-style-type: none"> • HDD — Hard Disk Drive. An HDD is a non-volatile storage device which stores digitally-encoded data on rapidly rotating platters with magnetic surfaces. • SSD — Solid-State Drive. An SSD is a data storage device that uses solid-state memory to store persistent data. • Unknown — Storage Management is unable to determine the media type of the physical disk.
Used RAID Disk Space	Displays the amount of the physical disk space that is used by the virtual disks on the controller. This property is not applicable for physical disks attached to non-RAID controllers. In certain circumstances, the Used RAID Disk Space displays a value of zero (0) even though a portion of the physical disk is actually being used. This occurs when the used space is 0.005GB or less. The algorithm for calculating the used disk space rounds off a figure of 0.005GB or less than zero. Used disk space that is between 0.006GB and 0.009GB is rounded off to 0.01GB.
Available RAID Disk Space	Displays the amount of available space on the disk. This property is not applicable for physical disks attached to non-RAID controllers.
Hot Spare	Indicates whether the disk has been assigned as a hot spare. This property is not applicable for physical disks attached to non-RAID controllers.
Vendor ID	Displays the hardware vendor of the disk.
Product ID	Displays the product ID of the device.
Firmware Revision	Displays the firmware version of the physical device.
Serial No.	Displays the serial number of the disk.
PCIe Negotiated Link Speed	Displays the current negotiated transfer speed of the physical device in GT/s.
PCIe Maximum Link Speed	Displays the capable transfer speed of the physical device in GT/s.
Manufacture Day	Displays the day of the month on which the physical disk was manufactured.
Manufacture Week	Displays the week of the year during which the physical disk was manufactured.
Manufacture Year	Displays the year in which the physical disk was manufactured.
SAS Address	Displays the SAS address of the physical disk. The SAS address is unique to each SAS disk.
After Import Status	Displays the status of the physical disk after importing the foreign configuration. Possible values are: <ul style="list-style-type: none"> • Foreign • Online • Offline • Replaced • Rebuild
Encryption Capable	Displays whether the physical disk is a Self Encryption Disk (SED). The possible values are Yes and No .
Encrypted	Displays whether the physical disk is encrypted to the controller. The possible values are Yes and No . For a non-SED the value is N/A .

Property	Definition
Part Number	Displays the unique Bill Of Materials assignment number for a physical disk. The numbers four through eight represent the service provider part number for that model drive.
PCIe Negotiated Link Width	Displays the current negotiated link width of the physical device.
PCIe Maximum Link Width	Displays the capable link width of the physical device.

Setting Background Initialization Rate

Does my controller support this feature? See [Supported Features](#).

The **Set Background Initialization Rate** task changes the amount of system resources dedicated to the background initialization task.

The background initialization rate, configurable between 0% and 100%, represents the percentage of the system resources dedicated for running the background initialization task. At 0%, the background initialization has the lowest priority for the controller, takes maximum time to complete, and has the least impact to system performance. A background initialization rate of 0% does not mean that the background initialization is stopped or paused.

At 100%, the background initialization is the highest priority for the controller. The background initialization time is minimized and has the most impact to system performance.

Related links

[Background Initialization On PERC Controllers](#)

To Change The Background Initialization Rate For A Controller

1. Type a numerical value in the **Set New BGI Rate (0-100)** field. The value must be within the 0–100 range.
2. Click **Apply Changes**. To exit and cancel your changes, click **Return to Previous Page**.

To Locate Background Initialization Rate In Storage Management

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Select a controller object.
3. Click **Information/Configuration** on the controller **Properties** page.
4. Select **Set Background Initialization Rate** from the **Available Tasks** drop-down menu.
5. Select **Execute**.

You can also locate this task from the **Change Controller Properties** drop-down menu.

Related links

[Changing The Controller Properties](#)

Setting The Check Consistency Rate

Does my controller support this feature? See [Supported Features](#).

The **Set Check Consistency Rate** task changes the amount of system resources dedicated to the check the consistency rate.

The check consistency rate, configurable between 0% and 100%, represents the percentage of the system resources dedicated for running the check consistency task. At 0%, the check consistency has the lowest priority for the controller, takes maximum time to complete, and has least impact to system performance. A check consistency rate of 0% does not mean that the check consistency is stopped or paused.

At 100%, the check consistency is the highest priority for the controller. The check consistency time is minimized and has most impact to system performance.

Related links

[Performing A Check Consistency](#)

To Change The Check Consistency Rate For A Controller

1. Type a numerical value in the **Set New Check Consistency Rate (0-100)** field. The value must be within the 0–100 range.
2. Click **Apply Changes**.
To exit and cancel your changes, click **Return to Previous Page**.

To Locate Set Check Consistency Rate In Storage Management

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Select a controller object.
3. Click **Information/Configuration** on the controller **Properties** page.
4. Select **Set Check Consistency Rate** from the **Available Tasks** drop-down menu.
5. Click **Execute**.
You can also locate this task from the **Change Controller Properties** drop-down menu.

Related links

[Changing The Controller Properties](#)

Setting The Reconstruct Rate

Does my controller support this feature? See [Supported Features](#).

The **Set Reconstruct Rate** task changes the amount of system resources dedicated for the reconstruct task.

The reconstruct task recreates the virtual disk after you have changed the RAID level or reconfigured the virtual disk. The reconstruct rate, configurable between 0% and 100%, represents the percentage of the system resources dedicated for running the reconstruct task. At 0%, the reconstruct has the lowest priority for the controller, takes maximum time to complete, and has least impact to system performance. A reconstruct rate of 0% does not mean that the reconstruct is stopped or paused.

At 100%, the reconstruct is the highest priority for the controller, the reconstruct time is minimized, and has most impact to system performance.

Related links

[Reconfiguring Or Migrating Virtual Disks](#)

To Change The Reconstruct Rate For A Controller

1. Type a numerical value in the **Set New Reconstruct Rate (0-100)** field. The value must be within the 0–100 range.
2. Click **Apply Changes**. To exit and cancel your changes, click **Return to Previous Page**.

To Locate Set Reconstruct Rate In Storage Management

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Select a controller object.
3. Click **Information/Configuration** on the controller **Properties** page.
4. Select **Set Reconstruct Rate** from the **Available Tasks** drop-down menu.
5. Select **Execute**.
You can also locate this task from the **Change Controller Properties** drop-down menu.

Related links

[Changing The Controller Properties](#)

Setting The Redundant Path Configuration


Does my Controller Support This Feature? See [Supported Features](#).

Redundant path is supported only on external PERC cards that use firmware version 6.1 and later. A redundant path internal to the system is not supported. MD1xxx enclosures are supported.

For redundant paths, the enclosure must be in the **Unified** mode however specific port connections are not required. A connection from either controller ports to either EMM **In** ports creates the redundant path as long as two cables are used. But, Redundant Path

continues to appear in the field even as the redundant path is removed. The redundant path ceases to appear only when it is cleared in the storage management.

In a daisy-chain scenario, more than one enclosure can be connected to a controller in a redundant path mode. You can connect up to four MD1400 and MD1420 to a PERC H840 controller and SAS 12 Gbps adapter in a daisy-chain enclosure. For an example of a daisy-chain configuration (for PERC 6/E controller), see the following image:



If the communication channel between the connector and the first enclosure is lost, the redundant path configuration is lost. In this case, the health of the logical connector is displayed as critical. Navigate to the **Information/Configuration** page of the logical connector to view details of the [Path Health](#). For a brief outline of this scenario, see the following table:

Table 13. Path between Controller and Enclosure 1

Health of Logical Connector	Path between Controller and Enclosure 1	
	Connector 0 (C0)	Connector 1 (C1)
	Available	Available
	Available	Disconnected
	Disconnected	Available

However, if the communication channel between any two enclosures is lost, the redundant path configuration is degraded and the health of the logical connector is displayed as degraded. For a brief outline of this scenario, see the following table.

Table 14. Path between Enclosure *n* and Enclosure *n*+1

Health of Logical Connector	Path between Enclosure <i>n</i> and Enclosure <i>n</i> +1	
	Connector 0 (C0)	Connector 1 (C1)
	Available	Available
	Available	Disconnected
	Disconnected	Available

In the above scenario, the enclosure status is displayed in warning mode. Clicking **Information/Configuration** in the **Enclosures** page displays all enclosure components (EMMs, Fans, Physical Disks, Power Supplies, and Temperature) in normal condition. To view the

Path Failure message to indicate that the enclosure has lost a communication path to the controller, indicating that the enclosure is no longer in redundant path mode.

Related links

[Path Health](#)

[Clearing The Connectors Redundant Path View](#)

Clearing The Redundant Path View

Consider a case where you reboot your system and Storage Management, displays the logical connector with a path failure message. It is possible that you may have intentionally unplugged the second connector. In this case, the path failure message is not relevant. There could be a fault in the connected cable or the cable may not be connected properly to the controller. In both cases, Storage Management displays that the system was in redundant path configuration before reboot and is no longer in this configuration. If you are sure you, do not want the redundant path mode, clear the existing redundant path view using **Clear Redundant Path View** provided in the [Changing The Controller Properties](#) controller task. Selecting this option clears the redundant path view and the connectors are represented on the user interface as **Connector 0** and **Connector 1**.

Related links

[Setting The Redundant Path Configuration](#)

[Logical Connector Properties And Tasks](#)

Setting The Patrol Read Mode

 **NOTE: This task is not supported on PERC hardware controllers running in HBA mode.**

Does my controller support this feature? See [Supported Features](#).

Patrol read identifies disk errors in order to avoid disk failures and data loss or corruption. The **Set Patrol Read** task is applicable only for disks used as virtual disks or hot spares.

The **Set Patrol Read** task runs in the background and corrects, when possible. When the **Set Patrol Read mode** is set to **Auto**, patrol read is initiated when the controller is idle for a period of time and when no other background tasks are active. In this scenario, the patrol read enhances the system performance as disk errors can be identified and corrected when there is not input/output activity on the disk.

The controller adjusts the amount of system resources dedicated for patrol read based on the amount of controller activity that is competing with the Patrol Read task. When the controller activity is high, fewer system resources are dedicated to the patrol read task.

Patrol Read does not run on a physical disk in the following circumstances:

- The physical disk is not included in a virtual disk or is assigned as a hot spare.
- The physical disk is included in a virtual disk that is currently undergoing one of the following:
 - Rebuild
 - Reconfiguration or reconstruction
 - Background initialization
 - Check consistency

In addition, the Patrol Read suspends during heavy I/O activity and resumes when the I/O is finished.

Related links

[Starting And Stopping Patrol Read](#)

To Set The Patrol Read Mode

Select the desired Patrol Read Mode option. The options available are:

- **Auto** — Initiates the Patrol Read task. After the task is complete, it automatically runs again within a specified period. For example, on some controllers the Patrol Read runs every four hours and on other controllers, the Patrol Read runs every seven days. The Patrol Read task runs continuously on the system starting again within the specified period after each iteration of the task completes. If the system reboots while the Patrol Read task is running in **Auto** mode, the Patrol Read restarts at zero percent (0%). When the Patrol Read task is set to **Auto** mode, you cannot start or stop the task. **Auto** mode is the default setting.

 **NOTE: For more information on how often the Patrol Read task runs when in Auto mode, see your controller documentation.**

- **Manual** — Enables you to start and stop the Patrol Read task using **Start and Stop Patrol Read**. Setting the mode to **Manual** does not initiate the Patrol Read task. If the system reboots while Patrol Read is running in **Manual** mode, Patrol Read does not restart.
- **Disabled** — Prevents the Patrol Read task from running on the system.

To Locate Set Patrol Read Mode In Storage Management

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Select a controller object.
3. Click **Information/Configuration** on the controller **Properties** page.
4. Select **Set Patrol Read** mode from the **Available Tasks** drop-down menu.
5. Click **Execute**.

You can also locate this task from the **Change Controller Properties** drop-down menu.

Related links

[Changing The Controller Properties](#)

Starting And Stopping Patrol Read

Does my controller support this feature? See [Supported Features](#).


 **NOTE: The Patrol Read operation is not supported on solid-state drives (SSD).**

 **NOTE: This task is not supported on PERC hardware controllers running in HBA mode.**

When the **Set Patrol Read** mode is set to **Manual**, you can start the Patrol Read task or stop the task when it is running. There are certain conditions under which the Patrol Read task cannot be run.

To start or stop the Patrol Read Task:

Click **Start Patrol Read** or **Stop Patrol Read**.

 **NOTE: On PERC 9 family of hardware controllers, the Stop Patrol Read task displays: Patrol Read Aborted. On controllers before PERC 9, the Stop Patrol Read task displays: Patrol Read Stopped.**

To exit without starting or stopping the Patrol Read, click **Go Back to Previous Page**.

Related links

[Setting The Patrol Read Mode](#)

Locate Start Stop Patrol Read In Storage Management

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Select a controller object.
3. Click **Information/Configuration** on the controller **Properties** page.
4. Select **Start Patrol Read** or **Stop Patrol Read** from the **Available Tasks**.
5. Click **Execute**.

You can also locate this task from the **Change Controller Properties** drop-down menu.

Related links

[Changing The Controller Properties](#)

Changing The Controller Properties

 **NOTE: This task is not supported on PERC hardware controllers running in HBA mode.**

Does my controller support this feature? See [Supported Features](#).

The **Change Controller Properties** task provides you the option to change multiple controller properties simultaneously. This task is available only on SAS controllers with firmware version 6.1 and later.

You can change any or all of the following properties using the **Change Controller Properties** task:

- Rebuild Rate
- BGI Rate
- Check Consistency Rate
- Reconstruct Rate
- Abort check consistency on error
- Reversible Hot Spare
- Load Balance
- Auto replace member on predictive failure
- Redundant path view
- Persistent hot spare

 **NOTE:** You can also set these properties using the command-line interface. For more information, see the *Server Administrator Command-Line Interface User's Guide*.

To Locate Change Controller Properties In Storage Management

1. In the **Server Administrator** window, under the system tree, select **Storage**.
2. On the **Storage Dashboard** page, select **Change Controller Properties** from the **Available Tasks** drop-down menu.
3. Click **Execute**.


To Locate Change Controller Properties In Storage Management: Method 2

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Select a controller object.
3. Click **Information/Configuration** on the controller **Properties** page.
4. Select **Change Controller Properties ...** from the **Controller Tasks** drop-down menu.
5. Click **Execute**.

Managing The Physical Disk Power

Does my controller support this feature? See [Supported Features](#).

The **Manage Physical Disk Power** task allows you to manage the power consumed by the physical disks.

 **NOTE:** The **Manage Physical Disk Power** task is provided with H330 cards by spinning down the hot spares and unconfigured disks. PERC H730P, H730, H740P, and H840 cards also support the **Manage Physical Disk Power** task with additional power-saving modes, **Spin Down Configured Drives**, and **Automatic Disk Power Saving (Idle C)** options.

 **NOTE:** This task is not supported on PERC hardware controllers running in HBA mode.

You can enable one of the following modes to manage the power consumed:

- **No Power Savings Mode** — This is the default mode for the controller. In this mode, all power-saving features are disabled.
- **Balanced Power Savings Mode** — Provides good power savings while limiting I/O latency.
- **Maximum Power Savings Mode** — Provides maximum power savings for all drives.
- **Customized Power Savings Mode** — Allows you to customize the power-saving settings. The default values are populated when you select this power mode. You can select or clear the features you want to enable. You can select **Quality of Service (QoS)** to customize power saving on configured drives by setting the **Start Time** and **Time interval** to spin up.

To enable the **Quality of Service (QoS)** feature:

1. Select **Customized Power Savings Mode**.
2. Select **Enable** for the **Spin Down Configured Drives** option


Related links

- [Properties In Manage Physical Disk Power Option](#)
- [Physical Disk Power On Unconfigured Drives And Hot Spares](#)
- [Managing Physical Disk Power Using The Customized Power Savings Mode](#)
- [Managing Physical Disk Power Using The QoS Option](#)
- [Managing The Time Interval For The QoS Option](#)

Properties In Manage Physical Disk Power Option

The following table displays the properties in the **Manage Physical Disk Power** option:

Table 15. Manage Physical Disk Power Properties

Property	Definition
Spin Down Unconfigured Drives	The Enabled option spins down the unconfigured disks if they are unattended for a specified interval of time.
Spin Down Hot Spares	The Enabled option spins down the hot spares if no read-write operation takes place on the hot spare in a specified interval of time.
Spin Down Configured Drives	The Enabled option spins down the configured disks if they are unattended for a specified interval of time.
Automatic Disk Power Saving (Idle C)	Enables or disables the Automatic Disk Power Saving (Idle C) feature for additional power saving. When enabled, it enables the new generation power savings without affecting the legacy drives.
Time Interval for Spin Down	Sets the time interval after which the hot spares and unconfigured drives spins down.
Quality Of Service (QOS)	
Enable Quality Of Service Settings	Select to set the start time and the time interval for spin up activity at the virtual disk level.  NOTE: This option is available only if the Spin Down Configured Drives option is selected.
Start Time (HH:MM)	Displays the start time of the battery learn cycle. This option is enabled only if Enable Quality Of Service Settings is selected.
Time Interval for Spin Up (in Hours)	Displays the spin up time interval for the battery learn cycle. The time interval can range from 1-24 hours.

Related links

- [Managing The Physical Disk Power](#)

Physical Disk Power On Unconfigured Drives And Hot Spares

1. Select **Enabled** for the **Spin Down Unconfigured Drives** and **Spin Down Hot Spares** options.
2. Click **Apply Changes**. To exit and cancel your changes, click **Go Back to Previous Page**.

Related links

- [Managing The Physical Disk Power](#)

Managing Physical Disk Power Using The Customized Power Savings Mode

To manage physical disk power through the **Customized Power Savings** Mode:

1. Select the **Customized Power Save Mode** option.
2. Edit the remaining parameters on the **Manage Physical Disk Power** page. You can also configure the options in the QoS section as described in the following section.

Related links

[Managing The Physical Disk Power](#)

Managing Physical Disk Power Using The QoS Option

To manage physical disk power through the QoS option:

1. Select the **Customized Power Save Mode** option.
2. In the **Spin Down Configured Drives** drop-down menu, select **Enabled**.
3. The **Quality of Service (QoS)** option is enabled.
Enter the **Start Time** and **Time Interval** to spin up.
4. Click **Apply**.

Related links

[Managing The Physical Disk Power](#)

Managing The Time Interval For The QoS Option

To manage the time interval for the QoS option at the virtual disk level:

1. In the **Quality Of Service (QoS)** page, select **Enable Quality of Service Settings**.
2. Set the Start Time.
The start time can range from 1 to 24 hours.
3. Click **Apply Changes**.

 **NOTE: The Enable Quality of Service Settings option is enabled only if the Spin Down Configured Drives option is enabled.**

Related links

[Managing The Physical Disk Power](#)

To Locate Manage Physical Disk Power In Storage Management

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Select a controller object.
3. Click **Information/Configuration** on the controller **Properties** page.
4. Select **Manage Physical Disk Power** from the **Available Tasks** drop-down menu.
5. Click **Execute**.

Managing The Preserved Cache


The **Managed Preserved Cache** feature provides you the option to ignore or restore the controller cache data.

In the write-back policy, data is written to the cache before being written to the physical disk. If the virtual disk goes offline or is deleted for any reason, the data in the cache is lost.

Data in the cache may also be lost in case of unintended cable or power failure. If such a failure, Storage Management preserves the data written on the preserved or dirty cache until you recover the virtual disk or clear the cache.

This feature is available only on SAS controllers with firmware version 6.1 and later.

The status of the controller is affected by the preserved cache. The controller status is displayed as degraded if the controller has preserved cache.

 **CAUTION: You may not be able to use Storage Management to manage the preserved cache in some cases. As an example, consider you have a RAID 1 level with two disks — D1 and D2. If you now remove D2, the virtual disk is degraded and the data in the controller cache is written to D1. So, D1 has the latest data. Now, if you re-insert D2 and pull out D1, the virtual disk is still degraded and does not have the latest data.**

You can discard the preserved cache only if *all* of the following conditions are met:

- The controller does not have any foreign configuration. Select **Click for Preview** to view details of the foreign configuration. See [Foreign Configuration Operations](#).

- The controller does not have any offline or missing virtual disks. If there are offline or missing virtual disks, ensure that you have a backup of these virtual disks.
- Cables to any virtual disk are not disconnected.

Encryption Key

The controller uses the encryption key to lock or unlock access to SED. You can create only one encryption key for each encryption-capable controller.

If you are using LKM, you must create the encryption key by providing the **Encryption Key Identifier** and the **Passphrase**.

Encryption Key Identifier

An **Encryption Key Identifier** is a user-supplied text label for the **Passphrase**. The identifier helps you determine which **Passphrase** to provide while authenticating import of foreign encrypted SED drives.

Passphrase

A **Passphrase** is a user supplied string that the controller uses to create the encryption key.



NOTE: For more information on Encryption Key and Passphrase guidelines, click the  icon on the Manage Encryption Key page.

Related links

[Managing The Encryption Key](#)

Creating An Encryption Key And Enabling LKM

To create an encryption key on the selected controller:

1. Select the **Enable Local Key Management (LKM)** option.
2. Type **Encryption Key Identifier**.

An **Encryption Key Identifier** can contain numerals, alphabets both lower and upper case are allowed, non-alphanumeric characters, or a combination of any of these.



NOTE: For the Encryption Key Identifier and Passphrase guidelines, click the  icon on the page.

3. Type a **Passphrase**.

A **Passphrase** must contain at least one numeral, alphabets both lower and upper case are allowed, and one non-alphanumeric character (except space).



NOTE: Server Administrator Storage Management provides a suggested passphrase below the Passphrase text box.

4. If you want to save the Encryption Key credentials in a file on the system where Distributed Web Server is running, select the **Escrow** check box.

The **Path** field is displayed. Provide the path where you want to save the file. The path should contain a filename with an **.xml** extension. The saved file contains the information: SAS address, Encryption Key Identifier, Passphrase, and modified date. You can use this file for future reference.



CAUTION: It is important to understand that if you lose the Passphrase, you cannot recover it. If you move the physical disks associated with the lost Passphrase to another controller or if the controller fails or is replaced, you cannot access data from that disk.



NOTE: If Encryption Key Identifier or Passphrase contain special characters such as & , " , < , and > , in the file, they are written as & ; , " ; , < ; and > ; respectively.



NOTE: If the system crashes while saving the file, the backup file is saved in the specified location.

5. Select the check-box indicating that you understand the implications of using a passphrase and click **Apply Changes**.

In the controller **Information/Configuration** page, the **Encryption Key Present** is set to **Yes** and the **Encryption mode** is set to **LKM**.

Changing Or Deleting The Encryption Key

You can change an encryption key of a controller if the controller already has a configured encryption key. You can delete an encryption key for encrypted controllers only if there are no encrypted virtual disks.

To change the encryption key, type the **New Encryption Key Identifier** and **Passphrase**. You are prompted to authenticate the current **Passphrase**. Ensure that you read the note on the importance of passphrase and consequences of not saving the same, before applying the changes.




When you change the encryption key, the existing configuration on the controller is updated to use the new encryption key. If you have removed any of the encrypted drives previously, you must authenticate with the old passphrase to import the encrypted drives.

When changing the encryption key, you can also save or update the new encryption key credentials to a file in the system where Distributed Web Service is running. Select the **Escrow** check box. If you have already saved the encryption key credentials for a controller, providing the path of the file updates the credentials for that controller. If the credentials are for a new controller, the details are appended in the same file.

If you have not saved the credentials to a file, you can provide the path on which the file must be saved. The path must contain a filename with an **.xml** extension. On applying changes, this file is created with the credentials.

If you delete the encryption key, you cannot create encrypted virtual disks and all encrypted unconfigured self-encrypting drives are erased. However, deleting an encryption key does not affect encryption or data in foreign disks. If you have saved the encryption key credentials to a file, deleting the encryption key does not delete the file. Managing the file is the responsibility of the administrator.

Managing The Encryption Key

-  **NOTE: This task is not supported on PERC hardware controllers running in HBA mode.**
-  **NOTE: To configure encryption, SED is not required. The encryption settings are used to configure the virtual disk and the SED.**
-  **NOTE: On controller when encryption is disabled, manually enable encryption for virtual disks created using SED drives. Even if the virtual disk is created after a controller has had encryption enabled, to create an encrypted virtual disk the encryption option must still be selected from the Advanced Wizard during virtual disk creation.**

On an encryption-capable controller, the **Manage Encryption Key** task allows you to enable encryption in LKM mode. If you enable LKM, you can create an encryption key on an encryption-capable controller and save it locally. You can also change or delete the encryption key.

-  **NOTE: This task is available only on PERC H7x0 and H8x0 controllers.**

Related links

[Passphrase](#)

Manage Encryption Key Task In Storage Management

To go to the Manage Encryption Key task in Storage Management:

1. In the **Server Administrator** window, under the system tree, select **Storage**.
2. Go to **Storage Dashboard** → **Available Tasks** drop-down menu → **Manage Encryption Key...**
3. Click **Execute**.

Manage Encryption Key Task In Storage Management — Method 2

Alternatively to go to **Manage Encryption Key** task in Storage Management

1. Expand the **Storage** tree object to display the controller objects.
2. Select an encryption-capable controller object.
3. Click **Information/Configuration**.
4. Select **Manage Encryption Key....** from the **Controller Tasks** drop-down menu.
5. Click **Execute**.

If the controller is encryption-capable and an encryption key is not present, then the **Create Encryption Key** page is displayed. Else, the **Change or Delete Encryption Key** page is displayed.

Converting To Non-RAID Disks

On supported PERC adapters:

Does my controller support this feature? See [Supported Features](#).

 **NOTE: This task is not supported on PERC hardware controllers running in HBA mode.**

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Select the Controller object.
3. Click **Information/Configuration** on the controller **Properties** page.
4. In the **Controller Tasks** drop-down menu, select **Convert to Non-RAID**.
The disks in **Ready** state are displayed.
5. Select the drives that you want to convert.
6. Click **Apply**.

An acknowledgement that the disks have been converted is displayed.

Related links

[Converting To RAID Capable Disks](#)

Converting To RAID Capable Disks

On supported PERC adapters:

Does my controller support this feature? See [Supported Features](#).

 **NOTE: This task is not supported on PERC hardware controllers running in HBA mode.**

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Select the Controller object.
3. Click **Information/Configuration** on the controller **Properties** page.
4. In the **Controller Tasks** drop-down menu, select **Convert to RAID Capable disks**.
The non-RAID disks are displayed.
5. Select the drives that you want to convert.
6. Click **Apply**.

An acknowledgement that the disks have been converted is displayed.

 **NOTE: This operation is not supported on the PERC 10 controllers.**

Related links

[Converting To Non-RAID Disks](#)

Changing The Controller Mode

You can change the controller mode to Redundant Array of Independent Disks (RAID) or Host Bus Adapter (HBA) mode. To change the controller mode, perform the following steps:

 **NOTE: This option is supported only on PowerEdge RAID Controller 9 (PERC 9) and later family of hardware controllers.**

 **NOTE: When you change the controller mode from RAID to HBA or from HBA to RAID, certain features or functionality may vary.**

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Select the Controller object.
3. Click **Information/Configuration** on the controller **Properties** page.
4. In the **Controller Tasks** drop-down menu, select **Change Controller Mode...**, and then click **Execute**.

The **Change Controller Mode** window is displayed. The **Current Controller Mode** section displays the mode of the controller – **RAID** or **HBA**.

5. Select **RAID** or **HBA** from the **Change Controller Mode** section, and then click **Apply Changes**.
6. After you click **Apply Changes**, if you receive one of the following error messages, click **OK**, and then click **Return to Previous Page**. Disable or remove the options listed in the error message and then repeat steps 1 to 5.
 - **The controller is already running in the selected mode.** — Check the controller mode in the **Current Controller Mode** section before you proceed.
 - **Unable to change controller mode while foreign configuration is present on the controller.** — Remove all foreign configurations and then repeat the steps to change the controller mode.
 - **Unable to change controller mode while there is preserved cache present on the controller.** — Remove the preserved cache and then repeat the steps to change the controller mode.
 - **Unable to change controller mode while Virtual Disks are present on the controller.** — Remove all virtual disks and then repeat the steps to change the controller mode.
 - **Unable to change controller mode while Hotspares are present on the controller.** — Remove all hot spares and then repeat the steps to change the controller mode.
 - **Unable to change controller mode while security key is assigned to the controller.** — Remove the security key and then repeat the steps to change the controller mode.
7. Click **OK**, when you are prompted with the following message, **To apply the changes, you must restart the server. Are you sure, you want to change the controller mode?**
8. Restart the server to successfully change the controller mode.

Related Links

- [Controller Tasks](#)

Auto Configure RAID0 Operation

Does my controller support this feature? See [Supported Features](#).

You can use this feature to automatically configure all physical disks in **Ready** state into RAID 0 virtual disks. This feature is supported on PERC 9 family of hardware controllers with the latest firmware version.


 **NOTE: The Auto Configure RAID0 operation is only supported on SAS and SATA HDDs. If there are any SSDs available in the server, this feature ignores the SSDs when performing the Auto Configure RAID0 operation.**

 **NOTE: When performing an Auto Configure RAID0 operation on a single storage controller, the maximum number of physical disks supported is 192. Only physical disks in Ready state are configured into RAID 0 virtual disks.**

Before you proceed, browse to the physical disks details page and verify if the physical disks are in **Ready** state. To configure physical disks into RAID 0 virtual disks, perform the following steps:

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Select the Controller object.
3. Click **Information/Configuration** on the controller **Properties** page.
4. In the **Controller Tasks** drop-down menu, select **Auto Configure RAID0**, and then click **Execute**.
The **Auto Configure RAID0** window is displayed. Read the messages displayed on this window before you proceed.
5. Click **Confirm** to configure the physical disks.

The auto configure operation may take time to complete, depending on the configuration and available physical disks in **Ready** state. The **Auto Configure RAID0** operation is successful, only if the **Alert Log** is updated. You can check the **Alert Log** for more information.

 **NOTE: If there are no physical disks in Ready state and you repeat the procedure, the Auto Configure RAID0 operation automatically ignores your request without displaying any error message. For more information about alerts and their corrective actions, see the *Server Administrator Messages Reference Guide*. at dell.com/openmanagemanuals.**

System Configuration Lock Down Mode

'System Configuration Lock Down Mode' provide a configurable option to the customer so that after it is enabled, any configuration on that particular system is not permitted with some exceptions. Typical usage of this setting is after the system is provisioned to a

certain level along with the rest of the systems in the domain then enabling this mode reduces the drift the systems usually see during the operation over a period of time.

When the server is in “System configuration lock down” mode you will not be allowed to initiate any of the configuration operation from either OMSS GUI or CLI. If this mode is set, you are not allowed to perform most of the operations in OMSS except some of the enumeration or reporting operations.

You can enable the lock down mode from the iDRAC GUI. For enabling the lock down mode, you have to select the **Turn on the System Lockdown mode** option from the **More Actions** drop-down menu on the iDRAC page. It may take up to one minute to reflect the changes on the OMSS GUI.

 **NOTE: OMSS does not support configuring or setting the lock down mode from either GUI or CLI. For enabling the system configuration lock down mode, you have to access iDRAC GUI or similar applications.**

 **NOTE: If you restart the data manager/Engine service after setting the system lock down mode on, then the immediate reflection on the OMSS page happens. You might have to refresh the OMSS GUI after the restart, if required.**

OMSS CLI: Only the tasks available or supported in the GUI will be available or supported in the OMSS CLI, when the system is in “System configuration lock down” mode. If you try to run any of the unsupported operation, the following error message is displayed.

`'Error! System Configuration Lockdown mode is turned ON. Configuration actions cannot be performed in this mode. Operation failed!!!'`

Global Tasks Supported When System Configuration Lock Down Mode Is On

Table 16. Global Tasks Supported When System Configuration Lock Down Mode Is On

Global Tasks Name	PERC H730P Adapter/Mini Monolithic	PERC H740P Adapter/Mini Monolithic	PERC H730 Adapter/Mini Monolithic	PERC H330 Adapter/Mini Monolithic	PERC H840 Adapter	PERC HBA 330 Adapter/Mini	PERC 12 GBPS SAS HBA	PERC S140
Check Alert Log	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Setting Hot Spare Protection Policy	No	No	No	No	No	No	No	No
Setting RRWE Threshold	No	No	No	No	No	No	No	No
Setting Available Spare Threshold	No	No	No	No	No	No	No	No

Controller Tasks Supported When System Configuration Lock Down Mode Is On

Table 17. Controller tasks Supported When System Configuration Lock Down Mode Is On

Controller Tasks Name	PERC H730P Adapter/Mini Monolithic	PERC H740P Adapter/Mini Monolithic	PERC H730 Adapter/Mini Monolithic	PERC H330 Adapter/Mini Monolithic	PERC H840 Adapter	PERC HBA 330 Adapter/Mini	PERC 12 GBPS SAS HBA	PERC S140	PCIe SSDs (NVMe)
Export Log	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No
Start Patrol Read	No	Yes	Yes	Yes	Yes	Yes	No	No	No
View Slot Occupancy Report	No	No	No	No	No	No	No	No	Yes

 **NOTE:** If the Start Patrol Read task is supported by the controller, it will be supported in the 'system configuration lock down' mode too.

Controller Reports Supported When System Configuration Lock Down Mode Is On

Table 18. Controller Reports Supported When System Configuration Lock Down Mode Is On

Controller Reports Name	PERC H730P Adapter/Mini Monolithic	PERC H740P Adapter/Mini Monolithic	PERC H730 Adapter/Mini Monolithic	PERC H330 Adapter/Mini Monolithic	PERC H840 Adapter	PERC HBA 330 Adapter/Mini	PERC 12 GBPS SAS HBA
View Patrol Read Report	Yes	Yes	Yes	Yes	Yes	No	No
View Check Consistency Report	Yes	Yes	Yes	Yes	Yes	No	No
View Slot Occupancy Report	Yes	Yes	Yes	Yes	Yes	Yes	Yes
View Physical Disk Firmware Version Report	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Physical Disk Tasks Supported When System Configuration Lock Down Mode Is On

Table 19. Physical Disk Tasks Supported When System Configuration Lock Down Mode Is On

Physical Disk Tasks Name	PERC H730P Adapter/Mini Monolithic	PERC H740P Adapter/Mini Monolithic	PERC H730 Adapter/Mini Monolithic	PERC H330 Adapter/Mini Monolithic	PERC H840 Adapter	PERC HBA 330 Adapter/Mini	PERC 12 GBPS SAS HBA	PERC S140	PCIe SSDs (NVMe)
Blink	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Unblink	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Export Log	No	No	No	No	No	No	No	Yes	Yes

Virtual Disk Tasks Supported When System Configuration Lock Down Mode Is On

Table 20. Virtual Disk Tasks Supported When System Configuration Lock Down Mode Is On

Virtual Disk Tasks Name	PERC H730P Adapter/Mini Monolithic	PERC H740P Adapter/Mini Monolithic	PERC H730 Adapter/Mini Monolithic	PERC H330 Adapter/Mini Monolithic	PERC H840 Adapter	PERC HBA 330 Adapter/Mini	PERC 12 GBPS SAS HBA	PERC S140
Check Consistency	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Blink	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Unblink	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

 **NOTE:** Depending on the RAID level, the Check Consistency task is allowed for different PERC controllers.

Viewing Available Reports

To view a report:

1. Expand the **Storage** tree object to display the controller objects.
2. Select a controller object.
3. Click **Information/Configuration**.
4. Select a report from the **Select Report** drop-down menu.
5. Click **Execute**.

Available Reports

- [Viewing Patrol Read Report](#)
- [Viewing Check Consistency Report](#)
- [Viewing Slot Occupancy Report](#)
- [Viewing Physical Disk Firmware Version Report](#)

Viewing Patrol Read Report

Does my controller support this feature? See [Supported Features](#).

The patrol read report provides information on all the patrol reads performed on the controller in the chronological order. It provides information such as last run time and result. If a patrol read fails, the reason for the failure is displayed.

Related links

[Setting The Patrol Read Mode](#)

To Locate View Patrol Read Report In Storage Management

1. In the **Server Administrator** window, under the system tree, click **Storage**.
2. Select **View Patrol Read Report** from the **Select Report** drop-down menu.
3. Click **Execute**.

Viewing Check Consistency Report

Does my controller support this feature? See [Supported Features](#).

The check consistency report provides information on all the consistency checks performed on the controller in a chronological order. It provides information such as last run time and result. If the consistency check fails, it provides the reason for the failure.

Related links

[Performing A Check Consistency](#)

To Locate View Check Consistency Report In Storage Management

1. In the **Server Administrator** window, under the system tree, click **Storage**.
2. Select **View Check Consistency Report** from the **Select Report** drop-down menu.
3. Click **Execute**.

Viewing Slot Occupancy Report

Does my controller support this feature? See [Supported Features](#).

The **View Slot Occupancy Report** allows you to view empty and occupied slot details of all enclosures and back planes. It provides a diagram that represents the occupancy of physical drive slots. Move the mouse over each slot to view details, such as physical disk ID, state, and size.

Viewing Physical Disk Firmware Version Report

 **NOTE: This option is not supported on PERC hardware controllers running in HBA mode.**

Does my controller support this feature? See [Supported Features](#)

The Physical Disk Firmware Version Report compares the current firmware against a list of currently available firmware and legacy driver models.

 **NOTE: To generate the HDD firmware version report, negotiable speed and the model number of the drives are used as keys for indexing entries in the hddfwwer.csv file. If the negotiable speed of the drive is not available from the controller, then the model number of the drive is used as the key for indexing entries in the hddfwwer.csv file.**

You can run the report on each controller basis or for the storage system.

For each controller report, click **Storage** → **Controller** → **Information/Configuration** → **Available Reports** → **View Physical Disk Firmware Version Report** → **Execute**.

For a storage system report, select **Storage** → **Information/Configuration** → **Global Tasks** → **View Physical Disk Firmware Version Report** → **Execute**.

If you do not have the latest compare file (**hddfwwer.csv**), contact your service provider to download the latest compare file. Replace the existing **hddfwwer.csv** file with the new file at the following location:

On systems running Windows:

```
C:\<Program Files (x86)>\Dell\SysMgt\sm
```

where C:\Program files may vary based on the system.

On systems running Linux:

```
/opt/dell/srvadmin/etc/srvadmin-storage/hddfver.csv
```

On systems running ESXi:

```
/etc/cim/dell/srvadmin/srvadmin-storage/hddfver.csv
```

If the existing firmware(s) for all physical disks is the latest, the following message is displayed:

```
There are no physical disks available that require firmware update.
```

Related links

[Physical Disk Firmware Version Report Properties](#)

Physical Disk Firmware Version Report Properties

The report displays information for the drives that require a firmware upgrade as listed in the table below:

Table 21. Physical Disk Firmware Version Report Properties

Property	Definition
Name	Displays the nexus or location of each drive that has to be updated. This nexus is represented as a two or three-digit mapping of the drive location. Example: Two-digit mapping: 0:1 = Controller 0; Slot 1 Three-digit mapping: 1:0:4 = Controller 1; Connector 0; Slot 4
Model Number	Displays the unique number associated with a specific OEM vendor's drives and drive capacity.
Firmware Version	Displays the current running version of the firmware on the specific drive in the system.
Latest Available Firmware Version	Displays the firmware version that is compared against the firmware version in the comparison file.
Nautilus EFI	Nautilus is the tool that is used for offline firmware updates. Nautilus EFI is the version of the tool that works on 11th generation of supported servers. This tool updates multiple drive types with a single scan and updates procedure boot and runs from a USB key. If the Nautilus EFI column has a part number, then that drive is shipped on an 11th generation server. When downloaded, this tool appears under the Drive Firmware Downloads with a filename of the format NautilusEFIA_{xx}_ZPE.exe .
Nautilus DOS	Nautilus is the tool that is used for offline firmware updates. Nautilus DOS is the version of the tool that works on 9th — 11th generation of supported servers with SAS and SATA drives. This tool updates multiple drive types with a single scan and updates procedure boot and runs from a USB key, preboot execution environment (PXE), or CD-ROM. If the Nautilus DOS column has a part number, then that drive is shipped on a 9th — 11th generation server. When downloaded, this tool appears under the Drive Firmware Downloads with a filename of the format NautilusA_{xx}_ZPE.exe .
DUP Reboot Required	If this field is set to Yes , then the Update Package (DUP) field is not blank. It indicates the availability of an online DUP. The DUP allows to be sent to the firmware payload through an online executable, but the firmware will not be committed to the disk until the next system reboot. Hence, you can perform one-to-many online deployments using applications or scripts that can launch the online executable.
DUP	It is a single executable that runs on a single family of drives. Unlike Nautilus, to update different drives you have to use different DUP packages. A single DUP package updates all drives applicable to that DUP package in an execution. You can run the DUP online without a reboot. It is recommended to stop or at least slow I/O operations during a DUP online firmware upgrade.

Property	Definition
Part Number	In the event of a drive failure, you can run the View Physical Disk Firmware Version Report to find out the part number of the failed drive and to check if any of the drives require an update.

Support For PERC 9 and PERC 10 Hardware Controllers

The PowerEdge RAID Controller (PERC) family of enterprise-class controllers is designed for enhanced performance, increased reliability and fault tolerance, and simplified management — providing a powerful, easy-to-manage way to create a robust infrastructure and help maximize server uptime. The introduction of the PERC 9 and PERC 10 family of hardware controllers also brings about improvements in storage solutions.

The new PERC 9 and PERC 10 family of hardware controllers support the following storage enhancements:


- [Support For RAID Level 10 Virtual Disk On PERC 9 and PERC 10 Hardware Controller](#)
- [Support For Advanced Format 4KB Sector Hard-Disk Drives](#)

 **NOTE: The order of the controllers displayed on Storage Management may differ with the order of the controllers displayed in the Human Interface (HII) and PERC Option ROM. The order of the controllers does not cause any limitation.**

Support For RAID Level 10 Virtual Disk Creation On PERC 9 and PERC 10 Hardware Controllers

RAID Level 10 is a solution for users who require high performance and redundancy with quickest recovery during drive failure. Though a RAID Level 10 setup is more expensive to maintain, it has several benefits as it combines the properties of RAID Level 1 and RAID Level 0.

RAID Level 10 virtual disk creations with PERC 9 and PERC 10 hardware controllers support Uneven Span feature. When you create RAID Level 10 virtual disks with PERC 9 and PERC 10 hardware controllers the firmware suggests the preferred span layout for the setup.

 **NOTE: A minimum of 4 physical disks and a maximum of 256 physical disks are allowed for a RAID Level 10 virtual disk setup with PERC 9 and PERC 10 hardware controllers.**

You can create RAID Level 10 Virtual Disk on PERC 9 and PERC 10 hardware controllers using the following Wizards:

- **Express Wizard**
- **Advanced Wizard**

 **NOTE: The procedure for creating virtual disks on PERC 9 hardware controllers is the same as other PERC hardware controllers.**

Related Tasks

- [Create Virtual Disk Express Wizard](#)
- [Create Virtual Disk Advanced Wizard](#)

RAID Level 10 Virtual Disk Creation With Uneven Span

RAID Level 10 virtual disk creation with uneven span feature is available on Storage Management User Interface (UI) and Command Line Interface (CLI). For information on Storage Management CLI, see *Server Administrator Command Line Interface Guide*.

- Based on the minimum (and even) number of physical disks selection the firmware on the PERC 9 hardware controller recommends the preferred span layout.

 **NOTE: RAID Level 10 virtual disk creation command from Storage Management CLI does not support optional parameter [spanlength=<n>] on PERC 9 hardware controllers.**

- The **Advanced Wizard** does not provide the option to select span length for RAID Level 10 virtual disk creation on PERC 9 hardware controllers.
- The span layout for RAID Level 10 virtual disk created from the **Express Wizard** on PERC 9 hardware controllers uses the span layout as recommended by the firmware of the PERC 9 hardware controller.

 **NOTE: Storage Management uses the PERC 9 hardware controller-firmware suggested span layout for RAID Level 10 virtual disk creation.**

- The PERC 9 hardware controller firmware suggested span layout is same for the same set of physical disks.
- **Intelligent Mirroring** feature is supported on PERC 9 hardware controllers.
- When you create a virtual disk using the **Advanced Wizard** on PERC 9 hardware controllers the span layout information under **Selected Physical Disks** will not be displayed.
- Sliced RAID Level 10 virtual disk creation on PERC 9 hardware controllers support uneven span.
- Grouping of selected physical disks for RAID Level 10 virtual disk through the **Advanced Wizard** is not supported on PERC 9 hardware controllers.
- When you import a foreign configuration from hardware controllers prior to PERC 9 hardware controllers to PERC 9 hardware controllers, the span layout for RAID Level 10 virtual disk remains the same.
- When you import a foreign configuration for RAID Level 10 virtual disks from PERC 9 hardware controllers to other PERC 9 hardware controllers the span layout does not change.

 **NOTE: Importing a foreign configuration (other than RAID Level 10 virtual disks) from PERC 9 hardware controllers to hardware controllers prior to PERC 9 is not supported.**


Related Tasks

- [Create Virtual Disk Express Wizard](#)
- [Create Virtual Disk Advanced Wizard](#)

Support For Advanced Format 4KB Sector Hard-Disk Drives

From the introduction of PERC 9 family of hardware controller support in Storage Management, users using legacy 512B sector hard-disk drives can now transition to Advanced Format 4KB sector hard-disk drives. 4KB sector hard-disk drives utilize the storage-surface media more efficiently by combining data that would have been stored in eight 512B sector hard-disk drives into one sector of 4096B (4KB). This data-combining feature in 4KB sector hard-disk drives results in improved data efficiency and error correction capabilities.

Storage Management supports virtual disk creation on 4KB sector hard-disk drives connected to PERC 9 hardware controllers.

 **NOTE: 4KB sector hard-disk drives are not supported on controllers prior to PERC 9 family of hardware controllers. If the 4KB sector hard-disk drive is connected to any hardware controller prior to PERC 9, the 4KB sector hard-disk drive is displayed as Unsupported.**

- When you create a virtual disk using the **Advanced Wizard**, you can select the physical disk sector from the **Sector Size** drop-down list. The options available are:
 - **512B**
 - **4KB**
- You cannot use both 4KB sector hard-disk drives and 512B sector hard-disk drives for a virtual disk creation as mixing of hard-disk drive sectors is not allowed in Storage Management.

 **NOTE: If the system contains any 512e sector hard-disk drives, the 512e sector hard-disk drives are identified/ reported as 512B sector hard-disk drives and follows the behavior of the 512B sector hard-disk drives.**

Hot Spare Considerations — 4KB Sector Hard-Disk Drives

The following is the Hot Spare considerations (Dedicated and Global Hot Spare) for 4KB sector hard-disk drives supported on PERC 9 hardware controllers:

- You cannot assign 4KB sector hard-disk drives as Dedicated Hot Spare for virtual disks created with 512B sector hard-disk drives and vice versa.
- You cannot assign 4KB sector hard-disk drives as Global Hot Spare, if the virtual disk created consists of 512B sector hard-disk drives only and vice versa.
- You can assign 4KB sector hard-disk drives as Global Hot Spare, if there are 4KB sector hard-disk drives and 512B sector hard-disk drives in the virtual disks created and vice versa.

 **NOTE: If you perform this action a warning message is displayed.**

Related Tasks

- [Create Virtual Disk Advanced Wizard](#)

Reconfiguration Considerations — 4KB Sector Hard-Disk Drives

You cannot reconfigure 4KB sector hard-disk drives with virtual disks that consist of 512B sector hard-disk drives and vice versa.

Support for Marvell RAID Controllers

All Operating Systems compatible to OM 9.0.1 are supported by Marvell RAID Controllers except ESXi family of OSes.

The Marvell RAID Controllers supports the following enumeration and monitoring operations:

- The physical disks are (M.2 devices) are directly connected to the Controller
- Enumeration of Physical Disks (M.2 devices) are supported
- Enumeration of Virtual Disks on M.2 devices are supported

The following tasks are not supported:

- Configuration options are not supported for this controller.
- Enclosures and connectors are not applicable for the Controller.
- No physical disks operations / tasks are supported for this Controller
- No configuration operations are supported for virtual disks including creation, deletion, reconfiguration and so on.
- No controller tasks are supported

The following physical disks properties are enumerated for this controllers: ID, Status, Name, State, Bus Protocol, Media, Revision, Model Number, Capacity, Used RAID Disk Space, Available RAID Disk Space, Hot Spare, Vendor ID, Product ID, Serial No., Negotiated Speed, Capable Speed, Sector Size and SAS Address.

 **NOTE: SAS Address is not available for M.2 devices.**

The following controller properties are enumerated for this controllers: ID, Status, Name, Slot ID, State, Firmware Version, Rebuild Rate, BGI Rate, Check Consistency Rate, Cache Memory Size, and Patrol Read Rate.

The following virtual disk properties are enumerated for this controllers: Status, Name, State, Tasks, Layout, size, Device Name, Bus Protocol, Media, Read Policy, Write Policy, Stripe Element Size, and Disk Cache Policy.

 **NOTE: Layout option will not be available for non-RAID disks.**

 **NOTE: For Marvell RAID Controller, the physical disk is directly connected to the controller as no enclosure or connector is applicable for this controller.**

 **NOTE: Any configuration operations through OMSA are not supported for virtual disks, physical disks and controllers.**

Enclosures And Backplanes

Physical disks can be contained in an enclosure or attached to the backplane of a system. An enclosure is attached to the system externally while the backplane and its physical disks are internal.

Related links

[Enclosures](#)

[Backplanes](#)



Backplanes

You can view the **Backplane** object by expanding the controller in the Storage Management tree view. Storage Management displays the status of the backplane and the attached physical disks. Backplane is similar to an enclosure. In a Backplane, the controller connector and physical disks are attached to the enclosure, but it does not have the management features (temperature probes, alarms, and so on) associated with external enclosures.

Flexible Backplane Zoning

Flexible backplane zoning is a functionality introduced in Storage Management that allows you to connect two PERC hardware controllers to the backplane or internal drive-array by using a single expander. This configuration enables Storage Management to split the backplane between the two PERC hardware controllers, as a result of that increasing the performance of the system. When flexible backplane zoning is enabled, the backplanes display the same backplane ID for all backplanes connected to the two PERC hardware controllers. In flexible backplane zoning, physical disks and virtual disks connected to the first controller is not displayed on the second controller and the other way around. For example, If you have to create a virtual disk using the first controller, only the physical disks connected to the first controller is enumerated and available for the operation. The same is applicable when you view the **Slot Occupancy Report** for a particular controller.

The Flexible backplane zoning feature is only supported on backplanes of 24 slots—PowerEdge R630 and R730xd. If the backplane consists of 26 slots, in the case of PowerEdge R730xd, the additional two slots located next to the rear ports are not considered in this configuration.

-  **NOTE: Flexible backplane zoning can be configured only through RACADM and not through Storage Management.**
-  **NOTE: Flexible backplane zoning is supported only on PERC (internal) family of controllers—PERC H730 Adapter, PERC H730 Mini, HBA 330, and HBA 330 Mini. This feature is not supported on PERC H330 Adapter and PERC H330 Mini.**

Enclosures

Storage Management enables management of various enclosures and their components. In addition to managing the physical disks contained in the enclosure, you can monitor the status of the fans, power supply, and temperature probes in an enclosure. You can view these components by expanding the controller, connector, and enclosure objects in the Storage Management tree view.

Storage management allows hot plugging of enclosures. Hot plugging is defined as the adding of a component to a system while the operating system is still running.

-  **NOTE: This feature requires that the physical devices connected to the controller have the latest firmware. For the latest supported firmware, contact your service provider.**

After you hot plug or hot reconfigure an enclosure, refresh the left tree to display the change in status and configuration; a system reboot is not required.



NOTE: Storage Management does not allow hot removal of enclosures. Reboot the system to effect this change in Storage Management.

Storage Management notifies you of enclosure status changes through alerts that are displayed in the **Alert Log**.

The following sections provide more information on the enclosure components and management features provided by Storage Management:

- [Enclosure Physical Disks](#)
- [Enclosure Fans](#)
- [Enclosure Power Supplies](#)
- [Enclosure Temperature Probes](#)
- [Enclosure Management Modules \(EMMs\)](#)
- [Enclosure And Backplane Health](#)
- [Enclosure And Backplane Properties And Tasks](#)

Enclosure Physical Disks




The physical disks in an enclosure are displayed under the enclosure object in the tree view. Select a disk in the tree to view the status information.

Enclosure Fans

The fans are a component of the enclosure cooling module. The enclosure fans are displayed under the **Fans** object in the tree view. You can select a fan to view the status information.

Fan Properties

Table 22. Fan Properties

Property	Definition
Status	<p>These icons represent the severity or health of the storage component.</p> <p> — Normal/OK</p> <p> — Warning/Non-critical</p> <p> — Critical/Failure/Error</p> <p>For more information, see Storage Component Severity.</p>
Name	Displays the name of the fan.
State	<p>Displays the status of the fan. Possible values are:</p> <ul style="list-style-type: none"> • Ready — The fan is functioning normally. • Degraded — The fan has encountered a failure and is operating in a degraded state. • Offline — The fan or power supply has been removed from the enclosure. • Failed — The fan has encountered a failure and is no longer functioning. Storage Management is unable to communicate with the enclosure using SES commands. The Failed state is displayed when the enclosure does not respond to a status query from Storage Management. For example, disconnecting the cable. • Missing — The fan is not present in the enclosure.
Part Number	<p>Displays the part number of the fan.</p> <p>This property is not displayed for the 22xS enclosures that have E.17 firmware or later.</p>
Speed	Indicates the fan speed. The values are displayed in rpm . If the fan is in Offline state, the value for the Speed property is Unknown .




Property	Definition
	For information on events that cause the fan speed to change, see the hardware documentation.

Enclosure Power Supplies

The enclosure power supplies are displayed under the **Power Supplies** object in the tree view. You can select the **Power Supplies** object to view the status information.

Power Supply Properties

Table 23. Power Supply Properties

Property	Definition
Status	<p>These icons represent the severity or health of the storage component.</p> <p> — Normal/OK</p> <p> — Warning/Non-critical</p> <p> — Critical/Failure/Error</p> <p>For more information, see Storage Component Severity.</p>
Name	Displays the name of the power supply.
State	<p>Displays the status of the power supply.</p> <ul style="list-style-type: none"> • Ready — The power supply is functioning normally. • Degraded — The power supply has encountered a failure and is operating in a degraded state. • Failed — The power supply has encountered a failure and is no longer functioning. Storage Management is unable to communicate with the enclosure using SES commands. The Failed state is displayed when the enclosure does not respond to a status query from Storage Management for any reason. For example, This state is displayed when you disconnect the cable. • Missing — The power supply is not present in the enclosure.
Part Number	<p>Displays the part number of the power supply.</p> <p>This property is not displayed for the 22xS enclosures that have E.17 firmware or later.</p>
Firmware Version	Displays the firmware version number of the power supply.

Enclosure Temperature Probes

The enclosure temperature probes are displayed under the **Temperatures** object. You can select the **Temperatures** object to view the status information. The status information includes the current temperature in Celsius and the warning and failure thresholds for the temperature probe.

The failure threshold has a default value that cannot be changed. However, you can set the warning threshold.

Related links

[Temperature Probe Properties And Tasks](#)

[Checking The Temperature Of The Enclosure](#)

[Setting The Temperature Probe Values](#)

Setting The Temperature Probe Properties And Tasks

Click **Set Temperature Probe** to launch the wizard for changing the temperature probe's Warning threshold. You can change the Warning threshold for each of the temperature probes included in the enclosure.

To Launch The Set Temperature Probe Wizard





To launch the **Set Temperature Probe** wizard:

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Expand a controller object.
3. Expand the enclosure object.
4. Select the **Temperatures** object.
5. Click **Set Temperature Probe**.

Temperature Probe Properties And Tasks

Does my enclosure support this feature? See [Supported Features](#). Use this window to view information about the enclosure's temperature probes.

Table 24. Temperature Probe Properties

Property	Definition
Status	<p>These icons represent the severity or health of the storage component.</p> <ul style="list-style-type: none">  — Normal/OK  — Warning/Non-critical  — Critical/Failure/Error <p>For more information, see Storage Component Severity.</p> <p> NOTE: On some enclosures, Storage Management may experience a short delay before displaying the current enclosure temperature and temperature probe status. For more information, see Storage Management May Delay Before Updating Temperature Probe Status.</p>
Name	Displays the name of the temperature probe.
State	<p>Displays the status of the temperature probe.</p> <ul style="list-style-type: none"> • Ready — The temperature probe is functioning normally. • Degraded — The temperature probe has encountered a failure and is operating in a degraded state. • Failed — The temperature probe has encountered a failure and is no longer functioning. Storage Management is unable to communicate with the enclosure using SES commands. The Failed state is displayed when the enclosure does not respond to a status query from Storage Management for any reason. For example, disconnecting the cable causes a Failed state. • Minimum Warning Threshold Exceeded — The temperature of the enclosure has dropped below the minimum warning threshold. • Maximum Warning Threshold Exceeded — The temperature of the enclosure has exceeded above the maximum warning threshold. • Missing — The temperature probe is not present in the enclosure. • Inactive — The temperature probe is present in the enclosure, but the EMM that it monitors is not installed.
Reading	Displays the current temperature of the enclosure as reported by the temperature probe.
Warning Threshold	The Minimum and Maximum properties indicate the temperatures currently set for the Warning threshold.

Property	Definition
Failure Threshold	The Minimum and Maximum properties indicate the temperatures currently set for the Failure threshold.

Related links

- [Enclosure Temperature Probes](#)
- [Checking The Temperature Of The Enclosure](#)
- [Setting The Temperature Probe Values](#)

Enclosure Management Modules (EMMs)

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 view the individual EMM modules and their status information.

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 physical disk
- The LEDs on the enclosure

When the alarm on the enclosure 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 [Enabling The Enclosure Alarm](#). For more information on EMMs, see the enclosure hardware documentation.

All EMM modules in the enclosure should have the same version of firmware. You can view the properties of each individual EMM module to verify the firmware version.

Verifying The EMM Firmware Version Of The Enclosure

Does my controller support this feature? See Supported Features.





The firmware for the EMM must be at the same level as that of the enclosure. The status of the EMMs is displayed as degraded if there is a mismatch between the EMM firmware.


To verify the EMM firmware version:

1. In the **Server Administrator** window, under the system tree, click **Storage** dashboard.
2. Expand the tree view until the **EMMs** object is displayed.
3. Select the **EMMs** object. The firmware version for each EMM is displayed in the **Firmware Version** column in the right pane. For information related to the enclosure's EMMs, see Enclosure Management Modules (EMMs).

EMM Properties

Table 25. EMM Properties

Property	Definition
Status	<p>These icons represent the severity or health of the storage component.</p> <ul style="list-style-type: none">  — Normal/OK  — Warning/Non-critical  — Critical/Failure/Error  — Unknown <p>For more information, see Storage Component Severity.</p>
Name	Displays the name of the EMM.

Property	Definition
State	<p>Displays the current state of the EMMs. Possible values are:</p> <ul style="list-style-type: none"> • Ready — EMM is functioning normally. • Degraded — EMM has encountered a failure and is operating in a degraded state. • Failed — EMM has encountered a failure and is no longer functioning. Storage Management is unable to communicate with the enclosure using SES commands. The Failed state is displayed when the enclosure does not respond to a status query from Storage Management. For example, disconnecting the cable. • Missing — EMM is not present in the enclosure. • Not Installed — EMM is not present in the enclosure.
Part Number	Displays the part number of the EMM module.
Type	<p>Indicates whether the EMM is a SCSI SES Module or a SCSI Terminator.</p> <p>SCSI SES Module — A SCSI SES module provides SCSI Enclosure Services (SES) and SCSI Accessed Fault-Tolerant Enclosures (SAFTE) reporting to the host server, control of all system LED indicators, and monitoring of all environmental elements, such as temperature sensors, cooling modules, and power supplies.</p> <p>SCSI Terminator — The SCSI terminator card is only used if the 220S or 221S enclosure is not configured with a redundant SCSI SES module type of EMM. In systems equipped with two SCSI SES modules, the SCSI termination is done through the EMMs.</p>
Firmware Version	<p>Indicates the version of the firmware loaded on the EMM. All EMM modules in the enclosure should have the same level of firmware.</p> <p> NOTE: For multiple backplanes, the firmware version appears as upstream and downstream versions.</p>
SCSI Rate	Displays the maximum SCSI speed that the EMM in a SCSI enclosure supports.

Enclosure And Backplane Health

Displays the status of the enclosure or backplane and the components attached to the enclosure or backplane.

Enclosure And Backplane Information

For information on enclosures and backplanes, see:

- [Enclosures And Backplanes](#)
- [Enclosure And Backplane Properties](#)
- [Enclosure And Backplane Tasks](#)

Enclosure And Backplane Components

For information on attached components, see [Physical Disks Or Physical Devices](#).

Enclosure And Backplane Properties And Tasks





You can view information about the enclosure or backplane and execute enclosure tasks.


Related links

- [Enclosure And Backplane Properties](#)
- [Enclosure And Backplane Tasks](#)

Enclosure And Backplane Properties

Table 26. Enclosure and Backplane Properties

Property	Definition
ID	Displays the ID of the enclosure or backplane.
Status	<p>These icons represent the severity or health of the storage component. For more information, see Storage Component Severity.</p> <p> — Normal/OK</p> <p> — Warning/Non-critical</p> <p> — Critical/Failure/Error</p> <p> NOTE: If the enclosure is connected to the controller in redundant path mode (for more information, see Setting The Redundant Path Configuration), the loss of connection to any one EMM may cause the enclosure status to be displayed as degraded.</p>
Name	Displays the name of the enclosure or backplane.
State	<p>Displays the status of the enclosure or backplane. Possible values are:</p> <ul style="list-style-type: none"> • Ready — The enclosure or backplane is functioning normally. • Degraded — The enclosure has encountered a failure and is operating in a degraded state. This state does not apply to backplanes. • Failed — The enclosure or backplane has encountered a failure and is no longer functioning.
Connector	Displays the number of connectors attached to the enclosures or backplanes. This number matches the number connector on the controller hardware. Depending on the controller type, the connector can be either a SCSI channel or a SAS port.
Enclosure ID	Displays the enclosure ID assigned to the enclosure by Storage Management. Storage Management numbers the enclosures attached to the controller starting with zero. This number is the same as the enclosure ID number reported by the omreport command. For information on Command Line Interface, see the <i>Server Administrator Command Line Interface User's Guide</i> .
Target ID	Displays the SCSI ID of the backplane (internal to the server) or the enclosure to which the controller connector is attached. The default value is six.
Configuration	<p>Displays the mode in which the enclosure is operating. Possible values are:</p> <ul style="list-style-type: none"> • Joined — Indicates that the enclosure is operating in the joined-bus mode. • Split — Indicates that the enclosure is operating in the split-bus mode. • Unified — Indicates that enclosure is operating in the unified mode. • Clustered — Indicates that the enclosure is operating in cluster mode. Clustered mode is only available on cluster-enabled RAID controllers. For more information, see Background Initialization On PERC Controllers. <p>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 hardware documentation.</p>

Property	Definition
	The bus configuration switch when toggled on a 220S or 221S enclosure, the enclosure should be turned off. For more information, see Changing The Mode On 220S And 221S Enclosures . Older 200S enclosures with a version 1.8 kernel that only have one EMM may display Split Bus mode and not Joined Bus mode. Joined Bus or Clustered modes may display when applicable in these circumstances.
Firmware Version	Displays the version of the enclosure firmware.  NOTE: For systems supporting multiple backplanes, the firmware version appears as upstream and downstream versions.
Service Tag	Displays the enclosure Service Tag. The enclosure is identified by a unique Service Tag number. This information is used to route support calls to the appropriate personnel.
Express Service Code	The enclosure is identified by a unique Express Service Code number. This information is used to route support calls to the appropriate personnel.
Asset Tag	Displays the asset tag information for the enclosure. You can change this property using the Setting Asset Data task.
Asset Name	Displays the name assigned to the enclosure. You can change this property using the Setting Asset Data task.
Backplane Part Number	Displays the part number of the enclosure backplane.
SAS Address	Displays the SAS address of the SAS backplane.
Split Bus Part Number	Displays the part number of the enclosure split bus module. A split bus is indicated by a single triangle symbol on the back of the enclosure.
Enclosure Part Number	Displays the part number of the enclosure.
Enclosure Alarm	Displays whether the alarm on the enclosure is enabled or disabled.

Enclosure And Backplane Tasks

To execute a drop-down menu enclosure task:

1. In the **Server Administrator** window, under the system tree, expand **Storage** dashboard to display the controller objects.
2. Expand a controller object.
3. Select the Enclosure object.
4. Click **Information/Configuration** on the storage Properties screen.
5. Select a task from the **Enclosure Tasks** drop-down menu.
6. Click **Execute**.

Enclosure — Available Tasks

The drop-down menu enclosure tasks are:

- [Enabling The Enclosure Alarm](#)
- [Disabling The Enclosure Alarm](#)
- [Setting Asset Data](#)
- [Blinking The LED On The Enclosure](#)
- [Setting The Temperature Probe Values](#)

Enabling The Enclosure Alarm

Does my enclosure support this feature? See [Supported Features](#).

Select the **Enable Alarm** task to enable the enclosure alarm. When enabled, the alarm turns on when any of the following events occur:

- Enclosure temperature has exceeded the warning threshold.
- Power supply, fan, or enclosure management module — EMM has failed.
- Split bus is not installed. A split bus is indicated by a single triangle symbol on the back of the enclosure.

Related links

[Using Alarms To Detect Failures](#)

Disabling The Enclosure Alarm

Does my enclosure support this feature? See [Supported Features](#).

Select the **Disable Alarm** task to disable the enclosure alarm. The alarm is turned off 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 turned on, you can turn it off with this task.

Setting Asset Data

Does my enclosure support this feature? See [Supported Features](#).

You can change the asset tag and asset name of the enclosure. The new asset tag and asset name that you specify are displayed on the **Information/Configuration** of the enclosure.

Changing The Asset Tag And Asset Name Of An Enclosure

To change the asset tag and asset name of the enclosure:

1. Type the new asset tag name in the **New asset tag** field.
You can specify an inventory number or other useful information for your environment. The asset tag typically refers to the enclosure hardware.
2. Type the new asset name in the **New asset name** field.
You can specify a name useful to you for organizing your storage environment. For example, the asset name could refer to the type of data stored on the enclosure or to the enclosure's location.
3. Click **Apply Changes**.
To exit and cancel your changes, click **Go Back To Enclosure Information Page**.

Related links

[Setting Asset Data](#)

To Locate Set Asset Data In Storage Management

To locate this task in storage management:

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Expand a controller object.
3. Select the Enclosure object.
4. Click **Information/Configuration**.
5. Select **Set Asset Data** from the **Enclosure Tasks** drop-down menu.
6. Click **Execute**.

Blinking The LED On The Enclosure

Does my enclosure support this feature? See [Supported Features](#).

Select the **Blink** task to blink the LED on the enclosure. You may want to use this task to locate an enclosure. The LEDs on the enclosure may display different colors and blinking patterns. For more information on what the blink colors and patterns indicate, refer to the enclosure hardware documentation.

Setting The Temperature Probe Values

Does my enclosure support this feature? See [Supported Features](#).

The temperature probes monitor the temperature of the enclosure. Each temperature probe has a warning and a failure threshold. The warning threshold indicates that the enclosure is approaching an unacceptably warm or cool temperature. You can modify the warning threshold.

The failure threshold indicates that the temperature of the enclosure has gone below the minimum threshold or exceeded the maximum threshold which may cause data loss. You cannot change the default values for the failure threshold.

Related links

- [Enclosure Temperature Probes](#)
- [Temperature Probe Properties And Tasks](#)
- [Using Enclosure Temperature Probes](#)

Checking The Temperature Of The Enclosure

Does my controller support this feature? See [Supported Features](#).

To check the temperature:

1. In the **Server Administrator** window, under the system tree, click **Storage** dashboard.
2. Expand the tree view until the **Temperatures** object is displayed.
3. Select the **Temperatures** object. The temperature reported by the temperature probe is displayed in celsius in the **Reading** column in the right pane.

Related links

- [Enclosure Temperature Probes](#)
- [Temperature Probe Properties And Tasks](#)

Changing The Warning Threshold On The Temperature Probe

The temperature probes are listed in the **Temperature Probes** section. To change the temperature probe's warning threshold:

1. Select the probes that you want to change.
2. Select **Set new values** in **Set New Temperature Probe Values** page.
3. Type the lowest acceptable temperature in celsius for the enclosure in the **Minimum Warning Threshold** text box. The text box label indicates the range that you can specify.
4. Type the highest acceptable temperature in celsius for the enclosure in the **Maximum Warning Threshold** text box. The text box label indicates the range that you can specify.
5. Click **Apply Changes**.

To return to the default values, **Reset to default values** and click **Apply Changes**. The default values are displayed in the **Minimum Warning Threshold** and **Maximum Warning Threshold** text boxes.



NOTE: On some enclosures, Storage Management may experience a short delay before displaying the current enclosure temperature and temperature probe status.

Related links

- [Storage Management May Delay Before Updating Temperature Probe Status](#)

To Locate Set Temperature Probe Values In Storage Management

To locate this task in storage management:

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Expand a controller object.
3. Select the enclosure object.
4. Click **Information/Configuration**.
5. Select **Set Temperature Probe Values** from the **Enclosure Tasks** drop-down menu.
6. Click **Execute**.

Available Reports

- [View Slot Occupancy Report](#)
- [View Slot Occupancy Report](#)**



NOTE: This option is not supported on PERC hardware controllers running in HBA mode.

Does my controller support this feature? See [Supported Features](#).

The **View Occupancy Slot Report** task allows you to view empty slots, occupied slots, or slots split by the backplane zoning functionality of the selected enclosure. The report provides a diagram that represents the occupancy of physical drive slots. Move the mouse over each slot to view details, such as physical disk ID, state, and size.

For more information on flexible backplane zoning, see [Backplanes](#).

To Locate View Slot Occupancy Report In Storage Management

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Expand a controller object.
3. Select the enclosure object.
4. Click **Information/Configuration**.
5. Select **View Slot Occupancy Report** from the **Available Tasks** drop-down menu.
6. Click **Execute**.

If the backplane supports flexible backplane zoning, then only the slots, empty or occupied, of the controller you are viewing is displayed. The slots connected to the other controller are grayed-out and do not display any information. For more information on flexible backplane zoning, see [Backplanes](#)

Changing The Mode On 220S And 221S Enclosures

Make sure that you turn off the enclosure before you toggle the bus configuration switch on a 220S or 221S enclosure, the enclosure should be powered off. The bus configuration switch is used to change the enclosure to split bus, joined bus, or clustered mode. If you change the 220S or 221S enclosure mode with the enclosure turned on, the enclosure may no longer be displayed by Storage Management and you may notice erratic behavior. In addition, the bus configuration switch on these enclosures is not rated for frequent toggling.

Enclosure Management

In addition to the enclosure tasks, the following activities may be required for enclosure management:

- Identifying the service tag of the enclosure — To identify the service tag of the enclosure, select the enclosure in the tree view and click **Information/Configuration**. The **Information/Configuration** page displays the service tag and other enclosure properties.
- Identifying the Express Service Code of the enclosure — The Express Service Code is a numeric function of your service tag. You can key in the numeric Express Service Code for automated call-routing while calling technical support. To identify the Express Service Code of the enclosure, select the enclosure in the tree view and click the **Information/Configuration**. The **Information/Configuration** displays the Express Service Code and other enclosure properties.
- Preparing a physical disk for removal — To prepare a physical disk for removal is a physical disk command. See [Preparing To Remove](#).
- Troubleshooting — For more information on Troubleshooting, see [Troubleshooting](#).
- Removing the wrong physical disk — You can avoid removing the wrong physical disk by blinking the LED display of the disk you intend to remove. See [Blinking And Unblinking A Physical Disk](#).
- If you have already removed the wrong physical disk, see:
 - [Recovering From Removing The Wrong Physical Disk](#)
 - [Identifying An Open Connector On The Enclosure](#)
 - [Checking The Temperature Of The Enclosure](#)
 - [Verifying The EMM Firmware Version Of The Enclosure](#)

Identifying An Open Connector On The Enclosure

Does my controller support this feature? See [Supported Features](#).

If you have not attached the enclosure to an open connector, you may need to identify a connector on the controller that can be used for this purpose. To identify an open connector:

1. In the **Server Administrator** window, under the **System tree**, expand **Storage**.
2. Expand the controller object.
The available connectors are displayed. These connectors are numbered starting from zero.
3. Identify a connector that is not attached to storage. If the connector is already attached to storage it can be expanded to display an enclosure or backplane and the attached physical disks. A connector object that cannot be expanded in the tree view

is an open connector not currently attached to storage. Storage Management displays a number for each connector. These numbers correspond to the connector numbers on the controller hardware. You can use these numbers to identify which open connector displayed in the tree view with the open connector on the controller hardware.

Enclosure Components

For information on attached components, see:

- [Physical Disk Or Physical Device Properties](#)
- [EMM Properties](#)
- [Fan Properties](#)
- [Power Supply Properties](#)
- [Temperature Probe Properties And Tasks](#)

Connectors

A controller contains one or more connectors (channels or ports) to which you can attach disks. You can externally access a connector by attaching an enclosure (for external disks) to the system or internally access by attaching to the backplane (for internal disks) of a system. You can view the connectors on the controller by expanding the controller object in the tree view.

Channel Redundancy

You can create a virtual disk that uses physical disks that are attached to different controller channels. The physical disks may reside in an external enclosure or the backplane (internal enclosure). If the virtual disks maintain redundant data on different channels, then these virtual disks are channel redundant. Channel redundancy is when one of the channels fail, the data is not lost as redundant data resides on another channel.

Channel redundancy is implemented by selecting physical disks on different channels when using the **Create Virtual Disk Advanced Wizard**.

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

Creating A Channel-Redundant Virtual Disk

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

The following instructions provide information on creating a virtual disk that uses channel redundancy.

1. Launch the **Create Virtual Disk Advanced Wizard**:
 - a. In the **Server Administrator** window, under the system tree, click **Storage** dashboard.
 - b. Locate the controller on which you are creating a channel-redundant virtual disk and expand the controller object until the **Virtual Disks** object is displayed.
 - c. Select **Virtual Disks** and click **Go To The Create Virtual Disk Wizard**.
 - d. Click **Virtual Disk Advanced Wizard**.
2. Follow the steps in [Create Virtual Disk Advanced Wizard](#).
3. Click **Exit Wizard** to cancel the virtual disk creation. In this step, you select the channels and the disks to be used by the virtual disk. The selections you make determine whether the virtual disk is channel-redundant.

There are specific RAID level and configuration requirements for implementing channel redundancy. You must select the same number of physical disks on each channel that you use. For information on the number of physical disks that can be used for different RAID levels, see [Number Of Physical Disks Per Virtual Disk](#). For information on controller-specific implementations of the RAID levels, see [Controller - Supported RAID Levels](#).

Creating A Physical Disk For Channel Redundant Virtual Disks On PERC Controllers

The following sections describe creating a channel-redundant virtual disk using RAID 10 or RAID 50 on PERC controllers.

Creating A Channel-Redundant Virtual Disk Using RAID 10

To create a channel-redundant virtual disk using RAID 10:

1. Select one physical disk on each of the two channels.
2. Select an additional disk on each of the two channels. You have now selected the minimum number of disks for a RAID 10. Repeat step 2 until you have selected the desired number of disks.

3. Click **Continue** to exit.

Creating A Channel-Redundant Virtual Disk Using RAID 50

To create a channel-redundant virtual disk using RAID 50:

1. Select one physical disk on each of the three channels.
2. Select an additional disk on each of the three channels. You have now selected the minimum number of disks for a RAID 50. Repeat step 2 until you have selected the desired number of disks.
3. Click **Continue** to exit.

Connector Health

The connector health page displays the status of the connector and the components attached to the connector.

Controller Information

For information on the controller, see [Controllers](#)




Connector Components

For information on attached components, see [Enclosures And Backplanes](#).

Connector Properties And Tasks

To view information about the connector and execute connector tasks, use the connector properties and tasks page.

Table 27. Connector Properties




Property	Definition
	<p>These icons represent the severity or health of the storage component.</p> <p> — Normal/OK</p> <p> — Warning/Non-critical</p> <p> — Critical/Failure/Error</p> <p>For more information, see Storage Component Severity. A Warning or Critical severity may indicate that the connector is unable to communicate with attached devices such as an enclosure. Check the status of attached devices. For more information, see Cables Attached Correctly and Isolate Hardware Problems.</p>
Name	Displays the connector number.
State	<p>Displays the status of the connector. Possible values are:</p> <ul style="list-style-type: none"> • Ready — The connector is functioning normally. • Degraded — The connector has encountered a failure and is operating in a degraded state. • Failed — The connector has encountered a failure and is no longer functioning.
Connector Type	<p>Displays whether the connector is operating in RAID or SCSI mode. Depending on the controller type, the connector can be either a SCSI connector or a SAS port.</p>
Termination	<p>Indicates the termination type of the connector. Possible values are:</p> <ul style="list-style-type: none"> • Narrow — Indicates an 8-bit data bus.

Property	Definition
	<ul style="list-style-type: none"> • Wide — Indicates a 16-bit data bus. • Unknown — Indicates that the termination type is unknown. • Not Terminated — On a SCSI controller, this property indicates that the data bus is not terminated. This property is also displayed when the termination type is Unknown.
SCSI Rate	Displays the SCSI speed for a SCSI device.

Logical Connector Properties And Tasks

To display the information about the logical connector (connector in redundant path mode) and to execute connector tasks, use the logical connector properties and tasks page.

Table 28. Logical Connector Properties

Property	Definition
	<p>These icons represent the severity or health of the storage component.</p> <p> — Normal/OK</p> <p> — Warning/Non-critical</p> <p> — Critical/Failure/Error</p> <p>For more information, see Storage Component Severity.</p> <p>A Warning or Critical severity may indicate that the connector is unable to communicate with attached devices such as an enclosure. Check the status of attached devices. For more information, see Cables Attached Correctly and Isolate Hardware Problems.</p>
Name	Displays the connector number. The default value is 0 .
State	<p>Displays the status of the connector. Possible values are:</p> <ul style="list-style-type: none"> • Ready — The connector is functioning normally. • Degraded — The connector has encountered a failure and is operating in a degraded state. • Failed — The connector has encountered a failure and is no longer functioning.
Connector Type	Displays whether the connector is operating in RAID mode. The connector is always a SAS connector.

Related links

[Path Health](#)

Path Health

The path health of the connectors is represented as normal, warning, or critical. The possible values are displayed as **Available**, **Degraded**, or **Failed**.

If the enclosure health is displayed as degraded and further investigation shows all enclosure components (EMMs, Fans, Physical Disks, Power Supplies, and Temperature) to be in normal condition, select the **Information/Configuration** subtab of the enclosure to view details of the **Path Failure**.

Related links

[Setting The Redundant Path Configuration](#)

[Logical Connector Properties And Tasks](#)

Clearing The Connectors Redundant Path View

If you do not want the redundant path view, physically disconnect the connector port from the enclosure and reboot the system. After the system reboots, the user interface still displays the logical connector, but in a critical state. To clear the redundant path mode, select **Clear Redundant Path view** from the **Controller Tasks**.

Selecting this option clears the redundant path view and the connectors are represented on the user interface as Connector 0 and Connector 1.

Related links

[Setting The Redundant Path Configuration](#)

Connector Components

For information on attached components, see [Enclosure And Backplane Properties And Tasks](#).

Tape Drive

Tape drives contain several tape backup units (TBUs) on which data can be backed up. Storage Management enumerates the TBUs that are used for data back up. You can view the tape drives associated with a particular controller on the **Tape Drives** on **Controller** page.

Tape Drive Properties

Table 29. Tape Drive Properties

Property	Definition
ID	Displays the ID of the tape drive.
Name	Displays the name of the tape drive.
Bus Protocol	Displays the bus protocol type of the tape drive.
Media	Displays the media type of the tape drive.
Vendor ID	Displays the ID of the vendor.
Product ID	Displays the ID of the product.
SAS Address	Displays the SAS address of the tape drive.

RAID Controller Batteries

Some RAID controllers have batteries. If the controller has a battery, Storage Management displays the battery under the controller object in the tree view.

If there is a power outage, the controller battery preserves data that is in the volatile cache memory (SRAM) but not yet written to disk. The battery is designed to provide a minimum backup of 24 hours.

When a RAID controller is first installed in a server, the battery may need charging.

For information on alert messages, see the *Server Administrator Messages Reference Guide*.

Related links

[Battery Properties And Tasks](#)





[Starting A Learn Cycle](#)

[Initiating The Battery Delay Learn Cycle](#)

Battery Properties And Tasks

To view information about the battery and execute battery tasks, use the battery properties and tasks page.

Table 30. Battery Properties

Property	Definition
Status	<p>These icons represent the severity or health of the storage component.</p> <ul style="list-style-type: none"> •  — Normal/OK •  — Warning/Non-critical •  — Critical/Failure/Error <p>For more information, see Storage Component Severity.</p>
Name	Displays the name of the battery.
State	<p>Displays the status of the battery. Possible values are:</p> <ul style="list-style-type: none"> • Charging — The battery is undergoing the recharge phase of the battery learn cycle. For more information, see Starting A Learn Cycle. • Missing — The battery is missing on the controller. • Failed — The battery has failed and has to be replaced.
Learn Mode	<p>Displays the learn mode of the battery. Possible values are:</p> <ul style="list-style-type: none"> • Auto — Storage Management performs an automatic learn cycle based on the set time. • Warn — The learn cycle is past the default 90 days. <p>When the battery is in Warn mode, the state of the controller is displayed as degraded.</p> <p> NOTE: Warn is available only on SAS controllers with firmware version 6.1 and later.</p>
Next Learn Time	Displays the number of days and hours left before the controller firmware initiates the next learn cycle.

Property	Definition
Maximum Learn Delay	Displays the maximum number of days and hours that you can delay the battery learn cycle. The controller firmware automatically initiates the battery learn cycle. You cannot stop or pause the learn cycle, but you can delay it. For more information, see Initiating The Battery Delay Learn Cycle and Starting A Learn Cycle .

Related links

[Battery — Available Tasks](#)

Battery Tasks

To access the battery tasks:

1. In the **Server Administrator** window, under the system tree, expand the **Storage** dashboard to display the controller objects.
2. Expand a controller object.
3. Select **Battery**.
4. Select a task from the **Available Tasks** drop-down menu.
5. Click **Execute**.

Battery — Available Tasks

The available battery tasks are:

- [Starting A Learn Cycle](#)
- [Battery Transparent Learn Cycle](#)
- [Initiating The Battery Delay Learn Cycle](#)

Related links

[Battery Properties And Tasks](#)

Starting A Learn Cycle


Does my controller support this feature? See [Supported Features](#).

Select the **Start Learn Cycle** task to initiate the battery learn cycle.


The battery learn cycle discharges and then fully charges the controller battery.

The learn cycle recalibrates the battery-integrated circuit so that the controller can determine whether the battery can maintain the controller cache for the prescribed period in the event of a power loss.

While the learn cycle is in progress, the battery may not be able to maintain the cache during a power loss. If the controller is using **Write-Back Cache** policy, then the controller changes to **Write-Through Cache** policy until the learn cycle completes. The write-through cache policy writes the data directly to the disk and reduces the risk of data loss in the cache during a power outage.

 **NOTE: If you have set the controller to Force Write-Back Cache policy, then the cache policy is not changed during the learn cycle. When using Force Write-Back Cache policy, it is possible for data loss to occur if there is a power outage while the learn cycle is in progress.**

The controller firmware automatically initiates the learn cycle every 90 days. You can, however, delay the start time of the learn cycle for a further seven days, after which the firmware automatically initiates the learn cycle.

 **NOTE: The learn cycle cannot be performed while the battery is charging. If either a user or the controller firmware initiates the learn cycle while the battery is charging, then the battery learn state displays Requested. The learn cycle begins when the battery is fully charged.**

Related links

[Initiating The Battery Delay Learn Cycle](#)

Battery Transparent Learn Cycle

Transparent Learn Cycle (TLC) is a periodic operation that calculates the charge that is remaining in the battery to ensure that there is sufficient energy. The operation runs automatically, and causes no impact to the system or controller performance. The controller automatically performs TLC on the battery to calibrate and gauge its charge capacity once every 90 days. The operation can be performed manually, if required.

Initiating The Battery Delay Learn Cycle

Does my controller support this feature? See [Supported Features](#).

The controller firmware automatically initiates the battery learn cycle every 90 days. Although you cannot stop the firmware from running the learn cycle, you can delay the start time of the learn cycle for up to seven days.

Related links

[Starting A Learn Cycle](#)

[RAID Controller Batteries](#)

Delay The Battery Learn Cycle

1. Type a numerical value in the **Days** text box. The value must be within the 0–7 range. The value indicates the number of days for which you want to delay the battery learn cycle. The learn cycle can be delayed for a maximum of seven days.
2. Type a numerical value in the **Hours** text box. The value must be within the 0–23 range. The value indicates the number of hours for which you want to delay the battery learn cycle.
3. Click **Apply Changes**.
To exit and cancel your changes, click **Go Back To Battery Information Page**.

To Locate Delay Learn Cycle In Storage Management

To locate this task in Storage Management:

1. In the **Server Administrator** window, under the system tree, expand the **Storage** to display the controller objects.
2. Expand the controller object.
3. Select the Battery object.
4. Select **Delay Learn Cycle** from the **Available Tasks** drop-down menu.
5. Click **Execute**.

Related links

[Battery Properties And Tasks](#)

[Starting A Learn Cycle](#)

[Initiating The Battery Delay Learn Cycle](#)

Physical Disks Or Physical Devices

Physical disks or physical devices reside within an enclosure or are attached to the controller. On a RAID controller, physical disks or devices are used to create virtual disks.


Related links

[Physical Disk Tasks](#)

Guidelines To Replace A Physical Disk Or Physical Device

A replacement disk may not be the same model as the physical disks or devices in the storage enclosure. Use the following guidelines when replacing a disk:

- A drive within an array must be replaced with a drive of equal or greater capacity.
- Consider the bus speed while replacing a drive. Replacing drives with equal or greater bus speeds within the same array is supported. You can use both 3 GB and 6 GB drives in the same array. However, it is not recommended to replace a 6 GB drive with a 3 GB drive. Although it is fully functional, the performance could be impacted. Replacing 3 GB drives with 6 GB drives is supported. This occurs more frequently as legacy parts are depleted and warranty service is required.
- Consider the Revolutions Per Minute (RPM) while replacing a drive. Replacing drives with equal or greater spindle speeds within the same array is supported. You can use both 10000 rpm and 15000 rpm drives in the same enclosure. However, it is not recommended to replace a 15000 rpm drive with a 10000 rpm drive. Although it is fully functional, the performance could be impacted. Replacing 10000 rpm drives with 15000 rpm drives is supported. This scenario occurs while replacing parts from service inventory due to part unavailability.
- SAS and SATA drives on the same backplane but not within the same virtual disk.
- Solid-State Drives (SSD) and Hard-Disk Drives (HDD) on the same backplane but not within the same virtual disk.

 **NOTE: Except for combining SSD, SAS, and SATA drives, only upgrades are supported.**

Adding A New Disk To The System

To add a new physical disk to the system:


1. Install or attach the new physical disk (or disks) or physical devices. For more information on installing or attaching, see the documentation that came with the disk.
2. Perform the following steps:

For SAS Controllers

For SAS controllers, you must:

1. Check the alert log for an alert verifying that the system has identified the new disk. You may receive alert 2052 or 2294. For information on alert messages, see the *Server Administrator Messages Reference Guide*.

2. Refresh the display by clicking **Refresh** () or by changing pages.

 **NOTE: Clicking the Refresh button in the right pane refreshes only the right pane. To view the new physical disk in the left pane tree view, click the system name displayed at the top of the left pane, or click View → Refresh on the browser.**

The new physical disk or physical device is displayed in the tree view after refreshing the display. If the new disk is not displayed, restart the computer.


Related Information

- If you are replacing a disk that is part of a virtual disk, see [Replacing The Disk](#).
- If you want to include the new disk in a virtual disk, see [Virtual Disk Considerations For Controllers](#)
- For more information, see [RAID Controller Technology: SATA and SAS](#)

Replacing A Physical Disk Receiving SMART Alerts

Self Monitoring, Analysis and Reporting Technology (SMART) alerts are messages predicting that a disk may fail soon. If a physical disk is receiving SMART alerts replace the disk. Use the following procedures to replace a disk receiving SMART alerts:

Disk Is Part Of A Redundant Virtual Disk

 **CAUTION: To avoid potential data loss, perform a check consistency before removing a physical disk that is receiving SMART alerts. The check consistency verifies that all data is accessible within the redundant virtual disk and uses the redundancy to repair any bad blocks that may be present. In some circumstances, failure to perform a check consistency can result in data loss. This may occur, for example, if the physical disk receiving SMART alerts has bad disk blocks and you do not perform a check consistency before removing the disk.**

1. Select the redundant virtual disk that includes the physical disk that is receiving SMART alerts and perform the **Check Consistency** task. For more information, see [Performing A Check Consistency](#)
2. Select the disk that is receiving SMART alerts and execute the **Offline** task.
3. Manually remove the disk.
4. Insert a new disk. Make sure that the new disk is the same size or greater in size than the disk you are replacing. On some controllers, you may not be able to use the additional disk space if you insert a disk with a greater disk space than the one you are replacing. For more information on disk space considerations, see [Virtual Disk Considerations For Controllers](#). After you complete this procedure, a rebuild is automatically initiated because the virtual disk is redundant.

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 greater in size than the original virtual disk. For controller-specific information on creating virtual disks, see [Virtual Disk Considerations For Controllers](#).
5. Restore the backed up data from the original virtual disk onto the newly created virtual disk.

Related links

[Monitoring Disk Reliability On RAID Controllers](#)

Other Disk Procedures




- [Replacing A Failed Disk](#)
- [Recovering From Removing The Wrong Physical Disk](#)
- [Moving Physical And Virtual Disks From One System To Another](#)
- [Troubleshooting](#)


Physical Disk Or Physical Device Properties


To view information about physical disks or physical device and perform physical disk or physical device tasks, use the Physical Disk Or Physical Device Properties page.

 **NOTE: Physical devices are applicable only to PCIe SSD.**

Table 31. Physical Disk Properties

Property	Definition
Status	<p>These icons represent the severity or health of the storage component.</p> <p> — Normal/OK</p> <p> — Warning/Non-critical</p> <p> — Critical/Failure/Error</p> <p>For more information, see Storage Component Severity.</p>
Power Status	<p>Displays the power status of the physical drives. The power status is present only for H700 and H800 controllers and is not present in the Physical Disk for Virtual Disk page.</p>
Spun Down	<p>Indicates that the physical drive is in spun down state. Only hot spare and unconfigured disk can be in spun down state if there is no activity on the drives for a specified interval of time.</p>
Transition	<p>Indicates that the physical drive is changing from the spun down state to spun up state.</p>
Spun Up	<p>Indicates that the physical drive is in the spun up state.</p>
Name	<p>Displays the name of the physical disk or device. The name includes of the connector number followed by the disk number.</p>
State	<p>Displays the current state of the physical disk or device. Possible values are:</p> <ul style="list-style-type: none"> • Ready — The physical disk or device is functioning normally. If the disk is attached to a RAID controller Ready indicates that the disk is available to be used by a virtual disk. When the physical disk or device is used in a virtual disk, the state changes to Online. • Online — Indicates that the physical disk is part of a virtual disk and is functioning normally. For more information, see Setting Online And Offline. • Degraded — The physical disk or device has encountered a failure and is operating in a Degraded state. • Failed — The physical disk or device has encountered a failure and is no longer functioning. This state is also displayed when a physical disk or device that is part of a redundant virtual disk is offline or deactivated. For more information, see Setting Online And Offline. • Offline — The physical disk or device has failed or contains dead segments. Check to see if the Remove Dead Segments task appears on the physical disk drop-down menu. Refresh the system and then remove the dead segment (Removing Dead Segments) for the physical disk. If this task is not displayed, then the physical disk or device cannot be recovered. • Rebuilding — Data from a redundant virtual disk is being rebuilt onto the physical disk or device. • Incompatible — The physical disk or device is not suitable for a rebuild. The physical disk or device may be too small or it may be using an incompatible technology. For example, you cannot rebuild a SAS disk with a SATA disk or a SATA disk with a SAS disk. • Removed — The physical disk or device has been removed. This state applies only to physical disks that are part of a virtual disk. • Clear — The Clear task is being performed on the physical disk or device. A physical disk or device may also display the Clear state if the physical disk or device is a member of a virtual disk that is being slow initialized. For more information, see Performing A Clear Physical Disk And Cancel Clear and Slow And Fast Initialize. • SMART Alert Detected — A SMART alert (predictive failure) has been detected on the physical disk or device. The physical disk or device may fail and should be replaced. This state applies to physical disks or devices attached to non-RAID controllers.

Property	Definition
	<ul style="list-style-type: none"> • Unknown — The physical disk or device has failed or is in an unusable state. At times, the physical disk or device can be returned to a usable state by performing a Format, Initialize, Slow, And Fast Initialize task. If the Format, Initialize, Slow, And Fast Initialize task does not appear on the physical disk or device drop-down menu, then this disk or device cannot be recovered. • Foreign — The physical disk has been moved from another controller and contains all or some portions of a virtual disk (foreign configuration). A physical disk or device that has lost communication with the controller due to a power loss, faulty cable or other failure event may also display the Foreign state. For more information, see Foreign Configuration Operations. • Unsupported — The physical disk or device is using an unsupported technology or it may not be certified by your service provider. The physical disk cannot be managed by Storage Management. • Replacing — A Replace Member Disk task is being performed on the physical disk or device. For more information, see Replacing A Member Disk and Enabling Reversible Hot Spare. <p> NOTE: You can cancel the copying of data at any time during the execution of this task.</p> <ul style="list-style-type: none"> • Non-RAID — Non-RAID disks are exposed to the operating system unlike unconfigured disks and enables usage of disk in direct pass-through mode. Maximum number of non-RAID disks that can be supported on H310 controller are 64. <p>You can perform the following tasks on the Non-RAID disks:</p> <ul style="list-style-type: none"> • Identified as locate option. • Perform a blink or unblink operation. • Select the disk as a bootable device. <p>You cannot perform the following tasks on the disk:</p> <ul style="list-style-type: none"> • Force disk offline or online. • Select as part of a virtual disk. • Assign a hot spare. • Choose as source or target for rebuild, copyback, replace member, or reconstruct. • Spun down to save power. • Select as a bootable device.
Certified	Displays whether the physical disk or device is certified by your service provider.
Mirror Set ID	Displays the mirror set ID of the member physical disk or device that has duplicated data from another physical disk or device.
Capacity	Displays the capacity of the disk.
Failure Predicted	<p>Displays whether the physical disk or device has received a SMART alert and is therefore predicted to fail. For more information on SMART predictive failure analysis, see Monitoring Disk Reliability On RAID Controllers. For information on replacing the physical disk, see Replacing A Physical Disk Receiving SMART Alerts.</p> <p>You may also want to review the alert log to see whether the physical disk or device has generated alerts pertaining to a SMART predictive failure. These alerts can assist you in identifying the cause of the SMART alert. The following alerts may be generated in response to a SMART alert:</p> <ul style="list-style-type: none"> • 2094 • 2106 • 2107 • 2108 • 2109


Property	Definition
	<ul style="list-style-type: none"> · 2110 · 2111 <p>For information on alert messages, see the <i>Server Administrator Messages Reference Guide</i>.</p>
Progress	Displays the progress of an operation that is being an operation is that is being performed on the physical disk or device.
Encryption Capable	Displays whether the physical disk or device is a Self-Encryption Disk (SED). The possible values are Yes and No .
Encrypted	Displays whether the physical disk or device is encrypted to the controller. The possible values are Yes and No . For a non-SED, the value is N/A .
Bus Protocol	Displays the technology that the physical disk or device is using. Possible values are: <ul style="list-style-type: none"> · SAS · SATA
Device Protocol	Displays the device protocol of the physical device, such as Non-Volatile Memory Express (NVMe).
Media	Displays the media type of the physical disk or device. The possible values are: <ul style="list-style-type: none"> · HDD — Hard Disk Drive. An HDD is a nonvolatile storage device which stores digitally encoded data on rapidly rotating platters with magnetic surfaces. · SSD — Solid-State Drive. An SSD is a data storage device that uses solid-state memory to store persistent data. · Unknown — Storage Management is unable to determine the media type of the physical disk or device.
Remaining Rated Write Endurance	Displays information on the SSD renewal/replacement based on the amount of write workloads. This field indicates the total remaining programs or erase-cycles available on the SSD, based on the cumulative specification of the total NAND (Negated AND or NOT AND) Flash chips in the SSD. <p> NOTE: This option is applicable to Micron PCIe SSDs, Non-Volatile Memory Express (NVMe) PCIe SSDs, and SAS/SATA SSDs.</p>
Used RAID Disk Space	Displays the physical disk or device space that is used by the virtual disks on the controller. This property is Not Applicable for physical disks or devices attached to non-RAID controllers. <p>In certain circumstances, the Used RAID Disk Space displays a value of zero even though a portion of the physical disk or device is being used. This occurs when the used space is 0.005GB or less. The algorithm for calculating the used disk space rounds a figure of 0.005GB or less to 0. Used disk space that is between 0.006GB and 0.009GB is rounded up to 0.01GB.</p>
Available RAID Disk Space	Displays the amount of available space on the disk. This property is Not Applicable for physical disks attached to non-RAID controllers.
Hot Spare	Indicates whether the disk has been assigned as a hot spare. This property is Not Applicable for physical disks attached to non-RAID controllers.
Vendor ID	Displays the hardware vendor of the disk.
Product ID	Displays the product ID of the device.
Firmware Revision	Displays the firmware version of the physical device.

Property	Definition
Serial No.	Displays the serial number of the disk.
Part Number	Displays the Piece Part Identification (PPID) of the physical drive.
T10 Protection Information Capability	Indicates whether the physical disk supports data integrity. The possible values are Yes and No .
Sector Size	Displays the sector size of the physical disk. The possible options are 512B and 4KB .
PCIe Negotiated Link Speed	Displays the current negotiated transfer speed of the physical device in GT/s.
PCIe Maximum Link Speed	Displays the capable transfer speed of the physical device in GT/s.
Manufacture Day	Displays the day of the month when the physical disk was manufactured.
Manufacture Week	Displays the week of the year when the physical disk was manufactured.
Manufacture Year	Displays the year when the physical disk was manufactured.
SAS Address	Displays the SAS address of the physical disk. The SAS address is unique to each SAS disk.
PCIe Negotiated Link Width	Displays the current negotiated link width of the physical device.
PCIe Maximum Link Width	Displays the capable link width of the physical device.
ISE Capable	Displays whether the physical disk or device is an Instant Secure Erase (ISE) capable device. The possible values are Yes and No .

Physical Disk Or Physical Device Tasks

To execute a physical disk or physical device task:

1. Expand the **Storage** tree object to display the controller objects.
2. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
3. Expand the enclosure or backplane object.
4. Select **Physical Disks** or **Physical Devices** object.
5. Click **Information/Configuration** on the controller **Properties** page.
6. Select a task from the drop-down menu.
7. Click **Execute**.

 **NOTE: Different controllers support different features. For this reason, the tasks displayed on the tasks drop-down menu can vary depending on which controller is selected in the tree view. No Task Available is the only option in the drop-down list box when no tasks can be performed because of controller or system configuration limitations.**

Physical Disk Tasks

The following are the list of tasks that you can perform on a physical disk:

- [Blinking And Unblinking A Physical Disk](#)
- [Removing Dead Segments](#)
- [Assigning And Unassigning Global Hot Spare](#)
- [Preparing To Remove](#)
- [Setting Online And Offline](#)
- [Rebuilding Data](#)
- [Cancelling A Rebuild](#)
- [Performing A Clear Physical Disk And Cancel Clear](#)


- [Enabling Reversible Hot Spare](#)
- [Performing Cryptographic Erase](#)
- [Convert To RAID Capable Disk](#)
- [Convert To Non-RAID Disk](#)

Blinking And Unblinking A Physical Disk

Does my controller support this feature? See [Supported Features](#).

The **Blink** task allows you to find a disk within an enclosure by blinking one of the light-emitting diodes (LEDs) on the disk. You can use this task to locate a failed disk.

On most controllers, the **Blink** task automatically cancels after a short duration such as 30–60 seconds. Select **Unblink** to cancel the **Blink** task or to stop the LED on the physical device that is blinking indefinitely.

 **NOTE: The Blink and Unblink tasks are only supported for hotswap physical disks (disks that reside in a carrier). When using an Broadcom PCIe U320 controller, the Blink and Unblink tasks apply to physical disks contained in carriers that can be inserted into a server or an enclosure. If the physical 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 tasks are disabled.**

Removing Dead Segments

Does my controller support this feature? See [Supported Features](#).

The **Remove Dead Segments** task recovers disk space that is unusable. A *dead* or *orphaned* disk segment refers to an area of a physical disk or physical device that is unusable for any of the following reasons:

- The dead segment is an area of the physical disk or physical device that is damaged.
- The dead segment is included in a virtual disk, but the virtual disk is no longer using this area of the physical disk or physical device.
- The physical disk or physical device contains more than one virtual disk. In this case, disk space that is not included in one of the virtual disks may be unusable.
- The dead segment resides on a physical disk or physical device that has been disconnected and then reconnected to the controller.

Preparing To Remove

 **NOTE: Cryptographic Erase is not supported for NVMe devices connected to PERC S140 controller.**


Does my controller support this feature? See [Supported Features](#).

Use the **Prepare to Remove** task to spin down a physical disk or physical device so that it can be safely removed from an enclosure or backplane. It is recommended that you perform this task before removing a disk or device to prevent data loss.

This task causes the LED on the disk to blink. You can safely remove the disk or device by following the conditions listed below:

- Wait for about 30 seconds to allow the disk to spin down.
- Wait until you notice that the initial blink pattern has changed into a different pattern or the lights have stopped blinking.

A physical disk or physical device is no longer in ready state performing. Removing the physical disk or device from the enclosure or backplane and replacing it causes the physical disk or device to spin up and return to **Ready** state.

 **NOTE: This procedure is not available for physical disks or devices that have been assigned as a hot spare or physical disk or devices that are part of a virtual disk. In addition, this procedure is only supported for hot swap physical disks or devices (disks that reside in a carrier.)**

Rebuilding Data

Does my controller support this feature? See [Supported Features](#).

Select the **Rebuild** task to reconstruct data when a physical disk in a redundant virtual disk fails.

 **NOTE: Rebuilding a disk may take several hours.**


Related links

[Replacing A Failed Disk That Is Part Of A Redundant Virtual Disk](#)

Canceling A Rebuild

Does my controller support this feature? See [Supported Features](#).

Select the **Cancel Rebuild** task to cancel a rebuild that is in progress. If you cancel a rebuild, the virtual disk remains in a **Degraded** state. The failure of an additional physical disk can cause the virtual disk to fail and may result in data loss. It is recommended that you rebuild the failed physical disk as soon as possible.

 **NOTE: If you cancel the rebuild of a physical disk that is assigned as a hot spare, reinitiate the rebuild on the same physical disk in order to restore the data. Canceling the rebuild of a physical disk and then assigning another physical disk as a hot spare does not cause the newly assigned hot spare to rebuild the data. Reinitiate the rebuild on the physical disk that was the original hot spare.**

Assigning And Unassigning Global Hot Spare


Does my controller support this feature? See [Supported Features](#).

A global hot spare is an unused backup disk that is part of the disk group. Hot spares remain in standby mode. When a physical disk that is used in a virtual disk fails, the assigned hot spare is activated to replace the failed physical 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 physical disk.

You can change the hot spare assignment by unassigning a disk and choosing another disk as needed. You can also assign more than one physical disk as a global hot spare.

Global hot spares must be assigned and unassigned manually. They are not assigned to specific virtual disks. If you want to assign a hot spare to a virtual disk (it replaces any physical disk that fails in the virtual disk), then use the **Assign and Unassign Dedicated Hot Spare**.

 **NOTE: When deleting virtual disks, all assigned global hot spares may be automatically unassigned when the last virtual disk associated with the controller is deleted. When the last virtual disk of a disk group is deleted, all assigned dedicated hot spares automatically become global hot spares.**

 **NOTE: For PERC H330, H730, H730P, H740P, and PERC H840 controllers, if any of the drives you selected are in the Spun Down state, the following message is displayed: The current physical drive is in the spun down state. Executing this task on this drive takes additional time, because the drive needs to spun up.**

You should be familiar with the size requirements and other considerations associated with hot spares.

Related links

[Protecting Your Virtual Disk With A Hot Spare](#)

Setting The Physical Disk Online Or Offline

Does my controller support this feature? See [Supported Features](#).

The **Online** and **Offline** tasks apply only to physical disks that are included in a redundant virtual disk and attached to a PERC controller.

Select the **Offline** task to deactivate a disk before removing it. Select the **Online** task to reactivate an offline disk. In some cases, you may want to use the **Online** task on a failed disk in an attempt to recover data from the disk.

Related links

[Using The Physical Disk Online Command On Select Controllers](#)

Online Or Offline The Physical Disk

To set the Online or Offline Task the physical disk:

1. Review the physical disk that must be made online or offline. There may be data loss when performing this task. Back up your data, if necessary. If you want to blink the physical disk, click **Blink**.
2. Click **Online** or **Offline** when ready or click **Return to Previous Page**.

To Locate Online Or Offline In Storage Management


To locate this task in storage management:

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Expand a controller object.
3. Expand a connector object.
4. Expand the enclosure or backplane object.
5. Select the **Physical Disks** object.
6. Select **Online** or **Offline** from the **Tasks** drop-down menu of the physical disk you want to make online or offline.
7. Click **Execute**.

Performing A Clear Physical Disk And Cancel Clear

Does my controller support this feature? See [Supported Features](#).

Select the **Clear Physical Disk** task to erase data residing on a physical disk. The **Clear** task applies to physical disks that are in **Ready** state and that contain data or that are in **Clear** state.

 **NOTE: A physical disk may display the Clear state if it is a member of a virtual disk that is being slow initialized. Performing a Cancel Clear task on the physical disk causes the Slow Initialize task to be canceled for the entire virtual disk.**

To clear the physical disk:

1. Review the physical disk to be erased. Ensure that it does not contain necessary data and make a backup if necessary. If you want to blink the physical disk, click **Blink** button.
2. Click **Clear** when you are ready to erase all information on the physical disk. To exit without clearing the physical disk, click **Return to Previous Page**.

Related links

[Format, Initialize, Slow, And Fast Initialize Considerations For Slow Initialize](#)

To Locate Clear In Storage Management

To locate this task in storage management:

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Expand a controller object.
3. Expand the enclosure or **Backplane** object.
4. Select the **Physical Disks** object.
5. Select **Clear** from the **Tasks** drop-down menu of the physical device you want to clear.
6. Click **Execute**.

Enabling Revertible Hot Spare

Does my controller support this feature? See [Supported Features](#).

Select the **Revertible Hot Spare** task to copy data back from a hot spare to a physical disk.

If the physical disk in a virtual disk fails, the data on the failed disk is copied to the assigned hot spare. If you replace the failed disk with a new physical disk *and* if you have enabled the **Revertible Hot Spare** task, the data is copied from the hot spare to the new disk.

You can also use the **Revertible Hot Spare** task to copy data from a physical disk to the hot spare on a predictive failure event.

If Revertible Hot Spare is enabled *and* the physical disk is SMART-enabled, the controller firmware automatically starts copying data from the SMART-enabled disk in the virtual disk to the hot spare.

 **NOTE: To use the Revertible Hot Spare task, ensure that you have assigned a hot spare to the virtual disk.**

 **NOTE: If the disk is not SMART-enabled or if the Auto Replace on Predictive Failure option is disabled, the failed disk is not replaced automatically.**

Enable Revertible Hot Spare

To enable revertible hot spare:

1. On the **Change Controller Properties** page, enable **Allow Revertible Hot Spare** and **Replace Member Auto Replace Member on Predictive Failure**.
2. Click **Apply Changes**.

To Locate Controller Task In Storage Management

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Select a controller object on which you want to enable the revertible hot spare task.
3. Click **Information/Configuration** on the controller Properties page.
4. From the **Controller Task** drop-down list box, select **Change Controller Properties** and click **Execute**.

 **NOTE: The Rebuild rate for Revertible Hot Spare is the same as that defined for the controller.**

Performing Cryptographic Erase

Does my controller support this feature? See [Supported Features](#).

 **CAUTION: Cryptographic Erase permanently erases all data present on the disk.**

Select the **Cryptographic Erase** task to erase an encrypted physical disk. This task is available for:

- Instant Secure Erase (ISE) capable drives
- Unconfigured SED drive
- Foreign configured encrypted drives
- Unconfigured and foreign SED drive even when an Encryption Key is not present in the controller

Related links

[Managing The Encryption Key](#)

[To Locate Cryptographic Erase In Storage Management](#)

To Locate Cryptographic Erase In Storage Management

To locate this task in storage management:

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Expand a controller object.
3. Expand a connector object.
4. Expand the enclosure or backplane object.
5. Select the **Physical Disk** object.
6. Select **Cryptographic Erase** from the Tasks drop-down menu of the physical device you want to clear.
7. Click **Execute**.

After the task is completed, the Alert Log is populated.

Convert To RAID Capable Disk

This task enables a disk for all RAID operations.

 **NOTE: This task is not supported on PERC hardware controllers running in HBA mode.**

 **NOTE: This task is not supported on PERC 10 controllers.**

Convert To Non-RAID Disk

This task converts a disk to a Non-RAID disk. After converting a disk to non- RAID, the disk is exposed to the operating system unlike unconfigured good disks and it enables usage of disk in direct pass-through mode.

 **NOTE: This task is not supported on PERC hardware controllers running in HBA mode.**

Virtual Disks

A virtual disk refers to a storage created by a RAID controller from one or more physical disks. Although a virtual disk may be created from several physical disks, it is viewed by the operating system as a single disk. Depending on the RAID level used, the virtual disk may retain redundant data if there is a disk failure or have particular performance attributes.

 **NOTE: Virtual disks can only be created on a RAID controller.**

Related links

[Understanding RAID concepts](#)

[Deleting A Virtual Disk](#)

Considerations Before Creating Virtual Disks

Different controllers have particular characteristics in the way they implement virtual disks. These characteristics may include use of disk space, limitations on the number of virtual disks per controller, and so on. It can be helpful to understand these characteristics before creating virtual disks on the controller.

The following sections describe controller information that applies to virtual disks:

- [Virtual Disk Considerations For Controllers](#)
- [Problems Associated With Using The Same Physical Disks For Both Redundant And Non-Redundant Virtual Disks](#)
- [Virtual Disk Considerations On Systems Running Linux](#)
- [Number Of Physical Disks Per Virtual Disk](#)
- [Number Of Virtual Disks Per Controller](#)
- [Calculation For Maximum Virtual Disk Size](#)

You may also want to review the following sections:

- [RAID Controller Read, Write, Cache, And Disk Cache Policy](#)
- [Understanding Hot Spares](#)
- [Controller - Supported Stripe Sizes](#)
- [Time Delay In Displaying Configuration Changes](#)

 **NOTE: In addition to this document, review the hardware documentation that is provided with the controllers. Reviewing the hardware documentation along with this document may provide a better understanding of the controller limitations.**

Virtual Disk Considerations For Controllers

In addition to the considerations described in this section, you should also be aware of the controller limitations described in Number of Physical Disks per Virtual Disk for the following controllers:

- PERC H330 Adapter, PERC H330 Mini Monolithic
- PERC H730 Adapter, PERC H730 Mini Monolithic
- PERC H730P Adapter, PERC H730P Mini Monolithic
- PERC H740P Adapter, PERC H740P Mini Monolithic
- PERC H840 Adapter



NOTE: The order of the controllers displayed on Storage Management may differ with the order of the controllers displayed in the Human Interface (HII) and PERC Option ROM. The order of the controllers does not cause any limitation.

The following considerations apply when creating virtual disks:

- Creating virtual disks on controllers — When you create a virtual disk, you specify which physical disks are to be included in the virtual disk. The virtual disk you create spans the specified physical disks. Depending on the size of the virtual disk, the virtual disk may not consume all of the space on the physical disks. Any leftover space on the physical disks cannot be used for a second virtual disk unless the physical disks are of equal size. In addition, when the physical disks are of equal size and you use the leftover space for a second virtual disk, this new virtual disk cannot expand to include any physical disks not included in the original virtual disk.
- Space allocation when deleting and creating virtual disks on controllers — When you delete a virtual disk, you free up or make available space on the physical disks that was being used by the deleted virtual disk. If you have created several virtual disks on a disk group, then deleting virtual disks can result in pockets of free space residing in various locations on the physical disks. When you create a new virtual disk, the controller must decide which free space on the physical disks to allocate to the new virtual disk. The PERC controllers look for the largest area of free space and allocate this space to the new virtual disk.
- SCSI limitation of 2TB — Virtual disks created on a PERC controller cannot be created from physical disks with an aggregate size greater than 2TB. This is a limitation of the controller implementation. For example, you cannot select more than 30 physical 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 indicating that the 2TB limit has been reached, and that you should select a smaller number of physical disks. The 2TB limit is an industry-wide SCSI limitation.
- Expanding virtual disks — You can only use the **Reconfigure** task to expand a virtual disk that uses the full capacity of its member physical disks.
- Reconfiguring virtual disks — The **Reconfigure** task is not available when you have more than one virtual disk using the same set of physical disks. You can, however, reconfigure a virtual disk that is the only virtual disk residing on a set of physical disks.
- Virtual disk names not stored on controller — The names of the virtual disks that you create are not stored on the controller. 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.
- **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 is not lost because redundant data resides on another channel.



NOTE: For more information about channel redundancy, see [Channel Redundancy And Thermal Shutdown](#).

- Rebuilding data — An failed physical disk that is used by both redundant and nonredundant virtual disks cannot be rebuilt. Rebuilding a failed physical disk in this situation requires deleting the nonredundant virtual disk.
- Disk group concept consideration for S110 — Disk grouping is a logical grouping of disks attached to a RAID controller on which one or more virtual disks are created, such that all virtual disks in the disk group use all of the physical disks in the disk group. The current implementation supports the blocking of mixed disk groups during the creation of logical devices.

Physical disks are bound to disk groups, therefore, there is no RAID level mixing on one disk group.

Storage Management Server implements the disk group concept during virtual disk creation. Functionally, after a group of physical disks is used to create their first virtual disk, unused space in the disk is used only to expand the virtual disk, or create new virtual disks in the unused space. The virtual disks have identical RAID level.

Also, existing mixed configuration is not affected. However, you cannot create mixed configurations.

You can read or write to the virtual disks, rebuild, and delete the disks.

You cannot create virtual disks on a set of disks migrated from earlier software RAID versions and configured with multiple RAID levels.

Related links

[Channel Redundancy](#)

[Virtual Disk Task: Reconfigure \(Step 1 of 3\)](#)

Virtual Disk Considerations On Systems Running Linux

On some versions of the Linux operating system, the virtual disk size is limited to 1TB. Before creating a virtual disk that is larger than 1TB, you should make sure that your operating system supports this virtual disk size. The support provided by your operating system depends on the version of the operating system and any updates or modifications that you have implemented. In addition, you should

investigate the capacity of your peripheral devices to support a virtual disk that is larger than 1TB. For more information, see your operating system and device documentation.

Number Of Physical Disks Per Virtual Disk

There are limitations on the number of physical disks that can be included in the virtual disk. These limitations depend on the controller. When creating a virtual disk, the controllers support some stripes and spans (methods for combining the storage on physical disks). Since, the number of total stripes and spans is limited, the number of physical 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 RAID 10, RAID 50, and RAID 60.
- Maximum number of stripes affects RAID 0, RAID 5, RAID 50, RAID 6, and RAID 60.
- Number of physical disks in a mirror is always 2. This property affects RAID 1 and RAID 10.

If RAID 50 and RAID 60, you can use a greater number of physical disks than is possible for other RAID levels. RAID 10 on a SAS controller with firmware version 6.1 can use a maximum of 256 physical disks. However, the number of connectors on the controller imposes limitations on how many physical disks can be included in a virtual disk when using RAID 10, RAID 50, or RAID 60. This is because only a limited number of physical disks can be physically attached to the controller.

For information about the maximum number of physical disks supported by a virtual disk, see the virtual disk specifications for the controller in [Supported Features](#).

Number Of Virtual Disks Per Controller

There are limitations on the number of virtual disks that can be created on the controller. For information about the maximum number of virtual disks supported by a controller, see the virtual disk specifications for the controller in [Supported Features](#).

Calculation For Maximum Virtual Disk Size

The **Create Virtual Disk Express Wizard** displays the minimum and maximum values for the virtual disk size. This section describes how the maximum possible size for the virtual disk is calculated based on the controller type. To identify the controller type, see [RAID Controller Technology: SATA and SAS](#).

SATA RAID Controllers

When using a SATA RAID controller, the controller calculates the maximum virtual disk size based on your RAID level selection and the available physical disk space provided by all suitable physical disks attached to the controller. For example, if the controller contains 12 physical disks with available space and you have specified a RAID 5, then the controller calculates the maximum virtual disk size based on the disk space provided by the 12 physical disks, because the 12 physical disks can be included in a RAID 5.

SAS RAID Controllers

When using a SAS controller, the controller calculates the maximum virtual disk size based on the available disk space provided by the minimum number of physical disks required to create the RAID level you selected. For example, if you specified a RAID 5, then the controller calculates the maximum virtual disk size based on three physical disks, because only three physical disks are required to create a RAID 5.

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.

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

Related links

[Channel Redundancy](#)

[Creating A Channel-Redundant Virtual Disk](#)

Creating Virtual Disks

Does my controller support this feature? See [Supported Features](#).

In order to implement RAID functions, you must create a virtual disk. A virtual disk refers to storage created by a RAID controller from one or more physical disks. Although a virtual disk may be created from several physical disks, it is seen by the operating system as a single disk.

Before creating a virtual disk, you should be familiar with the information in [Considerations Before Creating Virtual Disks](#).

Storage Management provides wizards to help you create a virtual disk:

- **The Create Virtual Disk Express Wizard** calculates an appropriate virtual disk layout based on the available space and controller considerations. Using this wizard, you can quickly create a virtual disk using recommended selections.
- **The Create Virtual Disk Advanced Wizard** allows you to specify the read, write, and cache policy for the virtual disk. You can also select the physical disks and the controller connector to be used. You need a good knowledge of RAID levels and hardware to use the advanced wizard.

Related links

- [What is RAID?](#)
- [Deleting A Virtual Disk](#)
- [Channel Redundancy](#)
- [Reconfiguring Or Migrating Virtual Disks](#)
- [Physical Disk Or Physical Device Properties](#)
- [Physical Disk Or Physical Device Tasks](#)
- [Create Virtual Disk Advanced Wizard](#)
- [Create Virtual Disk Express Wizard](#)

Reconfiguring Or Migrating Virtual Disks

Does my controller support this feature? See [Supported Features](#).

Reconfiguring or migrating a virtual disk enables you to increase the capacity or change the RAID level of the virtual disk.

To reconfigure a virtual disk:

1. Review the information in Starting and Target RAID Levels for Virtual Disk Reconfiguration and Capacity Expansion.
2. Locate the controller on which the virtual disk resides in the tree view. Expand the controller object until the **Virtual Disks** object is displayed.
3. Select the **Reconfigure** task from the virtual disk's **Tasks** drop-down list box and click **Execute**.
4. Complete the **Reconfigure** task using the Reconfigure wizard.




Starting And Target RAID Levels For Virtual Disk Reconfiguration And Capacity Expansion

After you have created a virtual disk, the possibilities for reconfiguring the virtual disk depend on the controller, RAID level, and available physical disks.

Table 32. Possible scenarios for reconfiguring a virtual disk

Controller	Starting RAID Level	Target RAID Level	Comments
PERC H330 Adapter, PERC H330 Mini Monolithic, PERC H730 Adapter, PERC H730 Mini Monolithic, PERC H730P Adapter, PERC H730P Mini Monolithic, PERC H740P Adapter, PERC H740P Mini	RAID 0	RAID 1	Add a single disk
	RAID 0	RAID 0, RAID 5	Add at least one additional disk.
	RAID 0	RAID 6	RAID 6 requires a minimum of 4 disks.

Controller	Starting RAID Level	Target RAID Level	Comments
Monolithic, PERC H840 Adapter			Reconfiguration from RAID 0 to RAID 6 requires at least 2 additional disks even when this exceeds the 4-disk minimum required by RAID 6
	RAID 1	RAID 0	With or without adding additional disks
	RAID 1	RAID 5, RAID 6	Add at least one additional disk. RAID 6 requires a minimum of 4 disks.
	RAID 5	RAID 0	With or without adding additional disks
	RAID 5	RAID 5, RAID 6	Add at least one additional disk. RAID 6 requires a minimum of 4 disks.
	RAID 6	RAID 0, RAID 5	With or without adding additional disks
	RAID 6	RAID 6	Add at least one additional disk
	RAID 10	RAID 10	Without adding additional disks
PERC S140	RAID 0	RAID 0	With or without additional disks
	RAID 1	RAID 1	Without additional disks
	RAID 5	RAID 5	With or without additional disks
	RAID 10	RAID 10	Without additional disks

-  **NOTE:** The order of the controllers displayed on Storage Management may differ with the order of the controllers displayed in the Human Interface (HII) and PERC Option ROM. The order of the controllers does not cause any limitation.
-  **NOTE:** When creating virtual disks using software RAID controllers, the information related to the physical disks linked to the virtual disk is enumerated or displayed on Storage Management after a short delay. This delay in displaying the information does not cause any functional limitation. If you are creating partial virtual disks, it is recommended that you provide Storage Management adequate time between each partial virtual disk creation process.
-  **NOTE:** RAID 10 virtual disk reconfiguration operation is not supported with Intelligent Mirroring.

Related links

[Choosing RAID Levels And Concatenation](#)

Maintaining The Integrity Of Redundant Virtual Disks

Does my controller support this feature? See [Supported Features](#).

The virtual disk **Check Consistency** task verifies the accuracy of the redundant (parity) information. This task only applies to redundant virtual disks. When necessary, the **Check Consistency** task rebuilds the redundant data.

To verify redundant information of a virtual disk:

1. Locate the controller on which the virtual disk resides in the tree view. Expand the controller object until the **Virtual Disks** object is displayed.
2. Select the **Check Consistency** task from the virtual disk's **Tasks** drop-down list box and click **Execute**.

Rebuilding Redundant Information

Does my controller support this feature? See [Supported Features](#).

If you have a redundant virtual disk, you can reconstruct the contents of a failed physical disk onto a new disk or a hot spare. A rebuild can take place during normal operation, but it degrades performance.

Related links

- [Replacing A Failed Disk](#)
- [Setting The Rebuild Rate](#)
- [A Rebuild Does Not Work](#)

Managing Virtual Disk Bad Block Management

Virtual disk bad blocks are bad blocks on one or more member physical disks. The read operation on the virtual disks having bad blocks may fail.

Storage Management generates a critical alert (2387) to notify you of the bad blocks on the virtual disk.

Virtual disk bad blocks are discovered when the controller performs any operation that requires scanning the disk. Examples of operations that may result in this alert are:

- Consistency check
- Rebuild
- Virtual disk format
- I/O
- Patrol Read

Recovering a physical disk bad block depends on the RAID level and state of the virtual disk. If a virtual disk is redundant, the controller can recover a bad block on a physical disk. If a virtual disk is not redundant, then the physical disk bad block results in a virtual disk bad block.

Table 33. Sample Scenarios For Virtual Disk Bad Blocks

RAID Level Virtual Disk	State	Scenario	Result
RAID 0	Degraded	One bad block on a physical disk.	The controller cannot regenerate data from the peer disks as there is no redundancy. This results in a virtual disk bad block.
RAID 5	Ready	One bad block on a physical disk.	The controller regenerates data from the peer disks and sends a Write to the bad block. The disk then remaps the Logical Block Addressing (LBA) to another physical location. The problem is resolved.
RAID 5	Degraded	One bad block on a physical disk.	The controller cannot regenerate data from the peer disks because one drive is missing. This results in a virtual disk bad block.

RAID Level Virtual Disk	State	Scenario	Result
RAID 5	Ready	One bad block on two physical disks at the same location.	The controller cannot regenerate data from the peer disks. This results in a virtual disk bad block.
RAID 6	Partially degraded (one failed/missing physical disk)	One bad block on a physical disk.	The controller regenerates data from the peer disks and sends a Write to the bad block. The disk then remaps the LBA to another physical location. The problem is resolved.
RAID 6	Degraded (two failed/missing physical disks)	One bad block on a physical disk.	The controller cannot regenerate data from the peer disks. This results in a virtual disk bad block.
RAID 6	Ready	One bad block on a physical disk.	The controller regenerates data from peer disks and sends a Write to the bad block. The disk then remaps the Logical Block Addressing (LBA) to another physical location. The problem is resolved.

Recommendations For Clearing Bad Blocks

Storage Management allows you to clear the bad block warnings. To clear bad blocks, the following procedure is recommended:

1. Perform a backup of the virtual disk with the **Verify** option selected.
One of the following scenarios may occur:
 - Backup operation fails on one or more files. In this case, restore the file from a previous backup. After restoring the file, proceed to step 2.
 - Backup operation completes without error. This indicates that there are no bad blocks on the written portion of your virtual disk.

If you still receive bad block warnings, the bad blocks are in a non-data area. Proceed to step 2.

2. Run **Patrol Read** and check the system event log to ensure that no new bad blocks are found.
If bad blocks still exist, proceed to step 3. If not, the condition is cleared without the need for step 3.
3. To clear these bad blocks, execute the **Clear Virtual Disk Bad Blocks** task.

This **Clear Virtual Disk Bad Block** feature is applicable to PERC H330 Adapter, PERC H330 Mini Monolithic, PERC H730 Adapter, PERC H730 Mini Monolithic, PERC H730P Adapter, PERC H730P Mini Monolithic, PERC H740P Adapter, PERC H740P Mini Monolithic, and PERC H840 Adapter.

Virtual Disk Properties And Tasks

Use the virtual disk properties and tasks window to view information about the virtual disks and execute virtual disk tasks.

From the Options menu:

Refresh Partitions: Click to refresh the Virtual Disk Partitions data once the partition is created using the operating system commands for the available virtual disks.

Go to the Create Virtual Disk Wizard: Launches the Create Virtual Disk Wizard.




Related links



- [Virtual Disk Properties](#)
- [Virtual Disk Tasks](#)
- [Virtual Disk — Available Tasks](#)

Virtual Disk Properties

The virtual disk properties can vary depending on the model of the controller.

Table 34. Virtual Disk Properties


Property	Definition
Status	<p>These icons represent the severity or health of the storage component.</p> <ul style="list-style-type: none"> — Normal/OK — Warning/Non-critical — Critical/Failure/Error <p>For more information, see Storage Component Severity.</p>
Name	Displays the virtual disk name.
State	<p>Displays the status of the virtual disk. Possible values are:</p> <ul style="list-style-type: none">Ready — The virtual disk is functioning normally.Degraded — A physical disk in a redundant virtual disk is not online.Resynching — A consistency check is being performed on the virtual disk. Performing a Cancelling A Check Consistency on a virtual disk while it is in Resynching state causes the virtual disk to be in a Failed Redundancy state.Resynching Paused — A consistency check has been paused on the virtual disk.Regenerating — A physical disk in the virtual disk is rebuilding.Reconstructing — The virtual disk configuration has changed. The physical disks included in the virtual disk are being modified to support the new configuration.Failed — The virtual disk has encountered a failure of one or more components and is no longer functioning.Failed Redundancy — This state is displayed when the initial consistency check for the virtual disk is canceled or is not successful. This state may also be displayed when a RAID 1, RAID 10 or RAID 1-concatenated virtual disk suffers a physical disk failure. In addition, there are other conditions related to disk failures and the firmware that can cause a virtual disk to display a Failed Redundancy state. When a virtual disk is in Failed Redundancy state, performing a Check Consistency may return the virtual disk to a Ready state.Background Initialization — A background initialization is being performed on the virtual disk.Formatting — The virtual disk is being formatted. For more information, see Format, Initialize, Slow, And Fast Initialize.Initializing — The virtual disk is being initialized. For more information, see Format, Initialize, Slow, And Fast Initialize. On some controllers, the virtual disk state is not updated until the controller performs an I/O operation. For more information, see I/O and Reboot Requirements for Detecting Physical Disk Status Changes.Degraded Redundancy — This state is applicable to RAID 6 and RAID 60 in which a physical disk in a redundant virtual disk is not online, but the virtual disk is still accessible and functioning.
Partitions	This property displays whether the virtual disk has a partition. The possible values are Not Available and Link to the Partitions information page.
Virtual Disk Bad Block	Displays whether the virtual disk has bad blocks.
Encrypted	Displays whether the virtual disk is encrypted. The possible values are Yes and No .

Property	Definition
Hot Spare Policy Violated	Displays whether the Hot Spare Protection Policy has been violated.  NOTE: This property is displayed only if you set any Hot Spare Protection Policy. For more information, see Setting Hot Spare Protection Policy.
Layout	Displays the RAID level.
Size	Displays the total capacity of the virtual disk. The algorithm for calculating the virtual disk size rounds a value of 0.005 or less, down to 0.00 and a value between 0.006 and 0.009, up to 0.01. For example, a virtual disk size of 819.725 is rounded down to 819.72. A virtual disk size of 819.726 is rounded up to 819.73.
Device Name	Displays the operating system device name for this object.
Bus Protocol	Displays the technology that the physical disks included in the virtual disk are using. Possible values are: <ul style="list-style-type: none">· SAS· SATA
Media	Displays the media type of the physical disks present in the virtual disk. The possible values are: <ul style="list-style-type: none">· HDD· SSD· Unknown — Storage Management is unable to determine the media type of the physical disk.  NOTE: You cannot have a mix of HDD and SSD media on a virtual disk. Also, you cannot have a mix of SAS and SATA drives on the virtual disk.
Read Policy	Displays the read policy that the controller is using for the selected virtual disk. See RAID Controller Read, Write, Cache, and Disk Cache Policy .
Write Policy	Displays the write policy that the controller is using for the selected virtual disk. See RAID Controller Read, Write, Cache, and Disk Cache Policy .
Cache Policy	Displays the cache policy that the controller is using for the selected virtual disk. See RAID Controller Read, Write, Cache, and Disk Cache Policy .
Stripe Size	Displays the stripe size of the virtual disk.
Disk Cache Policy	Displays whether the disk cache policy of the physical disks that is part of the virtual disk is enabled or disabled. See RAID Controller Read, Write, Cache, and Disk Cache Policy .

Virtual Disk Tasks

To execute a virtual disk drop-down menu task:

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Expand a Controller object.
3. Select the Virtual Disks object.
4. Select a task from the **Available Tasks** drop-down menu.
5. Click **Execute**.

 **NOTE: Different controllers support different features. The tasks displayed on the Available Tasks drop-down box can vary depending on which controller is selected in the tree view. If no tasks can be performed because of controller or system configuration limitations, then the Available Tasks drop-down box displays No Task Available.**

Virtual Disk — Available Tasks

The following is a list of possible tasks in the virtual disk **Available Tasks** drop-down box.

- [Reconfiguring A Virtual Disk](#)
- [Format, Initialize, Slow, And Fast Initialize](#)
- [Canceling Background Initialization](#)
- [Restoring Dead Segments](#)
- [Deleting Data On the Virtual Disk](#)
- [Assigning And Unassigning Dedicated Hot Spare](#)
- [Performing A Check Consistency](#)
- [Canceling A Check Consistency](#)
- [Pausing A Check Consistency](#)
- [Resuming A Check Consistency](#)
- [Blinking And Unblinking A Virtual Disk](#)
- [Renaming A Virtual Disk](#)
- [Split Mirror](#)
- [Unmirror](#)
- [Canceling A Rebuild](#)
- [Changing The Virtual Disk Policy](#)
- [Replacing A Member Disk](#)
- [Clearing Virtual Disk Bad Blocks](#)
- [Encrypting A Virtual Disk](#)

Reconfiguring A Virtual Disk

Does my controller support this feature? See [Supported Features](#).

The **Reconfigure** task enables you to change the virtual disks properties. For example, you can use this task to add physical disks or change the RAID level.

Related links

[Virtual Disk Task: Reconfigure \(Step 1 of 3\)](#)

[Virtual Disk Task: Reconfigure \(Step 2 of 3\)](#)

[Virtual Disk Task: Reconfigure \(Step 3 of 3\)](#)

Format, Initialize, Slow, And Fast Initialize

Does my controller support this feature? See [Supported Features](#).

The **Format, Initialize, Slow Initialize, or Fast Initialize** task enables you to erase the files and remove the file systems on a virtual disk. Some controllers require that you initialize a virtual disk before it can be used.

Related links

[Slow And Fast Initialize](#)

[Considerations For Slow Initialize](#)

Canceling Background Initialization

Does my controller support this feature? See [Supported Features](#).

The **Cancel Background Initialization** task enables you to cancel a background initialization on a virtual disk.

On PERC controllers, background initialization of redundant virtual disks begins automatically after the virtual disk is created. Because the initialization is run in the background, other processes can continue while the initialization completes.

The background initialization of a redundant virtual disk prepares the virtual disk for parity information and improves write performance. It is important that the background initialization is allowed to run to completion. You can, however, cancel the background initialization. When you do so, the controller restarts the background initialization later.

Related links

[Background Initialization On PERC Controllers](#)

Restoring Dead Segments

Does my controller support this feature? See [Supported Features](#).

The **Restore Dead Segments** task recovers data from a RAID-5 virtual disk that has been corrupted. The **Restore Dead Segments** task attempts to reconstruct data from the corrupt portion of a physical disk included in a RAID-5 virtual disk. The **Restore Dead Segments** task uses the parity or redundant information to reconstruct the lost data. This task is not always able to recover lost data.

Deleting Data On the Virtual Disk

Does my controller support this feature? See [Supported Features](#).

The **Delete** task destroys all data on the virtual disk.

Related links

[Deleting A Virtual Disk](#)

Performing A Check Consistency

Does my controller support this feature? See [Supported Features](#).

The **Check Consistency** task verifies the accuracy of the redundant (parity) information. This task only applies to redundant virtual disks. When necessary, the **Check Consistency** task rebuilds the redundant data. If the virtual disk is in a Failed Redundancy state, running a check consistency may be able to return the virtual disk to a Ready state.

Related links

[Setting The Check Consistency Rate](#)

Canceling A Check Consistency


Does my controller support this feature? See [Supported Features](#).

The **Cancel Check Consistency** task stops a check consistency operation while it is in progress.

Pausing A Check Consistency

Does my controller support this feature? See [Supported Features](#).

The **Pause Check Consistency** task pauses a check consistency while it is in progress.

 **NOTE: The Pause Check Consistency task updates the virtual disk State property to Resyncing Paused immediately. The Progress property may continue to increment for up to three seconds. This time delay occurs because the polling task may take up to three seconds to query the task information and update the display.**

Resuming A Check Consistency

Does my controller support this feature? See [Supported Features](#).

The **Resume Check Consistency** task resumes a check consistency after it has been paused.

Blinking And Unblinking A Virtual Disk

Does my controller support this feature? See [Supported Features](#).

The **Blink** and **Unblink** tasks blink or unblink the lights on the physical disks included in the virtual disk.

 **NOTE: The blink operation is available only for removable disks.**

Renaming A Virtual Disk

Does my controller support this feature? See [Supported Features](#).

The **Rename** task enables you to change name of the virtual disk.


Related links

[Renaming A Virtual Disk](#)

Canceling A Rebuild

Does my controller support this feature? See [Supported Features](#).

Select the **Cancel Rebuild** task to cancel a rebuild that is in progress. If you cancel a rebuild, the virtual disk remains in a **Degraded** state. The failure of an additional physical disk can cause the virtual disk to fail and may result in data loss. It is recommended that you rebuild the failed physical disk as soon as possible.

 **NOTE: If you cancel the rebuild of a physical disk that is assigned as a hot spare, reinitiate the rebuild on the same physical disk in order to restore the data. Canceling the rebuild of a physical disk and then assigning another physical disk as a hot spare does not cause the newly assigned hot spare to rebuild the data. Reinitiate the rebuild on the physical disk that was the original hot spare.**

Changing The Virtual Disk Policy

Does my controller support this feature? See [Supported Features](#).

The **Change Policy** task enables you to change the read, write, or cache policy of a virtual disk. Changes to the read, write, and cache policy only apply to the virtual disk that you have selected. This task does not change the policy for all virtual disks on the controller.

Related links

[RAID Controller Read, Write, Cache, And Disk Cache Policy](#)

Replacing A Member Disk

Does my controller support this feature? See [Supported Features](#).

The **Replace Member Disk** task allows you to copy data from a physical disk, which is a member of a virtual disk, to another physical disk by providing the **Replace Member Configuration** option. You can initiate multiple copies of data from different array groups.

Related links

[Virtual Disk Task: Replace Member Disk \(Step 1 of 2\)](#)

[Virtual Disk Task: Replace Member Disk \(Step 2 of 2\)](#)

Clearing Virtual Disk Bad Blocks

Does my controller support this feature? See [Supported Features](#).

The **Clear Virtual Disk Bad Blocks** task enables you to clear bad blocks on your virtual disk. This feature is applicable only to P PERC H330 Adapter, PERC H330 Mini Monolithic, PERC H730 Adapter, PERC H730 Mini Monolithic, PERC H730P Adapter, PERC H730P Mini Monolithic, PERC H740P Adapter, PERC H740P Mini Monolithic, and PERC H840 Adapter.

Encrypting A Virtual Disk

Does my controller support this feature? See [Supported Features](#).

The **Encrypt Virtual Disk** task encrypts an unencrypted virtual disk. This feature is only applicable to controllers that:

- Are encryption capable (PERC H730 Adapter, PERC H730 Mini Monolithic, PERC H730P Adapter, PERC H730P Mini Monolithic, PERC H740P Adapter, PERC H740P Mini Monolithic, and PERC H840 Adapter)
- Have an Encryption Key

- Have Self Encryption Drives (SEDs) virtual disks

Create Virtual Disk Express Wizard

Does my controller support this feature? See [Supported Features](#).

The **Create Virtual Disk Wizard** allows you to select the wizard type and the RAID level. The **Create Virtual Disk Express Wizard** calculates the appropriate virtual disk configuration based on the available space and controller considerations. To make your own selections for the virtual disk configuration, select the **Advanced Wizard** option.

Before creating a virtual disk, you should be familiar with the information in [Considerations Before Creating Virtual Disks](#). You may also want to review [Choosing RAID Levels And Concatenation](#).

 **NOTE: This task is not supported on PERC hardware controllers running in HBA mode.**

To create a virtual disk using the **Create Virtual Disk Express Wizard**:

1. In the Server Administrator window, expand **System** → **Storage**.
2. Click **Connector 0 (RAID)**.
The **Virtual Disk(s) on Controller <Controller Name>** page is displayed.
3. Click **Go to the Create Virtual Disk Wizard**.
The **Create Virtual Disk Wizard (Step 1) <Controller Name>** page is displayed.
4. Select the **Express Wizard** option and the **RAID level** from the drop-down list.
 - Depending on the controller, **Concatenated** enables you to combine the storage capacity of several disks or to create a virtual disk using a single physical disk. For more information about the number of supported physical disks when using **Concatenated**, see [Number Of Physical Disks Per Virtual Disk](#). **Concatenated** does not provide data redundancy or affect the read/write performance.
 - Select **RAID 0** for striping. This selection groups n disks together as a large virtual disk with a total capacity of n disks. Data is alternately stored in the disks so that they are evenly distributed. Data redundancy is not available in this mode. Read/write performance is enhanced.
 - Select **RAID 1** for mirroring disks. This selection groups two disks together as one virtual disk with a capacity of one single disk. The data is replicated on both disks. When a disk fails, the virtual disk continues to function. This feature provides data redundancy and good read performance, but slightly slower write performance. The system must have at least two physical disks to use RAID 1.
 - Select **RAID 5** for striping with distributed parity. This selection groups n disks together as one large virtual disk with a total capacity of $(n-1)$ disks. When a disk fails, the virtual disk continues to function. This feature provides better data redundancy and read performance, but slower write performance. The system must have at least three physical disks to use RAID 5.
 - Select **RAID 6** for striping with additional distributed parity. This selection groups n disks as one large virtual disk with a capacity of $(n-2)$ disks. The virtual disk remains functional with up to two disk failures. RAID 6 provides better read performance, but slower write performance. The system must have at least four physical disks to use RAID 6.
 - Select **RAID 10** for striping over mirror sets. This selection groups n disks together as one large virtual disk with a total capacity of $(n/2)$ disks. Data is striped across the replicated mirrored pair disks. When a disk fails, the virtual disk continues to function. The data is read from the surviving mirrored pair disk. This feature provides the best failure protection, read and write performance. The system must have at least four physical disks to use RAID 10.
 - Select **RAID 50** to implement striping across more than one span of physical disks. RAID 50 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.
 - Select **RAID 60** to implement striping across more than one RAID 6 span. RAID 60 Groups $n*s$ disks as one large virtual disk with a capacity of $s*(n-2)$ disks, where s is the number of spans and n is the number of disks within each span. RAID 60 provides increased data protection and better read performance, but slower write performance.
5. Click **Continue**.

Related links

[Creating A Virtual Disk](#)

[Create Virtual Disk Advanced Wizard](#)

Create Virtual Disk Express Wizard (Step 2)

The **Create Virtual Disk Express Wizard - <Controller Name>** page displays the summary of attributes of the selected RAID level. The attributes include **Bus Protocol**, **Stripe Element Size**, **Read Policy**, and the selected physical disk. The default values of the virtual disk attributes excluding the RAID level are recommended for the selected RAID level.

1. In the **Name** field, type a name for the virtual disk.

The virtual disk name can contain only alphanumeric characters, spaces, dashes, and underscores. The maximum length depends on the controller. In most cases, the maximum length is 15 characters. The name cannot start with a space or end with a space.

 **NOTE: It is recommended that you specify a unique name for each virtual disk. If you have virtual disks with the same name, it is hard to differentiate the generated alerts.**

 **NOTE: If a physical disk is receiving a SMART alert, it cannot be used in a virtual disk. For more information on SMART alerts, see [Monitoring Disk Reliability On RAID Controllers](#).**

2. In the **Size** field, type the size of the virtual disk.

The virtual disk size must be within the minimum and maximum values displayed near the **Size** field. For information on how the maximum virtual disk size is calculated, see [Calculation For Maximum Virtual Disk Size](#).

In some cases, the virtual disk is slightly larger than the size you specify. The **Create Virtual Disk Wizard** adjusts the size of the virtual disk to avoid rendering a portion of the physical disk space unusable.

3. Click **Finish** to complete the virtual disk creation.

For PERC S140 controller, if removable and non-removable devices are used to create a virtual disk, the following message is displayed:

```
A combination of removable and non-removable NVMe devices have been used to create this virtual disk. Are you sure you want to proceed?
```

Click **Finish** to create the virtual disk.

Click **Go Back To Previous Page** to return to **Create Virtual Disk Express Wizard** page, if you want to change your selections.

Click **Exit Wizard** to cancel the virtual disk creation.

The virtual disk is displayed in the **Virtual Disk(s) on Controller <Controller Name>** page.

Create Virtual Disk Advanced Wizard

Does my controller support this feature? See [Supported Features](#).

The **Create Virtual Disk Advanced Wizard** allows you to specify the read, write, and cache policy, and parameters such as RAID level, bus protocol, media type, and encrypted disk for the virtual disk. You can also select the physical disks and the controller connector. You must have a good knowledge of RAID levels and hardware, to use the Advanced Wizard. If you want the wizard to choose a recommended virtual disk configuration for you, select the **Express Wizard** option.

Before creating a virtual disk, you should be familiar with the information in [Considerations Before Creating Virtual Disks](#). You may also want to review [Choosing RAID Levels And Concatenation](#).

 **NOTE: This task is not supported on PERC hardware controllers running in HBA mode.**

To create a virtual disk using the **Create Virtual Disk Advanced Wizard**:

1. On the left-hand side of the **Server Administrator** page, expand **Storage**.
2. Click on the **<PERC Controller>**.
3. Click on **Virtual Disks**.
The **Virtual Disk(s) on Controller <Controller Name>** page is displayed.
4. Click **Go to the Create Virtual Disk Wizard**.
The **Create Virtual Disk Wizard (Step 1) <Controller Name>** page is displayed.
5. Select the **Advanced Wizard** option.

6. To make sure that only encrypted physical disks are used to create the virtual disk, select **Yes** from the **Create Encrypted Virtual Disk** drop-down list.

The RAID levels are available for selection based on the number of encrypted physical disks.

If you select **No**, the RAID levels are available based on the total number of physical disks present on the system.

7. Select the required RAID level from the drop-down list.

- Depending on the controller, **Concatenated** enables you to combine the storage capacity of several disks or to create a virtual disk using a single physical disk. For more information on the number of supported physical disks supported by **Concatenated**, see [Number Of Physical Disks Per Virtual Disk](#). Using **Concatenated** does not provide data redundancy or affect the read/write performance.
- Select **RAID 0** for striping. This selection groups n disks together as a large virtual disk with a total capacity of n disks. Data is alternately stored in the disks so that they are evenly distributed. Data redundancy is not available in this mode. Read/write performance is enhanced.
- Select **RAID 1** for mirroring disks. This selection groups two disks as one virtual disk with the capacity of a single disk. The data is replicated on both disks. When a disk fails, the virtual disk continues to function. This feature provides data redundancy and good read performance, but slightly slower write performance. The system must have at least two physical disks to use RAID 1.
- Select **RAID 5** for striping with distributed parity. This selection groups n disks together as one large virtual disk with a capacity of $(n-1)$ disks. When a disk fails, the virtual disk continues to function. This feature provides better data redundancy and read performance, but slower write performance. The system must have at least three physical disks to use RAID 5.
- Select **RAID 6** for striping with additional distributed parity. This selection groups n disks as one large virtual disk with a capacity of $(n-2)$ disks. The virtual disk remains functional with up to two disk failures. RAID 6 provides better read performance, but slower write performance. The system must have at least four physical disks to use RAID 6.
- Select **RAID 10** for striping over mirror sets. This selection groups n disks together as one large virtual disk with a capacity of $(n/2)$ disks. Data is striped across the replicated mirrored pair disks. When a disk fails, the virtual disk continues to function. The data is read from the surviving mirrored pair disk. This feature provides the best failure protection and read/write performance. The system must have at least four physical disks to use RAID 10. For PERC controllers with firmware version 6 and later, RAID 10 also allows you to create a single span virtual disk with 22 or 26 physical drives.
- **Intelligent Mirroring** — Automatically calculates the span composition based on the physical disks you select.

Spans are not displayed on this page. Select **Continue** to view the span grouping on the Summary page. Click **Exit Wizard** to cancel the virtual disk creation.

Storage Management calculates the optimum span composition as follows:

- Determining span calculation:
 - * Calculating the number of disks that can be utilized from the selected disks.
 - * Maximizing the number of spans to increase the I/O performance
- Determining the mirror for the physical disks: The mirror is determined in a way that ensures maximum redundancy. The algorithm also tries to match a physical disk for its mirror, to a disk that is almost the same size. However, **Intelligent Mirroring** gives priority to size over redundancy.

The algorithm determines the candidate mirror in the following order:

- * Across connectors at the same level of enclosure and same size.
- * Across connectors in the enclosure that are not at the same level but of same size.
- * Across enclosures connected to the same connector and to a disk of the same size.
- * Within the enclosure with a physical disk of acceptable size difference.
- * Across connectors at the same level of enclosure and of acceptable size difference.
- * Across connectors in the enclosure that are not at the same level of the enclosure but with a physical disk of acceptable size difference.

If the size difference is not acceptable, the disk is not mirrored and hence dropped from the span. The number of spans and disks in the span is recalculated.

 **NOTE: It is recommended that you use Intelligent Mirroring to create RAID 10 across enclosures for simple and optimum configuration.**

 **NOTE: To view the redundancy across enclosures achieved through Intelligent Mirroring, click the virtual disk and view the physical disk IDs in each span, which are from alternate enclosures.**

- Select **RAID 50** to implement striping across more than one span of physical disks. RAID 50 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.

- Select **RAID 60** to implement striping across more than one RAID 6 span. RAID 60 groups $n*s$ disks as one large virtual disk with a capacity of $s*(n-2)$ disks, where s is the number of spans and n is the number of disks within each span. RAID 60 provides increased data protection and better read performance, but slower write performance.¹

8. Select the Bus Protocol.

The possible options are:

- SAS
- SATA

9. Select the Media Type.

The available options are:

- HDD
- SSD

For more information on **Bus Protocol** and **Media Type**, see [Virtual Disk Properties And Tasks](#).

10. Click Continue.

The **Create Virtual Disk Advanced Wizard (Step 2 of 3) — <Controller Name>** page is displayed with the connectors and physical disks available.

Related links

[Create Virtual Disk Express Wizard](#)

Create Virtual Disk Advanced Wizard (Step 2)

Depending on the RAID level you selected in the **Create Virtual Disk Wizard (Step 1) <Controller Name>** page, the **Create Virtual Disk Advanced Wizard (Step 2 of 3) - <Controller Name>** page displays the disks and connectors (channels or ports) available for configuring the virtual disk.

 **NOTE: If you are creating an encrypted virtual disk, only the encrypted physical disks are displayed. Else, both encrypted and unencrypted physical disks are displayed.**

The following is an example of what might be displayed:

Connector 0

The **Connector** section of the page displays the connectors on the controller and the disks attached to each connector. Select the disks you want to include in the virtual disk. In this example, the controller has a single connector with five disks.

- Physical disk 0:0
- Physical disk 0:1
- Physical disk 0:2
- Physical disk 0:3
- Physical disk 0:4

Physical Disks Selected


The **Physical Disks Selected** section of the page displays the disks you have chosen. In this example, two disks are selected.

- Physical disk 0:0
- Physical disk 0:1

Each RAID level has specific requirements for the number of disks that must be selected. RAID 10, RAID 50, and RAID 60 also have requirements for the number of disks that must be included in each stripe or span.

If the controller is a SAS controller with firmware versions 6.1 and later and you selected RAID 10, RAID 50, and RAID 60, the user interface displays the following:

- **All Disks** — Enables you to select all the physical disks in all the enclosures.
- **Enclosure** — Enables you to select all physical disks in the enclosure.

 **NOTE: The All Disks and Enclosure options enable you to edit spans after selecting the physical disks that comprise them. You can remove a span and specify a span again with different physical disks before proceeding.**

- **Number of Disks per Span** — Enables you to select the number of disks in each span (default = 2). This option is available only on SAS controllers with firmware version 6.1 and later.

 **NOTE: This option is available only if the Intelligent Mirroring on the Create Virtual Disk Advanced Wizard (Step 2 of 3) page.**

 **NOTE: On a SAS controller with firmware version 6.1 and later, RAID 10 supports only even number of disks per span and a maximum of 8 spans with 32 disks in each span.**

Let us consider that the controller has three enclosures with six physical disks each (total number of available disks = 3 x 6 = 18 disks). If you select four disks per span, the controller creates four spans (18 disks/4 disks per span = 4 spans). The last two disks of the last enclosure are not part of RAID 10.

- Select the **number of disks to create a single spanned virtual disk** list box — Enables you to create a single span virtual disk with 22 or 26 physical drives for PERC controllers. This list box option appears only if you have selected RAID 10 in step 1 and the system has 22 or more physical drives.

 **NOTE: Only physical disks that comply with the virtual disk parameters, selected in the Create Virtual Disk Wizard page are displayed.**

Select the required connector and the corresponding physical disk, and click **Continue**.

Create Virtual Disk Advanced Wizard (Step 3)

The **Create Virtual Disk Advanced Wizard (Step 3 of 3) - <Controller Name>** page displays the summary of attributes of the selected RAID level. The attributes include **Bus Protocol, Stripe Element Size, Read Policy** and the selected physical disk. The default values of the virtual disk attributes excluding the RAID level are recommended for the selected RAID level.

1. In the **Name** field, type a name for the virtual disk.

The virtual disk name can contain only alphanumeric characters, spaces, dashes, and underscores. The maximum length depends on the controller. In most cases, the maximum length is 15 characters.


The name cannot start with a space or end with a space. It is recommended that you specify a unique name for each virtual disk. If you have virtual disks with the same name, it is hard to differentiate the generated alerts.

2. In the **Size** field, type the size of the virtual disk.


The virtual disk size must be within the minimum and maximum values displayed near the **Size** field.

In some cases, the virtual disk is slightly larger than the size you specify. The **Create Virtual Disk Wizard** adjusts the size of the virtual disk to avoid rendering a portion of the physical disk space unusable.

 **NOTE: If a physical disk is receiving a SMART alert, it cannot be used in a virtual disk.**

 **NOTE: For a controller that has more than one channel, it may be possible to configure a virtual disk that is channel-redundant.**

Depending on the RAID level you selected and the virtual disk size, this page displays the disks and connectors (channels or ports) available for configuring the virtual disk.

 **NOTE: The Size field displays the default size of the virtual disk depending upon the RAID configuration you selected. You can specify a different size. The virtual disk size must be within the minimum and maximum values displayed near the Size field. In some cases, the virtual disk is slightly larger than the size you specify. The Create Virtual Disk Wizard adjusts the size of the virtual disk to avoid rendering a portion of the physical disk space unusable.**

3. Select a stripe size from the **Stripe Element Size** drop-down list box. The stripe size refers to the amount of space that each stripe consumes on a single disk.
4. Select the required read, write, and disk cache policy from the respective drop-down list box. These selections can vary depending on the controller.

 **NOTE: There is limited support for write policy on controllers that do not have a battery. The cache policy is not supported on controllers that do not have a battery.**

 **NOTE: If you have selected the Create Encrypted Virtual Disk option in Create Virtual Disk Advanced Wizard, then in the Summary of Virtual Disk Attributes, an Encrypted Virtual Disk attribute is displayed with a value Yes.**

Click **Span Edit** to edit the spans created in [Create Virtual Disk Advanced Wizard](#).

The **Span Edit** option is available only if the controller is a SAS controller with firmware 6.1 and later and if you selected RAID 10.

 **NOTE: If you click Span Edit, Intelligent Mirroring that has already been applied becomes invalid.**

5. Click **Finish** to complete the virtual disk creation.

The virtual disk is displayed in the **Virtual Disk(s) on Controller <Controller Name>** page.

Click **Go Back To Previous Page** to return to **Create Virtual Disk Advanced Wizard** page, if you want to change your selections.

Click **Exit Wizard** to cancel the virtual disk creation.

The **Create Virtual Disk Advanced Wizard (Step 3 of 3) - <Controller Name>** page displays a check box next to each physical disk that is suitable as a dedicated hot spare. Select a **Physical Disk** check box if you want to assign a dedicated hot spare

The **Physical Disk** check box is not available if the controller does not have a physical disk that is a suitable hot spare for the virtual disk you are creating. For example, the available physical disks may be too small to protect the virtual disk. If the **Physical Disk** check box is not available, you may need to specify a smaller virtual disk, use a different RAID level, or change the disk selection.

 **NOTE: If you are creating an encrypted virtual disk, then encrypted physical disks are listed as candidates for hot spare.**

Related links

[Calculation For Maximum Virtual Disk Size](#)

[Monitoring Disk Reliability On RAID Controllers](#)

[Channel Redundancy](#)

[What is RAID?](#)

[RAID Controller Read, Write, Cache, And Disk Cache Policy](#)

[Write Policy](#)

Span Edit

In the edit mode, you cannot alter the number of physical disks per span. If there are enough available physical disks, you can reduce or increase the number of spans. You can also alter the contents of a span by removing the span and selecting a new physical disk to comprise that span.

To successfully create a virtual disk, a minimum of two spans must exist always.

 **NOTE: If you select Span Edit, Intelligent Mirroring that is already been applied becomes invalid.**

Virtual Disk Task: Reconfigure (Step 1 of 3)


Does my controller support this feature? See [Supported Features](#).

The **Reconfigure** task enables you to change the virtual disk configuration. Using this task, you can change the RAID level and increase the virtual disk size by adding physical disks. On some controllers, you can also remove physical disks.

Before continuing with the virtual disk reconfiguration, you should be familiar with the information in [Starting And Target RAID Levels For Virtual Disk Reconfiguration And Capacity Expansion](#) and [Choosing RAID Levels And Concatenation](#).

 **NOTE: You cannot reconfigure a virtual disk on a controller that is operating in cluster mode.**

 **NOTE: You can create no more than 64 virtual disks on a controller. After you have reached this limit, you can no longer reconfigure any of the virtual disks on the controller.**

 **NOTE: On systems running the Linux operating system, if you reconfigure the same controller on which the operating system resides, you may experience extremely slow system performance until the reconfiguration is complete.**

 **NOTE: You may want to review [Virtual Disk Considerations For Controllers](#). This section contains considerations that also apply to reconfiguring a virtual disk on these controllers.**


Related links


- [Reconfiguring A Virtual Disk](#)
- [What is RAID?](#)
- [Deleting A Virtual Disk](#)
- [Channel Redundancy](#)
- [Reconfiguring Or Migrating Virtual Disks](#)
- [Physical Disk Or Physical Device Properties](#)
- [Physical Disk Or Physical Device Tasks](#)
- [Create Virtual Disk Advanced Wizard](#)
- [Create Virtual Disk Express Wizard](#)

To Reconfigure A Virtual Disk: Step 1 of 3

1. Select the physical disks that you want to include in the virtual disk. You can expand the capacity of the virtual disk by adding additional physical disks. On some controllers, you can also remove physical disks.

The changes you make to the physical disk selection are displayed in the **Selected Physical Disks** table.

 **NOTE:** For a controller that has more than one channel, it may be possible to configure a virtual disk that is channel-redundant.

 **NOTE:** For PERC H330 Adapter, PERC H330 Mini Monolithic, PERC H730 Adapter, PERC H730 Mini Monolithic, PERC H730P Adapter, PERC H730P Mini Monolithic, PERC H740P Adapter, PERC H740P Mini Monolithic, and PERC H840 Adapter controllers, if any of the selected physical drives are in the spun down state, the system displays a message that indicates the IDs of the spun down drives and the delay in executing tasks on them.

 **NOTE:** For PERC H330 Adapter, PERC H330 Mini Monolithic, PERC H730 Adapter, PERC H730 Mini Monolithic, PERC H730P Adapter, PERC H730P Mini Monolithic, PERC H740P Adapter, PERC H740P Mini Monolithic, and PERC H840 Adapter if the disk group has free space available, you can expand the virtual disk capacity. To expand virtual disk capacity, click **Expand Capacity**.

2. Click **Continue** to go to the next page or **Exit Wizard** if you want to cancel.

Related links

- [Channel Redundancy](#)

To locate Reconfigure In Storage Management

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Expand a controller object.
3. Select the **Virtual Disks** object.
4. Select **Reconfigure** from the **Available Tasks** drop-down box.
5. Click **Execute**.

Virtual Disk Task: Reconfigure (Step 2 of 3)

Does my controller support this feature? See [Supported Features](#).

The virtual disk reconfigure Step 2 page enables you to select the RAID level and size for the reconfigured virtual disk.

If you selected **Expand Capacity** in [Step 1](#), this page allows you to expand the capacity of the virtual disk.

To Reconfigure a Virtual Disk (Changing the RAID Level and Size): Step 2 of 3

1. Select the new RAID level for the virtual disk — The available RAID levels depend on the number of physical disks selected and the controller. The following describes the possible RAID levels:
 - Depending on the controller, Concatenated enables you to combine the storage capacity of several disks or to create a virtual disk using only a single physical disk. For information on whether the controller supports a single physical disk or two or more when using **Concatenated**, see [Number Of Physical Disks Per Virtual Disk](#). Using Concatenated does not provide data redundancy nor does it affect the read and write performance.

- Select **RAID 0** for striping — This selection groups **n** disks together as one large virtual disk with a total capacity of **n** disks. Data is stored to the disks alternately so that they are evenly distributed. Data redundancy is not available in this mode. Read and write performance is enhanced.
- Select **RAID 1** for mirroring disks — This selection groups two disks together as one virtual disk with a capacity of one single disk. The data is replicated on both disks. When a disk fails, the virtual disk continues to function. This feature provides data redundancy and good read performance, but slightly slower write performance. The system must have at least two physical disks to use RAID 1.
- Select **RAID 5** for striping with distributed parity — This selection groups **n** disks together as one large virtual disk with a total capacity of **(n-1)** disks. When a disk fails, the virtual disk continues to function. This feature provides better data redundancy and read performance, but slower write performance. The system must have at least three physical disks to use RAID 5.
- Select **RAID 6** for striping with additional parity information — This selection groups **n** disks as one large virtual disk with a capacity of **(n-2)** disks. Two sets of parity information are alternately stored on all disks. The virtual disk remains functional with up to two disk failures.
- Select **RAID 10** for striping over mirror sets — This selection groups **n** disks together as one large virtual disk with a total capacity of **(n/2)** disks. Data is striped across the replicated mirrored pair disks. When a disk fails, the virtual disk continues to function. The data is read from the surviving mirrored pair disk. This feature provides the best failure protection, read and write performance. The system must have at least four physical disks to use RAID 10.

 **NOTE:**

- RAID 10 virtual disk reconfiguration is applicable only for PERC 9 family of hardware controllers with the latest firmware version.
- Other RAID level virtual disks cannot be converted to RAID 10 virtual disks.
- RAID 10 virtual disks cannot be converted to any other RAID level virtual disks.
- RAID 10 virtual disk reconfiguration is supported only for even number of physical disks.
- RAID 10 virtual disk reconfiguration is allowed for a maximum of 32 disks.
- RAID 10 virtual disk reconfiguration operation is applicable only for single span with 32 disks.
- RAID 10 virtual disk reconfiguration operation is not supported with Intelligent Mirroring.

 **NOTE: For RAID 10 partial virtual disks, there are only two methods to increase the disk size or capacity:**

- Expanding the capacity by increasing the percentage value.
- Adding disks into the original RAID 10 virtual disk.

 **NOTE: If you create RAID10 with uneven spans in PERC9 controller and then tries to import the RAID10 VD from PERC9 to PERC10, the import fails.**

2. In the **Size** field, type the size of the reconfigured virtual disk. The minimum and maximum allowable size is displayed below the Size field. These values reflect the new capacity of the virtual disk after any addition or deletion of physical disks which you may have chosen in [Virtual Disk Task: Reconfigure \(Step 1 of 3\)](#).
3. Click **Continue** to go to the next page or **Exit Wizard** if you want to cancel.

Related links

[Reconfiguring A Virtual Disk](#)

To Reconfigure A Virtual Disk Expand Virtual Disk Capacity: Step 2 of 3

 **NOTE: For RAID 10 partial virtual disks, there are only two methods to increase the disk size or capacity:**

- Expanding the capacity by increasing the percentage value.
- Adding disks into the original RAID 10 virtual disk.

1. Enter the percentage of the free disk size available by which you want to expand the virtual disk capacity. The page displays the maximum free size available and the description of the selected RAID level.
2. Click **Continue** to go to the next page or click **Exit Wizard** if you want to cancel.

Virtual Disk Task: Reconfigure (Step 3 of 3)

Does my controller support this feature? See [Supported Features](#).

The virtual disk reconfigure step 3 page enables you to review your changes before completing the virtual disk reconfiguration.

1. Review your changes. The **New Virtual Disk Configuration** table displays the changes you have made to the virtual disk. The **Previous Virtual Disk Configuration** displays the original virtual disk prior to reconfiguration.
2. Click **Finish** to complete the virtual disk reconfiguration. To exit without changing the original virtual disk, click **Exit Wizard**.



NOTE: After the virtual disk reconfiguration process starts, the read and write policies temporarily revert to the default settings until the process is completed.

Related links

[Reconfiguring A Virtual Disk](#)

Slow And Fast Initialize

Considerations For Fast Initialize

The **Fast Initialize** task initializes all physical disks included in the virtual disk. The **Fast Initialize** task updates the metadata on the physical disks so that all disk space is available for future write operations. The initialize task can be completed quickly because the existing information on the physical disks is not erased, although future write operations overwrite any information that remains on the physical disks.



CAUTION: Performing a Fast Initialize causes existing data to be inaccessible. This task should be considered data destructive.

In comparison with the **Slow Initialize** task, the **Fast Initialize** task has the following advantages:

- The **Fast Initialize** task takes less time to complete.
- The **Fast Initialize** task does not write zeroes to the disk blocks on the physical disks. It is because the **Fast Initialize** task does not perform a write operation, it causes less degradation to the disk than does the **Slow Initialize** task.

If you have had trouble with a physical disk or suspect that it has bad disk blocks, you may want to perform a **Slow Initialize** task, as this task remaps bad blocks and writes zeroes to all disk blocks.

Related links

[Format, Initialize, Slow, And Fast Initialize](#)

[Considerations For Slow Initialize](#)

[Format, Initialize, Slow, And Fast Initialize](#)

Considerations For Slow Initialize

The **Slow Initialize** task initializes all physical disks included in the virtual disk. The **Slow Initialize** task updates the metadata on the physical disks and erases all existing data and file systems.

In comparison with the **Fast Initialize** task, you may want to use the **Slow Initialize** task if you have had trouble with a physical disk or suspect that it has bad disk blocks. The **Slow Initialize** task remaps bad blocks and writes zeroes to all disk blocks.

The **Slow Initialize** task initializes one physical disk at a time. Each physical disk displays the Clear state while being initialized. While the physical disk is in the Clear state, the **Cancel Clear** physical disk task is available. Performing a **Cancel Clear** task on the physical disk causes the **Slow Initialize** task to be canceled for the entire virtual disk and all member physical disks.

Related links

[Format, Initialize, Slow, And Fast Initialize](#)

Formatting Or Initializing A Disk

To format or initialize a disk:

1. Review the virtual disk that is destroyed by the Format or Initialize and make sure that vital data is not lost. Click **Blink** at the bottom of the page to blink the physical disks included in the virtual disk.
2. Depending on the task you are initiating, click the following option when ready:
 - **Format**

- **Initialize**
- **Slow Initialize**
- **Fast Initialize**

To exit without formatting or initializing the virtual disk, click **Go Back To Virtual Disk Page**.

To Locate Virtual Disks Task In Storage Management

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Expand a Controller object.
3. Select the Virtual Disks object.
4. Depending on the task you want to initiate, select one of the following from the Available Tasks drop-down menu.
 - **Format**
 - **Initialize**
 - **Slow Initialize**
 - **Fast Initialize**
5. Click **Execute**.

Deleting A Virtual Disk

Does my controller support this feature? See [Supported Features](#).

Deleting a virtual disk destroys all information including file systems and volumes residing on the virtual disk.

 **NOTE: When deleting virtual disks, all assigned global hot spares may be automatically unassigned when the last virtual disk associated with the controller is deleted. When deleting the last virtual disk of a disk group, all assigned dedicated hot spares automatically become global hot spares.**

Related links

- [Deleting Data On the Virtual Disk](#)
- [Virtual Disks](#)

To Delete A Virtual Disk

To identify which physical disks are included in the virtual disk, click **Blink**. The LED lights on the physical disks blink for 30-60 seconds.

When deleting a virtual disk, the following considerations apply:

- There are particular considerations for deleting a virtual disk from a cluster-enabled controller.
- It is recommended that you reboot the system after deleting the virtual disk. Rebooting the system ensures that the operating system recognizes the disk configuration correctly.
- 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 recognizes 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.

To Locate Delete In Storage Management

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Expand a Controller object.
3. Select the Virtual Disks object.
4. Select **Delete** from the **Available Tasks** drop-down menu.
5. Click **Execute**.

Renaming A Virtual Disk

Does my controller support this feature? See [Supported Features](#).

Renaming a virtual disk enables you to change the name of the virtual disk. The numbering format for the virtual disk remains unchanged.

Depending on the controller you have, there are different considerations regarding the controller BIOS:

The virtual disk name can contain only alphanumeric characters as well as spaces, dashes and underscores. The maximum length of the name depends on the controller. In most cases, the maximum length is 15 characters. The name cannot start with a space, end with a space, or be left blank.

Related links

[Renaming A Virtual Disk](#)

To Rename A Virtual Disk

1. Type the new name in the text box.
2. Click **Rename**.

To exit without renaming the virtual disk, click **Go Back To Virtual Disk Page**.

To Locate Rename In Storage Management

To locate this task in storage management:

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Expand a controller object.
3. Select the **Virtual Disks** object.
4. Select **Rename** from the **Available Tasks** drop-down menu.
5. Click **Execute**.

Changing The Policy Of A Virtual Disk

Does my controller support this feature? See [Supported Features](#).

Use the **Change Policy** task to change the read, write, or cache policy of a virtual disk.

Related links

[RAID Controller Read, Write, Cache, And Disk Cache Policy](#)

Changing The Read, Write, Or Disk Cache Policy Of A Virtual Disk

1. Select the new policy from the **Read Policy**, **Write Policy**, and **Disk Cache Policy** drop-down menus.
2. Click **Apply Changes**.

To exit without changing the virtual disk policy, click **Go Back To Virtual Disk Page**.


To Locate Change Policy In Storage Management

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Expand a Controller object.
3. Select the Virtual Disks object.
4. Select **Change Policy** from the **Available Tasks** drop-down box.
5. Click **Execute**.

Split Mirror

Does my controller support this feature? See [Supported Features](#).

Use the **Split Mirror** task to separate mirrored data originally configured as a RAID 1, RAID 1-concatenated, or RAID 10 virtual disk. Splitting a RAID 1 or RAID 1-concatenated 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: On Linux operating systems, a Split Mirror cannot be performed on a mounted virtual disk. For this reason, a Split Mirror cannot be performed on the boot drive.**

Splitting A Mirror

Click **Split Mirror**. To exit without splitting the mirror, click **Go Back To Virtual Disk Page**.

 **CAUTION: Your virtual disk is longer redundant after performing a Split Mirror operation.**

To Locate Split Mirror In Storage Management

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Expand a controller object.
3. Select the **Virtual Disks** object.
4. Select **Split Mirror** from the **Available Tasks** drop-down menu.
5. Click **Execute**.


Related links

[Time Delay In Displaying Configuration Changes](#)

Unmirror

Does my controller support this feature? See [Supported Features](#).

Use the **Unmirror** task to separate mirrored data and restore one half of the mirror to free space. Unmirroring a RAID 1 or RAID 1-concatenated virtual disk results in a single, nonredundant concatenated virtual disk. Unmirroring a RAID 10 virtual disk results in a single, nonredundant RAID 0 (striped) virtual disk. Data is not lost during this operation.

 **NOTE: On systems running Linux operating systems, an Unmirror cannot be performed on a mounted virtual disk.**

To Unmirror

Click **Unmirror**. To exit without unmirroring, click **Go Back To Virtual Disk Page**.

 **CAUTION: Your virtual disk is longer redundant.**

To Locate Unmirror In Storage Management

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Expand a controller object.
3. Select the **Virtual Disks** object.
4. Select **Unmirror** from the **Available Tasks** drop-down menu.
5. Click **Execute**.

Related links

[Time Delay In Displaying Configuration Changes](#)

Assigning And Unassigning Dedicated Hot Spare

Does my controller support this feature? See [Supported Features](#).

A dedicated hot spare is an unused backup disk that is assigned to a single virtual disk. When a physical disk in the virtual disk fails, the hot spare is activated to replace the failed physical disk without interrupting the system or requiring your intervention.

For more information on hot spares including size requirements, see [Protecting Your Virtual Disk With A Hot Spare](#).

Related links

[Dedicated Hot Spare Protection Policy](#)

[Assigning A Dedicated Hot Spare](#)

[Unassigning A Dedicated Hot Spare](#)

Assigning A Dedicated Hot Spare

1. Select the disk in the **Connector** (channel or port) table that you want to use as the dedicated hot spare. On some controllers, more than one disk can be selected. The disks you have selected as dedicated hot spares are displayed in the **Disks currently configured as dedicated hot spare** table.
2. Click **Apply Changes** when ready.

Unassigning A Dedicated Hot Spare

1. Select the disk in the **Disks currently configured as dedicated hot spare** table to unassign it. On some controllers, more than one disk can be selected. Clicking the disk removes the disk from the **Disks currently configured as dedicated hot spare** table and returns it to the Connector (channel or port) table.
2. Click **Apply Changes** when ready.

To Locate Assign Or Unassign Dedicated Hot Spare In Storage Management

1. Expand the **Storage** tree object to display the controller objects.
2. Expand a **Controller** object.
3. Select the **Virtual Disks** object.
4. Select **Assign Dedicated Hot Spare** or **Unassign Dedicated Hot Spare** from the **Available Tasks** drop-down menu.
5. Click **Execute**.

Virtual Disk Task: Replace Member Disk (Step 1 of 2)

 **NOTE: This feature is supported only on SAS and SATA controllers with firmware versions 6.1 and later.**

Does my controller support this feature? See [Supported Features](#).

You can copy data from a physical disk, which is a member of a virtual disk, to another physical disk by providing a Replace Member Configuration option. You can initiate multiple copies of data from different array groups.

The source physical disk should be part of a virtual disk and in the Online state. Also, the virtual disk should not be fully degraded.

 **NOTE: You must also enable the Revertible Hot Spare option to use Replace Member Disk task.**

The destination physical disk should be in the Ready state, available to be written to, and of the appropriate size and type.

 **NOTE: The destination physical disk can also be an available hot spare.**

Related links

[Replacing A Member Disk](#)

Replacing A Member Disk: (Step 1 of 2)

1. Select the physical disk in the **Connector** table that you want to replace.
2. Select the destination disk in the **Disks available for replace member operation** table.


 **CAUTION:** If you choose a hot spare as the destination physical disk, your virtual disk does not have a hot spare, unless you assign one.

 **NOTE:** You can select only one source or destination physical disk at a time.

3. Click **Apply Changes**.

To exit without replacing the member disk, click **Go Back To Virtual Disk Page**.

You can view the progress of the **Replace Member Disk** task on the **Physical Disk Details** page.

 **NOTE:** For PERC H330, H730, H730P, H740P and H840 controllers, if any of the drives you selected is in the spun down state, the following message is displayed: The below listed physical drive(s) are in the spun down state. Executing this task on these drive(s) takes additional time, because the drive(s) need to spun up. The message displays the ID(s) of the spun down drive(s).

Related links

[Physical Disk Or Physical Device Properties](#)

[Physical Disk Or Physical Device Tasks](#)

To Locate Replace Member Disk In Storage Management

1. In the **Server Administrator** window, under the system tree, expand **Storage** to display the controller objects.
2. Expand a controller object.
3. Select the **Virtual Disks** object.
4. Select **Replace Member Disk** from the **Available Tasks** drop-down menu.
5. Click **Execute**.

Virtual Disk Task: Replace Member Disk (Step 2 of 2)

This page displays the summary of the attributes of the virtual disk in which you replaced the member disk. Use this page to review your changes before completing the virtual disk replace member task.

To Replace a Member Disk: Step 2 of 2

1. Review your changes. The source **Physical Disk** table displays details of the source physical disk. The destination **Physical Disk** table displays details of the destination physical disk.
2. Click **Finish** to complete the replace member task.

If you want to change the replace member, click **Go Back to Previous Page**. To exit without making changes, click **Cancel**.

Related links

[Replacing A Member Disk](#)

Moving Physical And Virtual Disks From One System To Another

This section describes how to move physical and virtual disks from one system to another.

Related links

[Migrating SAS Virtual Disks To Another System](#)

Required Conditions

SAS Controllers

Virtual disks can only be migrated to a controller that is using the same technology. For example, virtual disks on a SAS controller must be migrated to a SAS controller.

Related links

[RAID Controller Technology: SATA And SAS](#)

SAS Controller

- On SAS controllers, you can migrate a virtual disk by moving the physical disks from one controller to another and then importing the foreign configuration on the receiving controller.
- When moving an enclosure on a SAS controller, you can move the enclosure to any connector number and then import the foreign configuration on the receiving controller.

Related links

[Foreign Configuration Operations](#)

Migrating SAS Virtual Disks To Another System

To migrate virtual disks and volumes from one system to another.

1. Turn off the system that the physical disks are being moved from.
2. If the receiving controller does not have a pre-existing virtual disk configuration, then turn off its server.
3. Place the physical disks into the new enclosure.
4. Start the system connected to the receiving controller.
5. Use the [Foreign Configuration Operations](#) task to import the migrated virtual disks on the receiving controller.
The migration is complete. The virtual disk is now manageable through Storage Management.

Protecting Your Virtual Disk With A Hot Spare

When you create a virtual disk using a RAID controller, it enables the system to continue functioning even during a disk failure. This feature can be enabled by assigning a hot spare to the virtual disk. When a disk fails, the redundant data is rebuilt onto the hot spare without interrupting system operations.

Related links

[Dedicated Hot Spare Protection Policy](#)

Understanding Hot Spares

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

The following sections describe procedures for assigning a hot spare:

- [Assigning And Unassigning Global Hot Spare](#)
- [Assigning And Unassigning Dedicated Hot Spare](#)

Setting Hot Spare Protection Policy

The **Hot Spare Protection Policy** is supported only on Serial Attached SCSI (SAS) controllers.

The **Hot Spare Protection Policy** provides you with a higher protection level for the virtual disks by enabling you to specify the number of dedicated/global hot spares to be assigned to the virtual disks/controller. You can also specify the severity levels for the protection policy. Storage Management sends alerts when the hot spare protection policy is violated.


Storage Management does not provide a default policy; however, you can determine the hot spare protection policy best suited for your environment.

Related links

[Assigning And Unassigning Dedicated Hot Spare](#)
[Assigning And Unassigning Global Hot Spare](#)

Dedicated Hot Spare Protection Policy

Table 35. Hot Spare Properties

Property	Definition
RAID Level	Indicates the RAID level(s) for which you want to configure the hot spare protection policy.
Minimum Number of Disks	Indicates the minimum number of physical disks to be assigned as dedicated hot spares for the selected RAID level(s).
Severity Level	Displays the severity level you want to assign to the generated alert, if this policy is violated.  NOTE: The status of the virtual disk is determined by the severity level you specify for this policy.

Related links

- [Hot Spare Protection Policy](#)
- [Setting Hot Spare Protection Policy](#)
- [Global Hot Spare Protection Policy](#)

Resetting The Hot Spare Protection Policy

Clear **RAID Layout** to reset the dedicated hot spare protection policy.

Global Hot Spare Protection Policy

Table 36. Global Hot Spare Protection Policy Properties

Property	Definition
Enable Global Hot Spare	Enables the Global Hot Spare Protection Policy.
Minimum Number of Disks	Displays the minimum number of physical disks to be assigned as global hot spares for the controller.
Severity Level	Displays the severity level that you must assign to the generated alert, if the Global Hot Spare policy is violated.
Consider Global Hot Spare Policy when determining Virtual Disk status	Storage Management considers the Global Hot Spare policy to determine the status of the virtual disk. The status of the virtual disk is determined by the severity level you specify for this policy.

 **NOTE:** When assigning a global hot spare, consider a physical disk that has higher capacity, which can replace a disk when a disk fails in the controller.

Considerations For Hot Spare Protection Policy

- RAID 0 does not support hot spares or protection policy.
- If the status of the virtual disk is displayed as **Degraded** or **Failed** because of the hot spare protection policy violation, then assign the required number of hot spares (as defined in the protection policies) for the status to be displayed as normal.

Considerations For Enclosure Affinity

- Enclosure affinity settings for dedicated hot spare are applicable only on PERC 6 family of controllers.
- Enclosure affinity settings for a global/dedicated hot spare are not automatically set when you upgrade to version 6.1.

Enclosure affinity settings for a global/dedicated hot spare are not automatically set when you import a foreign virtual disk.

CacheCade Using Solid-State Drives

CacheCade is used to improve random read performance of the solid-state drive (SSD) based virtual disks. A solid-state drive (SSD) is a data storage device that uses solid-state memory to store persistent data. SSDs significantly increase the I/O performance (IOPS) and/or write speed in Mbps from a storage device. With storage controllers, you can create a CacheCade using SSDs. The CacheCade is then used for better performance of the storage I/O operations. Use either Serial Attached SCSI (SAS) SSDs or Serial Advanced Technology Attachment (SATA) SSDs to create a CacheCade.

Create a CacheCade with SSDs in the following scenarios:

- Maximum application performance — Create a CacheCade using SSDs to achieve higher performance without wasted capacity.
- Maximum application performance and higher capacity — Create a CacheCade using SSDs to balance the capacity of the CacheCade with high performance SSDs.
- Higher capacity — If you do not have empty slots for additional HDDs, use SSDs and create a CacheCade. This reduces the number of HDDs required and increases application performance.

The CacheCade feature has the following restrictions:

- Only certified SSDs can be used to create a CacheCade.
- If you create a CacheCade using SSDs, the SSD properties are still retained. At a later time, you can use the SSD to create virtual disks.
- A CacheCade can contain either SAS drives or SATA drives but not both.
- Each SSD in the CacheCade does not have to be of the same size. The CacheCade size is automatically calculated as follows: CacheCade size = capacity of the smallest SSD * the number of SSDs. The unused portion of SSD is wasted and cannot be used as an additional CacheCade or an SSD-based virtual disk.
- The total amount of cache pool with a CacheCade is 512GB. If you create a CacheCade which is greater than 512GB, the storage controller still uses only 512GB.
- In a storage enclosure, the total number of logical devices including virtual disks and CacheCade(s) cannot exceed 64.

 **NOTE: In order to use CacheCade for the virtual disk, the Write and Read policy of the HDD-based virtual disk must be set to Write Back or Force Write Back and read policy must be set to Read Ahead.**

Related links

[Managing The CacheCade](#)

Managing The CacheCade

Ensure that SSDs are available in the storage enclosure before performing any CacheCade operations.

To create, resize, blink, or delete the CacheCade, go to the **CacheCade(s)** page:

The **CacheCade(s)** page is displayed.

1. In the **Server Administrator** window, under the **System** tree, expand **Storage**.
2. Click a storage controller.
3. Click a storage controller component.
4. Under **Available Tasks**, select **Manage CacheCade**.
5. Click **Execute**.

The **CacheCade(s)** page is displayed.




Related links

[CacheCade Using Solid-State Drives](#)

CacheCade Properties

The **CacheCade(s)** page displays a table of properties for each CacheCade.

Table 37. CacheCade Properties

Property	Definition
Status	<p>These icons represent the severity or health of the storage component.</p> <ul style="list-style-type: none"> — Normal/OK — Warning/Non-critical — Critical/Fatal <p>For more information, see Storage Component Severity.</p>
Name	Displays the name of the CacheCade. The link enables you to access the physical disks that constitute the CacheCade.
Tasks	Drop-down list provides the available tasks for the CacheCade.
Size	Provides the size of the CacheCade.
Bus Protocol	Displays the technology that the physical disk is using. Possible values are SAS and SATA.
Disk Cache Policy	Displays whether the disk cache policy of the physical disks that are part of the CacheCade is Enabled or Disabled. See RAID Controller Read, Write, Cache, And Disk Cache Policy .

Creating A CacheCade

To create a CacheCade:

1. In the **CacheCade(s)** page, click **Create CacheCade**.
The **Create CacheCade** page is displayed.
2. In the **CacheCade Group** section, select **SATA SSD** or **SAS SSD** if your system contains both SAS and SATA SSDs.
3. Type a name for the CacheCade. The size is calculated as follows:
CacheCade size = capacity of the smallest SSD * the number of SSDs
4. From the available free disks that are displayed, select the SSDs that you would like to include in the CacheCade. The selected disks are displayed in the **Physical Disks Selected** section.
5. Click **Finish**.



NOTE: The procedure of selecting physical disks while creating a CacheCade is similar to the selection of disks while creating a virtual disk.

Related links

[CacheCade Properties](#)

[Managing The CacheCade](#)

Resizing The CacheCade

To resize the CacheCade:

1. In the **CacheCade(s)** page, navigate to the CacheCade that you want to resize and select **Resize ...** from the **Tasks** drop-down menu.

The **Resize CacheCade(s)** page is displayed.

2. From the available CacheCade(s), add or remove additional CacheCade(s) as required. The disks that you select are displayed in the **Selected Physical Disks** section.



NOTE: While selecting and clearing the SSDs that must participate in the CacheCade, retain at least one of the original SSDs.

3. Click **Finish**.

Renaming The CacheCade

To rename the CacheCade:

1. In the **CacheCade(s)** page, navigate to the CacheCade that you want to rename and select **Rename** from the **Tasks** drop-down menu.
2. In the **Rename** page, type the new name for the CacheCade and click **Finish**.

Blinking And Unblinking The CacheCade

In the **CacheCade(s)** page, navigate to the CacheCade that you want to blink and select **Blink ...** from the available **Tasks** drop-down list box.

The LED of the physical disk(s) participating in the CacheCade turns on.

To unblink the physical disk in the **CacheCade(s)** page, go to the CacheCade disk that you want to unblink and select **Unblink ...** from the available **Tasks** drop-down list box.

Deleting The CacheCade

To delete the CacheCade:

1. In the **CacheCade(s)** page, navigate to the CacheCade that you want to delete and select **Delete ...** from the available **Tasks** drop-down list box.
The following warning is displayed: Warning! CacheCade pool size will be reduced. Do you want to continue with the delete operation?
2. Click **Delete**.

Troubleshooting

This section contains troubleshooting procedures for common situations as well as for specific problems.

Related links

- [Common Troubleshooting Procedures](#)
- [Virtual Disk Troubleshooting](#)
- [Specific Problem Situations And Solutions](#)
- [PCIe SSD 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 Hardware Problems](#)
- [Replacing A Failed Disk](#)
- [Recovering From Removing The Wrong Physical Disk](#)
- [Resolving Microsoft Windows Upgrade Problems](#)

Cables Attached Correctly

Verify that the power-supply cord and adapter cables are connected correctly. If the system is having trouble with read and write operations to a particular virtual disk or non-RAID physical disk (for example, the system hangs), then make sure that the cables connected to the corresponding enclosure or backplane are secure. If the problem persists, you may need to replace a cable.

On SAS controllers, verify that the cable configuration is valid. Refer to the SAS hardware documentation for valid cable configurations. If the cable configuration is invalid, you may receive alerts 2182 or 2356.

For information on alert messages, see the *Server Administrator Messages Reference Guide*.

Related links

- [Read And Write Operations Experience Problems](#)

System Requirements

Make sure that the system meets all system requirements. In particular, verify that the correct levels of firmware and drivers are installed on the system.

Related links

- [Drivers And Firmware](#)

Drivers And Firmware

Storage Management is tested with the supported controller firmware and drivers. In order to function properly, the controller must have the minimum required version of the firmware and drivers installed. The most current versions can be obtained from the support site.

 **NOTE: You can verify which firmware and drivers are installed by selecting the Storage object in the tree view and clicking Information/Configuration. You can also check the Alert Log for alerts relating to unsupported firmware and driver versions.**

It is also recommended to obtain and apply the latest server system BIOS on a periodic basis to benefit from the most recent improvements. For more information on server system BIOS, see the system documentation.

Isolate Hardware Problems

If you receive a **timeout** alert related to a hardware device or if you otherwise suspect that a device attached to the system is experiencing a failure, then to confirm the problem:

- Verify that the cables are correctly connected.
- If the cables are correctly connected 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 devices may be defective. For more information on hardware, refer to the hardware device documentation.

Related links

[Read And Write Operations Experience Problems](#)

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 Physical Disk That Is Part Of A Non-Redundant Virtual Disk](#)

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 will not result in data loss. You must 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 physical disk and the virtual disk is left without a hot spare. In this case, you must replace the failed disk and assign the replacement disk as a hot spare.

 **NOTE: If the redundant virtual disk does not have a hot spare assigned to it, then replace the failed disk using the procedure described in [Replacing a Physical Disk Receiving SMART Alerts](#).**

Related links

[Rebuilding Data](#)

Replacing The Disk

To replace the failed disk that is part of a redundant virtual disk:

1. Remove the failed disk.
2. Insert a new disk. Make sure that the new disk has equal or greater memory space than the one you are replacing. On some controllers, you may not be able to use the additional space by inserting a disk that has a higher memory space.

A rebuild is automatically initiated because the virtual disk is redundant.

Related links

[Virtual Disk Considerations For Controllers](#)

Assigning A Hot Spare

If a hot spare was already assigned to the virtual disk, then the data from the failed disk may already be rebuilt onto the hot spare. In this case, you need to assign a new hot spare.

Related links

[Assigning And Unassigning Dedicated Hot Spare](#)

[Assigning And Unassigning Global Hot Spare](#)

Replacing A Failed Physical Disk That Is Part Of A Non-Redundant Virtual Disk

If the failed physical disk is part of a nonredundant virtual disk (such as RAID 0), then the failure of a single physical disk causes the entire virtual disk to fail. To proceed, you need to verify when a backup was taken, and if there is any new data that has been written to the virtual disk since that time.

If you have backed up recently and there is no new data on the disks that would be missed, you can restore from backup.

Replacing The Disk

1. Delete the virtual disk which is in a failed state.
2. Remove the failed physical disk.
3. Insert a new physical disk.
4. Create a new virtual disk.
5. Restore from backup.

Using The Physical Disk Online Command On Select Controllers

Does my controller support this feature? See [Supported Features](#).

If you do not have a suitable backup available, and if the failed disk is part of a virtual disk on a controller that supports the **Online** physical disk task, then you can attempt to retrieve data by selecting **Online** from the drop-down task menu of the failed disk.

The online command attempts to force the failed disk back into an online state. If you are able to force the disk into an online state, you may be able to recover individual files. The amount of 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.

However, you may not 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.

After retrieving any viable data from the disk, replace the failed disk as described previously in [Replacing A Failed Disk That Is Part Of A Redundant Virtual Disk](#) or [Replacing A Failed Physical Disk That Is Part Of A Non-Redundant Virtual Disk](#).

Related links

[Setting The Physical Disk Online Or Offline](#)

Recovering From Removing The Wrong Physical Disk

If the physical disk 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 no longer has a hot spare since data has been rebuilt onto the disk previously assigned as a hot spare. In this case, assign a new hot spare.

If the physical disk that you removed is part of a redundant virtual disk that does not have a hot spare, then replace the physical disk and perform a rebuild task.

For information on rebuilding physical disks and assigning hot spares, see [Understanding Hot Spares](#).

You can avoid removing the wrong physical disk by blinking the LED display on the physical disk that you intend to remove. For information on blinking the LED display, see [Blinking And Unblinking A Physical Disk](#).

Resolving Microsoft Windows Upgrade Problems

If you upgrade the Microsoft Windows operating system on a server, you may find that Storage Management no longer functions after the upgrade. The installation process installs files and updates registry entries on the server that are specific to the operating system. Changing the operating system can disable Storage Management.

To avoid this problem, you should uninstall Storage Management before upgrading.

After you have uninstalled Storage Management and completed the upgrade, reinstall Storage Management using the Storage Management install media.

Virtual Disk Troubleshooting

The following sections describe troubleshooting procedures for virtual disks:

- [Replacing A Failed Disk That Is Part Of A Redundant Virtual Disk](#)
- [A Rebuild Does Not Work](#)
- [A Rebuild Completes With Errors](#)
- [Cannot Create A Virtual Disk](#)
- [Virtual Disk Errors On Systems Running Linux](#)

- [Problems Associated With Using The Same Physical Disks For Both Redundant And Non-Redundant Virtual Disks](#)

A Rebuild Does Not Work

A rebuild does not work in the following situations:

- The virtual disk is nonredundant — 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, to rebuild it:
 - Remove the failed physical disk and replace it. A rebuild automatically starts on the new disk.
 - Assign a hot spare to the virtual disk and then perform a rebuild.
- You are attempting to rebuild onto a hot spare that is too small — Different controllers have different size requirements for hot spares.
- The hot spare has been unassigned from the virtual disk — This occurs on some controllers if the hot spare is assigned to more than one virtual disk and is being used to rebuild a failed physical disk for another virtual disk.
- The virtual disk includes failed or corrupt physical disks — This situation may generate alert 2083. For information on alert messages, see the *Server Administrator Messages Reference Guide*.
- The rebuild rate setting is too low — If the rebuild rate setting is quite low and the system is processing a number of operations, then the rebuild may take an unusual amount of time to complete.
- The rebuild is canceled — Another user can cancel a rebuild that you have initiated.

Related links

[Setting The Rebuild Rate](#)

A Rebuild Completes With Errors

A rebuild completes with errors 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 indicates successful completion while also generating alert 2163. For information on alert messages, see the *Server Administrator Messages Reference Guide*.

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 drive.
 - 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 a tape drive.
3. Restore the virtual disk from the tape drive onto healthy physical disks.

Cannot Create A Virtual Disk

You may be attempting a RAID configuration that is not supported by the controller. Check the following:

- How many virtual disks already exist on the controller? Each controller supports a maximum number of virtual disks.
- Is there adequate available space on the disk? The physical disks that you have selected for creating the virtual disk must have an adequate amount of free space available.
- The controller may be performing other tasks, such as rebuilding a physical disk, that must run to completion before the controller can create the new virtual disk.

Related links

[Number Of Virtual Disks Per Controller](#)

A Virtual Disk Of Minimum Size Is Not Visible To Windows Disk Management

If you create a virtual disk using the minimum allowable size in Storage Management, the virtual disk may not be visible to Windows Disk Management even after initialization. This occurs because Windows Disk Management is only able to recognize extremely small virtual disks if they are dynamic. It is advisable to create virtual disks of larger size when using Storage Management.

Virtual Disk Errors On Systems Running Linux

On some versions of the Linux operating system, the virtual disk size is limited to 1TB. If you create a virtual disk that exceeds the 1TB limitation, your system may experience the following behavior:

- I/O errors to the virtual disk or logical drive.
- Inaccessible virtual disk or logical drive.
- Virtual disk or logical drive size is smaller than expected.

If you have created a virtual disk that exceeds the 1TB limitation, you must:

1. Back up your data.
2. Delete the virtual disk.
3. Create one or more virtual disks that are smaller than 1TB.
4. Restore your data from backup.

Irrespective of whether your Linux operating system limits the virtual disk size to 1TB, the virtual disk size depends on the version of the operating system and any updates or modifications that you have implemented. For more information on operating system, see your operating system documentation.

Problems Associated With Using The Same Physical Disks For Both Redundant And Non-Redundant Virtual Disks

When creating virtual disks, you should avoid using the same physical disks for both redundant and nonredundant virtual disks. This applies to all controllers. Using the same physical disks for both redundant and nonredundant virtual disks can result in unexpected behavior including data loss.

 **NOTE: SAS controllers do not allow you to create redundant and nonredundant virtual disks on the same set of physical disks.**

Specific Problem Situations And Solutions

This section contains additional troubleshooting information. Topics include:

- [Physical Disk Is Offline Or Displays An Error Status](#)
- [Receive A Bad Block Alert With Replacement, Sense, Or Medium Error](#)
- [Read And Write Operations Experience Problems](#)
- [A Task Menu Option Is Not Displayed](#)
- [A Corrupt Disk Or Drive Message Suggests Running Autocheck During A Reboot](#)
- [Erroneous Status And Error Messages After A Windows Hibernation](#)
- [Storage Management May Delay Before Updating Temperature Probe Status](#)
- [Storage Management May Delay Displaying Storage Devices After Reboot](#)
- [You Are Unable To Log Into A Remote System](#)
- [Cannot Connect To Remote System Running Microsoft Windows Server 2003](#)
- [Reconfiguring A Virtual Disk Displays Error In Mozilla Browser](#)
- [Physical Disks Are Displayed Under The Connector Object Instead Of The Enclosure Object](#)

Physical Disk Is Offline Or Displays An Error Status

A physical disk may display an error status if it has been damaged, is offline, or was a member of a virtual disk that has been deleted or initialized. The following actions may resolve the error condition:

- If a user has changed the status to **Offline**, then return the disk to **Online** status by performing the **Online** disk task.
- Investigate whether there are any cables, enclosure, or controller problems preventing the disk from communicating with the controller. If you find a problem and resolve it and disk does not return to **Online** or **Ready** status, reboot the system.
- If the disk is damaged, replace it.

Related links

[Replacing A Failed Disk](#)

Receive A Bad Block Alert With Replacement, Sense, Or Medium Error

The following alerts or events are generated when a portion of a physical disk is damaged:

- 2146
- 2147
- 2148
- 2149
- 2150

This damage is discovered when the controller performs an operation that requires scanning the disk. Examples of operations that may result in these alerts are as follows:

- Consistency check
- Rebuild
- Virtual disk format
- I/O

If you receive alerts 2146 through 2150 as the result of performing a rebuild or while the virtual disk is in a degraded state, then data cannot be recovered from the damaged disk without restoring from the backup. If you receive alerts 2146 through 2150 under circumstances other than a rebuild, then data recovery may be possible. The following describes each of these situations.

Alerts 2146 Through 2150 Received During A Rebuild Or While A Virtual Disk Is Degraded

Perform the following steps:

1. Replace the damaged physical disk.
2. Create a new virtual disk and allow the virtual disk to completely resynchronize.
While the resynchronization is in progress, the status of the virtual disk is **Resynching**.
3. Restore data to the virtual disk from the backup.

Alerts 2146 Through 2150 Received While Performing I/O, Consistency Check, Format, Or Other Operation

If you receive alerts 2146 through 2150 while performing an operation other than a rebuild, replace the damaged disk immediately to avoid data loss.

Perform the following steps:

1. Back up the degraded virtual disk to a fresh (unused) tape drive.
2. Replace the damaged disk.
3. Perform a rebuild.

Read And Write Operations Experience Problems

If the system hangs, times out, or experiences other problems with read and write operations, then there may be a problem with the controller cables or a device.

Related links

[Cables Attached Correctly](#)

[Isolate Hardware Problems](#)

A Task Menu Option Is Not Displayed

You may notice that the task menus do not always display the same task options. Some tasks are only valid for certain types of objects or at certain times. For example, a [Check Consistency](#) task can only be performed on a redundant virtual disk. Similarly, if a disk is already offline, the Offline task option is not displayed.

There may be other reasons why a task cannot be run at a certain time. For example, there may already be a task running on the object that must complete before additional tasks can be run.

A Corrupt Disk Or Drive Message Suggests Running Autocheck During A Reboot

You can ignore the message and perform autocheck. The reboot completes after autocheck is completed. Depending on the size of your system, it may take about 10 minutes.

Erroneous Status And Error Messages After A Windows Hibernation

Activating the Windows hibernation feature may cause Storage Management to display erroneous status information and error messages. This problem resolves when the Windows operating system recovers from hibernation.

Storage Management May Delay Before Updating Temperature Probe Status

In order to display the enclosure temperature and temperature probe status, Storage Management polls the enclosure firmware at regular intervals to obtain temperature and status information. On some enclosures, there is a short delay before the enclosure firmware reports the current temperature and temperature probe status. This may cause a delay in displaying the correct temperature and temperature probe status.

Related links

[Changing The Warning Threshold On The Temperature Probe](#)

Storage Management May Delay Displaying Storage Devices After Reboot

Storage Management requires time after a reboot to find and inventory all attached storage devices. You may experience a delay in storage controllers being displayed until this operation has completed.

You Are Unable To Log Into A Remote System

Ensure that you have connected to the system with administrator privileges and have provided the correct login details. The remote system may be turned off or there may be network problems.

Cannot Connect To Remote System Running Microsoft Windows Server 2003

When connecting to a remote system running Windows Server 2003, log in to the remote system using an account that has administrator privileges. By default, Windows Server 2003 does not allow anonymous (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.

Reconfiguring A Virtual Disk Displays Error In Mozilla Browser

When reconfiguring a virtual disk using the Mozilla browser, the following error message may display:

```
Although this page is encrypted, the information you
have entered is to be sent over an unencrypted
connection and could easily be read by a third party.
```

You can disable this error message by changing the settings in the browser. To disable this error message:

1. Select **Edit** → **Preferences**.
2. Click **Privacy and Security**.
3. Click **SSL**.
4. Clear the **Sending form data from an unencrypted page to an unencrypted** page option.

Physical Disks Are Displayed Under The Connector Object Instead Of The Enclosure Object

Storage Management surveys the status of physical disks at frequent intervals. When the physical disk is located in an enclosure, the data reported by the SCSI Enclosure Processor (SEP) is used to determine the status of the physical disk.

When the SEP is not functioning, Storage Management can survey the status of the physical disk, but cannot identify the physical disk as being located in the enclosure. In this case, the physical disk is displayed under the Connector object and not under the Enclosure object.

You can resolve this problem by restarting the Server Administrator service or by rebooting the system. For more information on restarting the Server Administrator service, see the *Server Administrator User's Guide*.

PCIe SSD Troubleshooting

The following sections describe troubleshooting procedures for PCIe SSD. Topics include:

- [Peripheral Component Interconnect Express \(PCIe\) Solid-State Drive \(SSD\) Is Not Visible In The Operating System](#)
- [PCIe SSD Is Not Visible In Disk Management In The Operating System](#)

Peripheral Component Interconnect Express (PCIe) Solid-State Drive (SSD) Is Not Visible In The Operating System

Probable Cause: Hardware is not installed correctly

Solution: Check the following components to ensure that they are connected:

- Devices: The PCIe SSDs are installed in a PCIeSSD backplane.
- Backplane: The cables for the PCIeSSD backplane are connected.
- Cables: PCIe cables are unique for the configuration. Ensure that the backplane cable connectors mate with the backplane and the extender card cable connectors mate with the extender card.
- Extender card: The PCIe extender card is plugged into the correct supported slot.

PCIe SSD Is Not Visible In Disk Management In The Operating System

Probable Cause: Device driver is not installed.

Solution:

1. Download the latest PCIe SSD driver from the support site.
2. Open **Device Manager** and double-click **Other Devices** where the PCIe device displays a Yellow mark.
3. Right-click and install the driver on the instance.

For more information on possible error conditions with your PCIe SSD, see the system-specific Owner's Manual.

Frequently Asked Questions

This section provides frequently asked questions that address situations commonly experienced in a storage environment.

- [Why is a Rebuild not Working](#)
- [How To Avoid Removing The Wrong Disk](#)
- [How Can I Safely Remove or Replace a Physical Disk](#)
- [How do I Recover from Removing the Wrong Physical Disk](#)
- [How do I Identify the Firmware Version that is installed](#)
- [Which Controllers do I Have](#)
- [How do I Turn off an Alarm](#)
- [Which RAID Level is Best For me](#)

Why Is A Rebuild Not Working?

For more information, see [A Rebuild Does Not Work](#).

How To Avoid Removing The Wrong Disk

You can avoid removing the wrong disk by blinking the LED display on the disk that you want to remove. For information on blinking the LED display:

- See [Blinking And Unblinking A Physical Disk](#) to blink the LED display on a physical disk.
- See [Blinking And Unblinking A Virtual Disk](#) to blink the LED display on all physical disks included in a particular virtual disk.

If you have already removed the wrong disk, see [Recovering From Removing The Wrong Physical Disk](#).

How Can I Safely Remove Or Replace A Physical Disk

The following sections provide information related to safely removing a healthy physical disk:

- For disks that reside in an enclosure, use the [Preparing To Remove A PCIe SSD](#) task to locate the disk within the enclosure and deactivate it prior to removal.
- For physical disks included in a virtual disk, use the **Offline** task to deactivate the disk prior to removal. If you need help locating the disk within the enclosure, you can blink LED on the disk.

Related links

[Moving Physical And Virtual Disks From One System To Another](#)

[Recovering From Removing The Wrong Physical Disk](#)

[Protecting Your Virtual Disk With A Hot Spare](#)

[Setting The Physical Disk Online Or Offline](#)

[Blinking And Unblinking A Physical Disk](#)

[Replacing A Failed Disk](#)

[Replacing A Physical Disk Receiving SMART Alerts](#)

How Do I Recover From Removing The Wrong Physical Disk

For more information, see [Recovering From Removing The Wrong Physical Disk](#).

How Do I Identify The Firmware Version That Is Installed

The properties information of the controller displays the firmware version that is installed on the controller. From the **Storage** tree view object, you can display the firmware version for all controllers attached to the system. You can also view this information in the **Information/Configuration** page of the controller.

To display the firmware version of all controllers:

1. Select the **Storage** object in the tree view.
2. Click **Information/Configuration** subtab. The **Firmware version** column on the **Information/Configuration** subtab displays the firmware version for all controllers attached to the system.

Which Controllers Do I Have?

Each controller attached to the system is displayed under the **Storage** object in the tree view.

In addition, the **Storage Health** and **Information/Configuration** pages display information for each controller.

To identify the controllers attached to the system:

1. Select the Storage tree view object. The **Health** page displays the name and status for each controller attached to the system.
2. Click **Information/Configuration** subtab to display more information for each controller.
3. To display more information about a particular controller including its tasks, click the controller name in the **Name** column on the **Information/Configuration** subtab. This action displays the controller's **Information/Configuration** subtab. The **Information/Configuration** subtab displays information about the controller's components and enables you to execute the controller tasks.

Related links

[Controllers](#)

[Controllers Properties And Tasks](#)

[Controller Health](#)

How Do I Turn Off An Alarm

Some storage components have alarms to indicate error conditions. For more information on turning off the alarm, see:

- [Turning Off The Controller Alarm](#)
- [Disabling The Controller Alarm](#)
- [Disabling The Enclosure Alarm](#)

Which RAID Level Is Best For Me?

For more information, see [Choosing RAID Levels And Concatenation](#) and [Comparing RAID Level And Concatenation Performance](#).

Supported Features

Different controllers support different features. The tasks displayed by the Storage Management menus and other features vary depending on whether the controller supports the feature. This chapter lists the features supported by the controllers. For more information on Controllers, see your hardware documentation.

 **NOTE: The order of the controllers displayed on Storage Management may differ with the order of the controllers displayed in the Human Interface (HII) and PERC Option ROM. The order of the controllers does not cause any limitation.**

Related links

- [Supported Features On PERC Hardware Controllers](#)
- [Supported Features On The PERC S140 Controller](#)
- [Supported Non-RAID Controllers](#)
- [Enclosure And Backplane Features](#)

Supported Features On PERC Hardware Controllers

This section identifies the controller-supported features and whether an enclosure can be attached to the following hardware controllers:

- PERC H330 Adapter, PERC H330 Mini Monolithic
- PERC H730 Adapter, PERC H730 Mini Monolithic
- PERC H730P Adapter, PERC H730P Mini Monolithic
- PERC H740P Adapter, PERC H740P Mini Monolithic
- PERC H840 Adapter

 **NOTE: The order of the controllers displayed on Storage Management may differ with the order of the controllers displayed in the Human Interface (HII) and PERC Option ROM. The order of the controllers does not cause any limitation.**

The features include the following:

- [Controller Tasks](#)
- [Battery Tasks](#)
- [Connector Tasks](#)
- [Physical Disk Tasks](#)
- [Virtual Disk Tasks](#)
- [Virtual Disk Specifications](#)
- [Supported RAID Levels](#)
- [Read, Write, Cache And Disk Cache Policy](#)
- [Enclosure Support](#)

For enclosure-supported tasks, see [Enclosure And Backplane Features](#).

Controller Tasks Supported On PERC Hardware Controllers

Table 38. Controller Tasks Supported On PERC Hardware Controllers

Controller Task Name	PERC H730P Adapter/Mini Monolithic	PERC H740P Adapter/Mini Monolithic	PERC H730 Adapter/Mini Monolithic	PERC H330 Adapter/Mini Monolithic	PERC H840 Adapter
Enable Alarm	No	No	No	No	No
Disable Alarm	No	No	No	No	No
Quiet Alarm	No	No	No	No	No
Test Alarm	No	No	No	No	No
Load Balance	No	No	No	No	Yes
Reset Configuration	Yes	Yes	Yes	Yes	Yes
Set Rebuild Rate	Yes	Yes	Yes	Yes	Yes
Set Background Initialization Rate	Yes	Yes	Yes	Yes	Yes
Set Check Consistency Rate	Yes	Yes	Yes	Yes	Yes
Set Reconstruct Rate	Yes	Yes	Yes	Yes	Yes
Rescan Controller	No	No	No	No	No
Create Virtual Disk	Yes	Yes	Yes	Yes	Yes
Export Log File	Yes	Yes	Yes	Yes	Yes
Clear Foreign Configuration	Yes	Yes	Yes	Yes	Yes
Import Foreign Configuration	Yes	Yes	Yes	Yes	Yes
Import/Recover Foreign Configuration	Yes	Yes	Yes	Yes	Yes
Set Patrol Read Mode	Yes	Yes	Yes	Yes	Yes
Start Patrol Read	Yes	Yes	Yes	Yes	Yes
Convert RAID To Non-RAID	Yes	No	Yes	Yes	No
Convert Non- RAID To RAID	Yes	No	Yes	Yes	No
Replace Member	Yes	Yes	Yes	Yes	Yes
Import Preview of Foreign Configuration	Yes	Yes	Yes	Yes	Yes
Hot-plug of Enclosures	No	No	No	No	Yes
Change Controller Properties	Yes	Yes	Yes	Yes	Yes
Intelligent Mirroring	Yes	Yes	Yes	Yes	Yes
Redundant Path Configuration	No	No	No	No	Yes
Disk Cache Policy	Yes	Yes	Yes	Yes	Yes
Managing Preserved Cache	Yes	Yes	Yes	No	Yes
Manage Encryption Key	Yes	Yes	Yes	No	Yes
Manage CacheCade	No	No	No	No	No

Controller Task Name	PERC H730P Adapter/Mini Monolithic	PERC H740P Adapter/Mini Monolithic	PERC H730 Adapter/Mini Monolithic	PERC H330 Adapter/Mini Monolithic	PERC H840 Adapter
Persistent Hot Spare	Yes	Yes	Yes	Yes	Yes
Manage Physical Disk Power for Unconfigured and Hot Spare Drives	Yes	Yes	Yes	Yes	Yes
Manage Physical disk power for configured drives	Yes	Yes	Yes	No	Yes
Auto Configure RAID0	Yes	Yes	Yes	No	Yes
Non RAID HDD Disk Cache Policy	Yes	Yes	Yes	Yes	Yes
Controller Reports					
Patrol Read Report	Yes	Yes	Yes	Yes	Yes
Check Consistency Report	Yes	Yes	Yes	Yes	Yes
Slot Occupancy Report	Yes*	Yes*	Yes*	Yes	Yes
Physical Disk Firmware Version Report	Yes	Yes	Yes	Yes	Yes

Slot Occupancy Report — PERC H730 and H730P support flexible backplane zoning functionality. The flexible backplane zoning functionality is available only on backplanes that support 24 slots—PowerEdge R630 and R730xd. For more information on flexible backplane zoning, see [Backplanes](#).

Battery Tasks Supported On PERC Hardware Controllers

Table 39. Battery Tasks Supported on PERC Hardware Controllers

Battery Task Name	PERC H730P Adapter/Mini Monolithic	PERC H740P Adapter/Mini Monolithic	PERC H730 Adapter/Mini Monolithic	PERC H330 Adapter/Mini Monolithic	PERC H840 Adapter
Recondition Battery	No	No	No	NA	No
Start Learn Cycle	No	No	No	NA	No
Delay Learn Cycle	No	No	No	NA	No

Connector Tasks Supported By PERC Hardware Controllers

Table 40. Controller Tasks Supported by PERC Hardware Controllers

Controller Task Name	PERC H730P Adapter/Mini Monolithic	PERC H740P Adapter/Mini Monolithic/	PERC H730 Adapter/Mini Monolithic	PERC H330 Adapter/Mini Monolithic	PERC H840 Adapter
Connector Rescan	No	No	No	No	No

Physical Disk Tasks Supported By PERC Hardware Controllers

Table 41. Physical Disk Tasks Supported by PERC Hardware Controllers

Physical Disk Task Name	PERC H730P Adapter/Mini Monolithic	PERC H740P Adapter/Mini Monolithic	PERC H730 Adapter/Mini Monolithic	PERC H330 Adapter/Mini Monolithic	PERC H840 Adapter
Blink/Unblink	Yes	Yes	Yes	Yes	Yes
Assign and Unassign Global Hot Spare	Yes	Yes	Yes	Yes	Yes
Prepare to Remove	No	No	No	No	No
Offline	Yes	Yes	Yes	Yes	Yes
Online	Yes	Yes	Yes	Yes	Yes
Initialize	No	No	No	No	No
Rebuild	Yes	Yes	Yes	Yes	Yes
Cancel Rebuild	Yes	Yes	Yes	Yes	Yes
Remove Dead Disk Segments	No	No	No	No	No
Format Disk	No	No	No	No	No
Clear	No	No	No	No	No
Cancel Clear	No	No	No	No	No
Cancel Replace Member	Yes	Yes	Yes	Yes	Yes
Convert to Raid Capable Disk	Yes	Yes	Yes	Yes	Yes
Convert to Non-Raid Disk	Yes	Yes	Yes	Yes	Yes
Cryptographic Erase	Yes	Yes	Yes	No	Yes
Revertible Hot Spare	Yes	Yes	Yes	Yes	Yes

Virtual Disk Tasks Supported By PERC Hardware Controllers

Table 42. Virtual Disk Tasks Supported by PERC Hardware Controllers

Virtual Disk Task Name	PERC H730P Adapter/ Mini Monolithic	PERC H740P Adapter/ Mini Monolithic	PERC H730 Adapter/ Mini Monolithic	PERC H330 Adapter/ Mini Monolithic	PERC H840 Adapter
Assign and Unassign Dedicated Hot Spare	Yes	Yes	Yes	Yes	Yes
Create Virtual Disk	Yes	Yes	Yes	Yes	Yes

Virtual Disk Task Name	PERC H730P Adapter/ Mini Monolithic	PERC H740P Adapter/ Mini Monolithic	PERC H730 Adapter/ Mini Monolithic	PERC H330 Adapter/ Mini Monolithic	PERC H840 Adapter
Create Virtual Disk Advanced Wizard	Yes	Yes	Yes	Yes	Yes
Create Virtual Disk Express Wizard	Yes	Yes	Yes	Yes	Yes
Rename	Yes	Yes	Yes	Yes	Yes
Blink/Unblink	Yes	Yes	Yes	Yes	Yes
Reconfigure	Yes	Yes	Yes	Yes	Yes
Change Policy	Yes	Yes	Yes	Yes	Yes
Split Mirror	No	No	No	No	No
Unmirror	No	No	No	No	No
Delete Virtual Disk	Yes	Yes	Yes	Yes	Yes
Check Consistency	Yes	Yes	Yes	Yes	Yes
Cancel Check Consistency	Yes	Yes	Yes	Yes	Yes
Pause Check Consistency	No	No	No	No	No
Resume Check Consistency	No	No	No	No	No
Cancel Background Initialization (BGI)	Yes	Yes	Yes	Yes	Yes
Format Virtual Disk	No	No	No	No	No
Cancel Format Virtual Disk	No	No	No	No	No
Restore Dead Disk Segments	No	No	No	No	No
Fast Initialize Virtual Disk	Yes	Yes	Yes	Yes	Yes
Slow Initialize Virtual Disk	Yes	Yes	Yes	Yes	Yes
Replace Member	Yes	Yes	Yes	Yes	Yes
Encrypt Virtual Disk	Yes	Yes	Yes	No	Yes
Clear Virtual Disk Bad Blocks	No	No	No	Yes	Yes

Virtual Disk Specifications For PERC Hardware Controllers

Table 43. Virtual Disk Specifications for PERC Hardware Controllers

Virtual Disk Specification	PERC H730P Adapter/Mini Monolithic	PERC H740P Adapter/Mini Monolithic	PERC H730 Adapter/Mini Monolithic	PERC H330 Adapter/ Mini Monolithic	PERC H840 Adapter
Maximum Number of Virtual Disks per Controller	64	64	64	32	64*
Minimum Virtual Disk Size	100MB	100MB	100MB	100MB	100MB

Virtual Disk Specification	PERC H730P Adapter/Mini Monolithic	PERC H740P Adapter/Mini Monolithic	PERC H730 Adapter/Mini Monolithic	PERC H330 Adapter/Mini Monolithic	PERC H840 Adapter
Maximum Virtual Disk Size	None	None	None	None	None
Maximum Number of Spans per Virtual Disk	8	8	8	8	8
Maximum Number of Physical Disks per Span	32	32	32	8	32
Minimum Stripe Size	64 K	64 K	64 K	64 K	64 K
Maximum Stripe Size	1MB	1MB	1MB	64 K	1MB
Maximum Number of Virtual Disks per Disk Group	16	16	16	16	16
Maximum Number of Physical Disks that can be Concatenated	NA	NA	NA	NA	NA
Maximum Number of Physical Disks in a RAID 0	32	32	32	16	32
Maximum Number of Physical Disks in a RAID 1	2	2	2	2	2
Maximum Number of Physical Disks in a RAID 5	32	32	32	16	32
Maximum Number of Physical Disks in a RAID 10	256	256	256	16	256
Maximum Number of Physical Disks in a RAID 50	256	256	256	NA	256
Minimum Number of Physical Disks that can be Concatenated	NA	NA	NA	NA	NA
Minimum Number of Physical Disks in a RAID 0	1	1	1	2	1

Virtual Disk Specification	PERC H730P Adapter/Mini Monolithic	PERC H740P Adapter/Mini Monolithic	PERC H730 Adapter/Mini Monolithic	PERC H330 Adapter/Mini Monolithic	PERC H840 Adapter
Minimum Number of Physical Disks in a RAID 1	2	2	2	2	2
Minimum Number of Physical Disks in a RAID 5	3	3	3	3	3
Minimum Number of Physical Disks in a RAID 10	4	4	4	4	4
Minimum Number of Physical Disks in a RAID 50	6	6	6	6	6
Maximum Number of Physical Disks in a RAID 6	32	32	32	NA	32
Maximum Number of Physical Disks in a RAID 60	256	256	256	NA	256
Minimum Number of Physical Disks in a RAID 6	4	4	4	NA	4
Minimum Number of Physical Disks in a RAID 60	8	8	8	NA	8

RAID Levels Supported By PERC Hardware Controllers

Table 44. RAID Levels Supported by PERC Hardware Controllers

RAID Level	PERC H730P Adapter/ Mini Monolithic	PERC H740P Adapter/ Mini Monolithic	PERC H730 Adapter/ Mini Monolithic	PERC H330 Adapter/ Mini Monolithic	PERC H840 Adapter
Concatenation	No	No	No	No	No
RAID 0	Yes	Yes	Yes	Yes	Yes
RAID 1	Yes	Yes	Yes	Yes	Yes
RAID 5	Yes	Yes	Yes	Yes	Yes
RAID 10	Yes	Yes	Yes	Yes	Yes
RAID 50	Yes	Yes	Yes	Yes	Yes
RAID 6	Yes	Yes	Yes	No	Yes
RAID 60	Yes	Yes	Yes	No	Yes

Read, Write, Cache And Disk Cache Policy Supported By PERC Hardware Controllers

Table 45. Read, Write, and Cache Policy Supported by PERC Hardware Controllers

Read, Write, and Cache Policy	PERC H730P Adapter/ Mini Monolithic	PERC H740P Adapter/ Mini Monolithic	PERC H730 Adapter/ Mini Monolithic	PERC H330 Adapter/ Mini Monolithic	PERC H840 Adapter
Cache settings	Yes	Yes	Yes	No	Yes
Read Policy	Yes	Yes	Yes	No	Yes
Read Ahead (Enabled)	Yes	Yes	Yes	No	Yes
Adaptive Read Ahead	No	No	No	No	No
No Read Ahead (Disabled)	Yes	Yes	Yes	Yes	Yes
Write Policy	Yes	Yes	Yes	No	Yes
Write Back (Enabled)	Yes	Yes	Yes	No	Yes
Write Through (Disabled)	Yes	Yes	Yes	Yes	Yes
Force Write Back (Enabled Always)	Yes	Yes	Yes	No	Yes
Write Cache Enabled Protected	No	No	No	No	No
Cache Policy	No	No	No	No	No
Disk Cache Policy	Yes	Yes	Yes	Yes	Yes
Cache I/O	No	No	No	No	No
Direct I/O	No	No	No	No	No

Enclosure Support On PERC Hardware Controllers

Table 46. Enclosure Support on PERC Hardware Controllers

Enclosure Support	PERC H730P Adapter/ Mini Monolithic	PERC H740P Adapter/ Mini Monolithic	PERC H730 Adapter/ Mini Monolithic	PERC H330 Adapter/ Mini Monolithic	PERC H840 Adapter
Can an enclosure be attached to this controller?	No	No	No	No	Yes

Supported Features On The PERC S140 Controller

This section identifies the controller-supported features and whether or not an enclosure can be attached to the controller.

- [Controller Tasks](#)

- [Physical Disk Tasks](#)
- [Virtual Disk Specifications](#)
- [Supported RAID Levels](#)

Controller Tasks Supported On the PERC S140 Controller

Table 47. Controller Tasks Supported On the PERC S140 Controller

Controller Task Name	PERC S140
Create Virtual Disk	Yes

Physical Disk Tasks Supported By The PERC S140 Controller

Table 48. Physical Disk Tasks Supported by the PERC S140 Controller

Physical Disk Task Name	PERC S140
Blink/Unblink	Yes
Assign and Unassign Global Hot Spare	Yes
Set Disk Cache Policy	Yes
Export log	Yes
Convert to RAID capable disk	Yes

PERC S140 support on Non-volatile memory express (NVMe) devices

PERC S140 support on NVMe devices allows you to create RAID volumes or virtual disks on PCIe SSD devices connected with a PERC S140 controller.

 **NOTE: This feature does not support NVDIMM's.**

The features supported by PERC S140 on NVMe devices are mentioned below:


- Enumerating virtual disks on NVMe
- Deleting virtual disks on NVMe
- Booting virtual disks on NVMe
- Importing foreign configuration on NVMe - Automatically imported by the software RAID controller.
- Assigning Hot Spare devices on NVMe
- Converting Non-RAID NVMe devices to RAID

Virtual Disk Specifications For PERC S140 Controller

Table 49. Virtual Disk Specifications for PERC S140 Controller

Virtual Disk Specification	PERC S140
Maximum Number of Virtual Disks per Controller	8
Minimum Virtual Disk Size	100MB
Maximum Virtual Disk Size	None
Maximum Number of Spans per Virtual Disk	NA
Maximum Number of Physical Disks per Span	NA
Minimum Stripe Size	64 K

Virtual Disk Specification	PERC S140
Maximum Stripe Size	64 K
Maximum Number of Virtual Disks per Physical Disk	8
Maximum Number of Physical Disks that Can Be Concatenated	NA
Maximum Number of Physical Disks in a RAID 0	8
Maximum Physical Disks in a RAID 1	2
Maximum Number of Physical Disks in a RAID 5	8
Maximum Number of Physical Disks in a RAID 10	4
Minimum Number of Physical Disks that Can Be Concatenated	NA
Minimum Number of Physical Disks in a RAID 0	2
Minimum Number of Physical Disks in a RAID 1	2
Minimum Number of Physical Disks in a RAID 5	3
Minimum Number of Physical Disks in a RAID 10	4

 **NOTE:** When creating virtual disks using software RAID controllers, the information related to the member physical disks of the virtual disk is enumerated or displayed on Storage Management after a short delay. This delay in displaying the information does not cause any functional limitation. If you are creating partial virtual disks, it is recommended that you provide Storage Management adequate time between each partial virtual disk creation process.

RAID Levels Supported By The PERC S140 Controller

Table 50. RAID Levels Supported by the PERC S140 Controller

RAID Level	PERC S140
RAID 0	Yes
RAID 1	Yes
RAID 5	Yes
RAID 10	Yes
RAID 50	No
RAID 6	No
RAID 60	No

Supported Features On The Non-RAID Controllers

This section identifies the controller-supported features and whether an enclosure can be attached to the controller.

- [Controller Tasks](#)
- [Battery Tasks](#)
- [Connector Tasks](#)
- [Physical Disk Tasks](#)
- [Virtual Disk Tasks](#)

· [Enclosure Support](#)

For enclosure supported tasks, see [Enclosure And Backplane Features](#).

Controller Tasks Supported On The Non-RAID Controllers

Table 51. Controller Tasks Supported on the Non-RAID Controllers

Controller Task Name	Non-RAID SCSI	Non-RAID SAS
Enable Alarm	No	No
Disable Alarm	No	No
Quiet Alarm	No	No
Test Alarm	No	No
Reset configuration	No	No
Set Rebuild Rate	No	No
Set Background Initialization Rate	No	No
Set Check Consistency Rate	No	No
Set Reconstruct Rate	No	No
Rescan Controller	No	No
Create Virtual Disk	No	No
Export Log File	No	No
Clear Foreign Configuration	No	No
Import Foreign Configuration	No	No
Import/Recover Foreign Configuration	No	No
Set Patrol Read Mode	No	No
Start Patrol Read	No	No
Stop Patrol Read	No	No

Battery Tasks Supported By Non-RAID Controllers

Table 52. Battery Tasks supported by Non-RAID Controllers

Battery Task Name	Non-RAID SCSI	Non-RAID SAS
Recondition Battery	No	No
Start Learn Cycle	No	No
Delay Learn Cycle	No	No

Connector Tasks Supported On The Non-RAID Controllers

Table 53. Connector Tasks Supported on the Non-RAID Controllers

Connector Task Name	Non-RAID SCSI	Non-RAID SAS
Connector Rescan	No	No

Physical Disk Tasks Supported On The Non-RAID Controllers

Table 54. Physical Disk Tasks Supported on the Non-RAID Controllers

Physical Disk Task Name	Non-RAID SCSI	Non-RAID SAS
Blink/Unblink	Yes	Yes
Assign and Unassign Global Hot Spare	No	No
Prepare to Remove	No	No
Offline	No	No
Online	No	No
Initialize	No	No
Rebuild	No	No
Cancel Rebuild	No	No
Remove Dead Disk Segments	No	No
Format Disk	No	No
Clear	No	No
Cancel Clear	No	No

Virtual Disk Tasks Supported By The Non-RAID Controllers

Table 55. Virtual Disk Tasks supported by the Non-RAID Controllers

Virtual Disk Task Name	Non-RAID SCSI	Non-RAID SAS
Assign and Unassign Dedicated Hot Spare	No	No
Create Virtual Disk	No	No
Create Virtual Disk Advanced Wizard	No	No
Create Virtual Disk Express Wizard	No	No
Rename	No	No
Blink/Unblink	No	No
Reconfigure	No	No
Change Policy	No	No
Split Mirror	No	No
Unmirror	No	No
Delete Last Virtual Disk	No	No
Delete (any) Virtual Disk	No	No
Check Consistency	No	No
Cancel Check Consistency	No	No
Pause Check Consistency	No	No

Virtual Disk Task Name	Non-RAID SCSI	Non-RAID SAS
Resume Check Consistency	No	No
Cancel Background Initialization (BGI)	No	No
Format Virtual Disk	No	No
Cancel Format Virtual Disk	No	No
Restore Dead Disk Segments	No	No
Initialize Virtual Disk	No	No
Fast Initialize Virtual Disk	No	No
Slow Initialize Virtual Disk	No	No
Cancel Initialize Virtual Disk	No	No

Enclosure Support On The Non-RAID Controllers

Table 56. Enclosure Support on the Non-RAID Controllers

Enclosure Support	Non-RAID SCSI	Non-RAID SAS
Can an enclosure be attached to this controller?	Yes	No

Enclosure And Backplane Features

This section identifies the features supported by the enclosure or backplane.

- Enclosure and Backplane Tasks

For information on controller-supported features, see:

- Supported Features on PERC 10 controllers, see [Supported Features On The PERC Hardware Controllers](#)

Enclosure And Backplane Tasks

Table 57. Enclosure Tasks for Controllers

Enclosure Tasks	MD1400 Storage	MD1420
Enable Alarm	No	No
Disable Alarm	No	No
Set Temperature Probe Values	No	No
Set Asset Data (includes asset tag and asset name)	Yes	Yes
Blink Enclosure	Yes	Yes
Controller Reports		
Slot Occupancy Report	Yes	Yes

Maximum Supported Configuration For SAS Controllers

Table 58. Maximum Supported Configuration for SAS Controllers

Maximum Supported Configuration	SAS (MD 1400, MD 1420)
External controllers on each Server (a)	2
External connectors on each controller (b)	2
Enclosures per connector (c)	4
Total numbers of enclosures on a server (a x b x c)	16

Determining The Health Status For Storage Components












The chapter describes how the status of lower-level storage components is "rolled up" into the combined status displayed for the controller or other higher-level component. The examples provided by these tables do not cover all scenarios, but they do indicate how status is rolled up when a particular component is in a healthy, degraded, or failed state.

Related links

[Storage Component Severity](#)




















Health Status Rollup: Battery Is Charging Or Dead

Table 59. Health Status Rollup: Battery is Charging or Dead (Enclosures Not Included)

	Storage Subsystem	Controller	Battery	Connector	Physical Disk(s)	Firmware/Driver	Virtual Disk(s)
Component Status							
							
Health Rollup			N/A	N/A	N/A	N/A	N/A

Health Status Rollup: Physical Disks In A Virtual Disk Are Failed Or Removed

Table 60. Health Status Rollup: Physical Disks in a Virtual Disk are Failed or Removed (Enclosures Not Included)

	Storage Subsystem	Controller	Battery	Connector	Physical Disk(s)	Firmware/Driver	Virtual Disks
Component Status							
							
Health Rollup							
							

Health Status Rollup: Physical Disks In A Virtual Disk Are Unsupported, Partially, Or Permanently Degraded

Table 61. Health Status Rollup: Physical Disks in a Virtual Disk are Unsupported, Partially, or Permanently Degraded (Enclosures Not Included)

	Storage Subsystem	Controller	Battery	Connector	Physical Disk(s)	Firmware/Driver	Virtual Disk(s)
Component Status							
Health Rollup							

Health Status Rollup: All Physical Disks In A Virtual Disk Are In Foreign State

Table 62. Health Status Rollup: All Physical Disks in a Virtual Disk are in Foreign State (Enclosures Not Included)

	Storage Subsystem	Controller	Battery	Connector	Physical Disk(s)	Firmware/Driver	Virtual Disk(s)
Component Status							N/A
Health Rollup							N/A

Health Status Rollup: Some Physical Disks In A Virtual Disk Are In Foreign State

Table 63. Health Status Rollup: Some Physical Disks in a Virtual Disk are in Foreign State (Enclosures Not Included)

	Storage Subsystem	Controller	Battery	Connector	Physical Disk(s)	Firmware/Driver	Virtual Disk(s)
Component Status							
Health Rollup							

Health Status Rollup: Virtual Disk Is Degraded; Physical Disks Are Failed Or Rebuilding

Table 64. Health Status Rollup: Virtual Disk is Degraded; Physical Disks are Failed or Rebuilding (Enclosures Not Included)

	Storage Subsystem	Controller	Battery	Connector	Physical Disk(s)	Firmware/Driver	Virtual Disk(s)
Component Status							
Health Rollup							

	Storage Subsystem	Controller	Battery	Connector	Physical Disk(s)	Firmware/Driver	Virtual Disk(s)

Health Status Rollup: Virtual Disk Is Failed

Table 65. Health Status Rollup: Virtual Disk is Failed (Enclosures Not Included)

	Storage Subsystem	Controller	Battery	Connector	Physical Disk(s)	Firmware/Driver	Virtual Disk(s)
Component Status							
Health Rollup							

Health Status Rollup: Unsupported Firmware Version

Table 66. Health Status Rollup: Unsupported Firmware Version (Enclosures Not Included)

	Storage Subsystem	Controller	Battery	Connector	Physical Disk(s)	Firmware/Driver	Virtual Disk(s)
Component Status							
Health Rollup							

Health Status Rollup: Enclosure Power Supply Failed Or Power Connection Removed

Table 67. Health Status Rollup: Enclosure Power Supply Failed or Power Connection Removed

	Storage Subsystem	Controller	Connector	Enclosure	Enclosure Power Supply	Virtual Disks	Physical Disks
Component Status							
Health Rollup					N/A	N/A	N/A

Health Status Rollup: One Enclosure Fan Is Failed

Table 68. Health Status Rollup: One Enclosure Fan is Failed

	Storage Subsystem	Controller	Battery	Connector	Physical Disk(s)	Firmware/Driver	Virtual Disk(s)
Component Status							
Health Rollup					N/A	N/A	N/A

Health Status Rollup: One Enclosure EMM Is Failed

Table 69. Health Status Rollup: One Enclosure EMM is Failed

	Storage Subsystem	Controller	Connector	Enclosure	Enclosure EMM	Virtual Disks	Physical Disks
Component Status							
Health Rollup					N/A		

Health Status Rollup: One Enclosure Temperature Probe Is Failed

Table 70. Health Status Rollup: One Enclosure Temperature Probe is Failed

	Storage Subsystem	Controller	Connector	Enclosure	Enclosure Temperature Probe	Virtual Disks	Physical Disks
Component Status							
Health Rollup					N/A		

Health Status Rollup: Lost Both Power Connections To The Enclosure









Table 71. Health Status Rollup: Lost Both Power Connections to the Enclosure

	Storage Subsystem	Controller	Connector	Enclosure	All Enclosure Components	Virtual Disks	Physical Disks
Component Status						N/A	N/A
Health Rollup					N/A	N/A	N/A

Health Status Rollup: One Or More Physical Disks Are Failed












Table 72. Health Status Rollup: One or More Physical Disks are Failed

	Storage Subsystem	Controller	Connector	Enclosure	Enclosure Physical Disk(s)	Virtual Disks	Physical Disks
Component Status							

	Storage Subsystem	Controller	Connector	Enclosure	Enclosure Physical Disk(s)	Virtual Disks	Physical Disks
Health Rollup					N/A		N/A
							

Health Status Rollup: Physical Disk Is Rebuilding

Table 73. Health Status Rollup: Physical Disk is Rebuilding

	Storage Subsystem	Controller	Connector	Enclosure	Enclosure Components	Virtual Disks	Physical Disks
Component Status							
Health Rollup					N/A		N/A