



USER'S GUIDE

ServeRAID-M Software

Eighth Edition (December 2012)

© Copyright International Business Machines Corporation 2012.

US Government Users Restricted Rights – Use, duplication or disclosure
restricted by GSA ADP Schedule Contract with IBM® Corp.

Preface

This document explains how to use the MegaRAID® Storage Manager™ software and the WebBIOS Configuration Utility to configure, monitor, and maintain the ServeRAID-M® SAS/SATA controllers and the storage-related devices connected to them.

Organization

This document has the following chapters and appendixes:

- [Chapter 1, “Overview,”](#) describes the SAS, Serial ATA (SATA) II, and Solid State Disk (SSD) technologies, new features, and configuration scenarios.
- [Chapter 2, “Introduction to RAID,”](#) describes RAID (Redundant Array of Independent Disks), RAID functions and benefits, RAID components, RAID levels, and configuration strategies. In addition, it defines the RAID availability concept, and offers tips for configuration planning.
- [Chapter 3, “Self-Encrypting Disk Feature,”](#) describes the Self-Encrypting Encryption (SED) features, terminology, and workflow.
- [Chapter 4, “WebBIOS Configuration Utility,”](#) explains how to use the pre-boot WebBIOS Configuration Utility to create and manage storage configurations.
- [Chapter 5, “MegaRAID Storage Manager Overview and Installation,”](#) introduces the main features of MegaRAID Storage Manager software and explains how to install it.
- [Chapter 6, “MegaRAID Storage Manager Screens and Menus,”](#) describes the layout of the MegaRAID Storage Manager window and lists the available menu options.

- [Chapter 7, “Configuration,”](#) describes how to use the MegaRAID Storage Manager software to configure or reconfigure storage devices, how to save configurations, and how to apply saved configurations to a controller.
- [Chapter 8, “Monitoring System Events and Storage Devices,”](#) explains how the MegaRAID Storage Manager software monitors the status of storage configurations and devices and displays information about them.
- [Chapter 9, “Maintaining and Managing Storage Configurations,”](#) describes the MegaRAID Storage Manager maintenance functions for virtual disks and other storage devices.
- [Chapter 10, “CacheCade 2.0 Software,”](#) describes the features supported by the CacheCade 2.0 Software feature..
- [Appendix A, “Events and Messages,”](#) provides descriptions of the MegaRAID Storage Manager events.
- [Appendix B, SNMP Extension Agent Trap Definitions,](#) describes the trap definitions that an SNMP-based management application uses to send information about system events
- [Appendix C, “Glossary,”](#) contains definitions of storage-related terms.
- [Appendix D, Battery Glossary,](#) contains definitions of battery-related terms.
- [Appendix E, “Notices,”](#) contains information about the warranty, patents, license inquiries, and trademarks.

Conventions

Note: Notes contain supplementary information that can have an effect on system performance.

Attention: Attention notices identify actions that might adversely affect equipment operation, system performance, or data integrity.

Contents

Chapter 1 Overview

1.1	SAS Technology	1-1
1.2	Serial-attached SCSI Device Interface	1-3
1.3	Serial ATA Features	1-3
1.4	Solid State Drive Features	1-4
1.4.1	Solid State Drive Guard	1-5
1.5	Integrated MegaRAID Mode and MegaRAID Mode	1-5
1.6	Feature on Demand (FoD) Upgrades	1-6
1.6.1	Feature on Demand: iMR RAID 5 + Self-Encrypting Disk Upgrade	1-7
1.6.2	Feature on Demand: MegaRAID RAID 6/60 Upgrade	1-7
1.6.3	Feature on Demand: FastPath Upgrade	1-8
1.6.4	Feature on Demand: CacheCade 2 Upgrade	1-9
1.7	UEFI 2.0 Support	1-10
1.8	Configuration Scenarios	1-10
1.8.1	Valid Drive Mix Configurations	1-12

Chapter 2 Introduction to RAID

2.1	RAID Description	2-1
2.2	RAID Benefits	2-2
2.3	RAID Functions	2-2
2.4	Components and Features	2-2
2.4.1	Physical Array	2-3
2.4.2	Virtual Drive	2-3
2.4.3	RAID Drive Group	2-3
2.4.4	Fault Tolerance	2-3

2.4.5	Consistency Check	2-5
2.4.6	Copyback	2-5
2.4.7	Background Initialization	2-7
2.4.8	Patrol Read	2-7
2.4.9	Disk Striping	2-7
2.4.10	Disk Mirroring	2-9
2.4.11	Parity	2-9
2.4.12	Disk Spanning	2-10
2.4.13	Hot Spares	2-12
2.4.14	Disk Rebuilds	2-14
2.4.15	Rebuild Rate	2-15
2.4.16	Hot Swap	2-15
2.4.17	Drive States	2-15
2.4.18	Virtual Drive States	2-16
2.4.19	Enclosure Management	2-17
2.5	RAID Levels	2-17
2.5.1	Summary of RAID Levels	2-17
2.5.2	Selecting a RAID Level	2-18
2.5.3	RAID 0	2-19
2.5.4	RAID 1	2-20
2.5.5	RAID 5	2-20
2.5.6	RAID 6	2-22
2.5.7	RAID 10	2-23
2.5.8	RAID 50	2-24
2.5.9	RAID 60	2-26
2.6	RAID Configuration Strategies	2-27
2.6.1	Maximizing Fault Tolerance	2-28
2.6.2	Maximizing Performance	2-29
2.6.3	Maximizing Storage Capacity	2-31
2.7	RAID Availability	2-32
2.7.1	RAID Availability Concept	2-32
2.8	Configuration Planning	2-33
2.8.1	Number of Drives	2-33
2.8.2	Drive Group Purpose	2-33

Chapter 3

Self-Encrypting Disk Feature

3.1	Overview	3-1
3.2	Purpose	3-2
3.3	Terminology	3-2
3.4	Workflow	3-3
3.4.1	Enable Security	3-3
3.4.2	Change Security	3-4
3.4.3	Create Secure Virtual Drives	3-5
3.4.4	Import a Foreign Configuration	3-6
3.5	Instant Secure Erase	3-7

Chapter 4

WebBIOS Configuration Utility

4.1	Overview	4-1
4.2	Starting the WebBIOS CU	4-2
4.3	WebBIOS Configuration Utility Main Screen Options	4-3
4.4	Creating a Storage Configuration	4-6
4.4.1	Selecting the Configuration with the Configuration Wizard	4-6
4.4.2	Using Automatic Configuration	4-9
4.4.3	Using Manual Configuration	4-10
4.5	Selecting Self-Encrypting Disk Security Options	4-55
4.5.1	Enabling the Security Key Identifier, Security Key, and Passphrase	4-55
4.5.2	Changing the Security Key Identifier, Security Key, and Pass Phrase	4-61
4.5.3	Disabling the Drive Security Settings	4-68
4.5.4	Importing Foreign Configurations	4-70
4.6	Viewing and Changing Device Properties	4-70
4.6.1	Viewing and Changing Controller Properties	4-70
4.6.2	Viewing and Changing Virtual Drive Properties	4-75
4.6.3	Viewing Drive Properties	4-77
4.6.4	Shield State	4-79
4.6.5	Viewing and Changing Battery Backup Unit Information	4-83
4.6.6	Managing Link Speed	4-87

4.6.7	Viewing Enclosure Properties	4-90
4.6.8	SSD Disk Cache Policy	4-92
4.7	Viewing and Expanding a Virtual Drive	4-93
4.8	Recovering/Clearing Punctured Block Entries	4-96
4.9	Suspending and Resuming Virtual Drive Operations	4-96
4.10	Non-SED Secure Erase	4-98
4.10.1	Erasing a Non-SED Physical Drive	4-98
4.10.2	Virtual Drive Erase	4-101
4.11	Viewing System Event Information	4-103
4.12	Managing Configurations	4-105
4.12.1	Running a Consistency Check	4-105
4.12.2	Deleting a Virtual Drive	4-105
4.12.3	Importing or Clearing a Foreign Configuration	4-106
4.12.4	Migrating the RAID Level of a Virtual Drive	4-110
4.12.5	New Drives Attached to a MegaRAID Controller	4-113

Chapter 5

MegaRAID Storage Manager Overview and Installation

5.1	Overview	5-1
5.1.1	Creating Storage Configurations	5-1
5.1.2	Monitoring Storage Devices	5-2
5.1.3	Maintaining Storage Configurations	5-2
5.2	Hardware and Software Requirements	5-2
5.3	Prerequisites to Running MegaRAID Storage Manager Remote Administration	5-3
5.4	Installing MegaRAID Storage Manager	5-4
5.4.1	Prerequisite for MegaRAID Storage Manager Installation	5-4
5.4.2	Installing MegaRAID Storage Manager Software on Microsoft Windows	5-4
5.4.3	Source Port and Destination Port for MSM Components	5-12
5.4.4	Prerequisites for Installing MegaRAID Storage Manager on the RHEL6.X x64 Operating System	5-13
5.4.5	Installing MegaRAID Storage Manager for Linux Operating System	5-14
5.4.6	Linux Error Messages	5-16
5.4.7	Kernel Upgrade	5-16

5.4.8	Uninstalling MegaRAID Storage Manager Software on Linux	5-17
5.5	MegaRAID Storage Manager Support and Installation on the VMware Operating System	5-18
5.5.1	Prerequisites for Installing MegaRAID Storage Manager for VMware	5-18
5.5.2	Uninstalling MegaRAID Storage Manager on the VMware Operating System	5-18
5.5.3	Limitations	5-19
5.6	Installing and Configuring an SNMP Agent	5-20
5.6.1	Prerequisite for LSI SNMP Agent RPM Installation	5-21
5.6.2	Prerequisite for Installing SNMP Agent on Linux Server	5-21
5.6.3	Installing and Configuring an SNMP Agent on Linux	5-22
5.6.4	Installing an SNMP Agent on Windows	5-24

Chapter 6

MegaRAID Storage Manager Screens and Menus

6.1	Starting MegaRAID Storage Manager Software	6-1
6.2	Discovery and Login	6-2
6.3	LDAP Support	6-7
6.4	Configuring LDAP Support Settings	6-9
6.5	MegaRAID Storage Manager Main Menu Screen	6-11
6.5.1	Dashboard/Physical View/Logical View	6-11
6.5.2	Dashboard	6-12
6.5.3	Physical View	6-13
6.5.4	Logical View	6-14
6.5.5	Shield State	6-17
6.5.6	Displaying the Virtual Drive Properties	6-21
6.5.7	SSD Disk Cache Policy	6-24
6.5.8	Non-SED Secure Erase Support	6-28
6.5.9	Rebuild Write Cache	6-34
6.5.10	Background Suspend or Resume Support	6-34
6.5.11	Enclosure Properties	6-36
6.5.12	Monitoring Battery Backup Units	6-36
6.5.13	Properties/Graphical View Tabs	6-40
6.5.14	Event Log Panel	6-41

Chapter 7

Configuration

7.1	Creating a New Storage Configuration	7-2
7.1.1	Selecting Virtual Drive Settings	7-2
7.1.2	Creating a Virtual Drive Using Simple Configuration	7-4
7.1.3	Creating a Virtual Drive Using Advanced Configuration	7-10
7.2	Selecting Self-Encrypting Disk Security Options	7-19
7.2.1	Enabling Drive Security	7-19
7.2.2	Changing the Security Key Identifier, Security Key, and Pass Phrase	7-25
7.2.3	Disabling Drive Security	7-31
7.2.4	Importing or Clearing a Foreign Configuration	7-33
7.3	Adding Hot Spare Drives	7-36
7.4	Changing Adjustable Task Rates	7-37
7.5	Recovering and Clearing Punctured Block Entries	7-40
7.6	Changing Virtual Drive Properties	7-40
7.7	Changing a Virtual Drive Configuration	7-43
7.7.1	Accessing the Modify Drive Group Wizard	7-43
7.7.2	Adding a Drive or Drives to a Configuration	7-45
7.7.3	Removing a Drive from a Configuration	7-48
7.7.4	Replacing a Drive	7-49
7.7.5	Migrating the RAID Level of a Virtual Drive	7-50
7.8	Deleting a Virtual Drive	7-53

Chapter 8

Monitoring System Events and Storage Devices

8.1	Monitoring System Events	8-1
8.1.1	System Log	8-2
8.1.2	Pop-up Notification	8-3
8.1.3	E-mail Notification	8-3
8.2	Configuring Alert Notifications	8-4
8.2.1	Setting Alert Delivery Methods	8-6
8.2.2	Changing Alert Delivery Methods for Individual Events	8-7
8.2.3	Changing the Severity Level for Individual Events	8-9

8.2.4	Reverting to a Default Individual Event Configuration	8-10
8.2.5	Entering or Editing the Sender Email Address and SMTP Server	8-10
8.2.6	Authenticating an SMTP Server	8-11
8.2.7	Saving Backup Configurations	8-12
8.2.8	Loading Backup Configurations	8-13
8.2.9	Adding Email Addresses of Recipients of Alert Notifications	8-13
8.2.10	Testing Email Addresses of Recipients of Alert Notifications	8-14
8.2.11	Removing Email Addresses of Recipients of Alert Notifications	8-15
8.3	Monitoring Server Events	8-15
8.4	Monitoring Controllers	8-15
8.5	Monitoring Drives	8-17
8.6	Running a Patrol Read	8-18
8.6.1	Patrol Read Task Rates	8-21
8.7	Monitoring Virtual Drives	8-21
8.8	Monitoring Enclosures	8-23
8.9	Monitoring Battery Backup Units	8-24
8.9.1	Battery Learn Cycle	8-26
8.10	Monitoring Rebuilds and Other Processes	8-27

Chapter 9

Maintaining and Managing Storage Configurations

9.1	Initializing a Virtual Drive	9-1
9.2	Running a Group Initialization	9-2
9.3	Running a Consistency Check	9-4
9.3.1	Setting the Consistency Check Properties	9-4
9.3.2	Scheduling a Consistency Check	9-6
9.3.3	Running a Group Consistency Check	9-7
9.4	Scanning for New Drives	9-9
9.5	Rebuilding a Drive	9-9
9.6	Making a Drive Offline or Missing	9-11
9.7	Removing a Drive	9-12
9.8	Upgrading the Firmware	9-12

Chapter 10

CacheCade 2.0 Software

10.1	Logical Drive Property Settings Required for CacheCade	10-1
10.2	Viewing a Logical Drive with CacheCade	10-2
10.3	WebBIOS Configuration for CacheCade	10-3
10.4	MegaRAID Storage Manager Configuration for CacheCade	10-8
10.5	Modifying the CacheCade Virtual Drive Properties	10-11
10.5.1	Enabling SSD Caching on a Virtual Drive	10-12
10.5.2	Disabling SSD Caching on a Virtual Drive	10-13
10.5.3	Enabling or Disabling SSD Caching on Multiple Virtual Drives	10-14
10.5.4	Modifying a CacheCade Drive Group	10-16
10.5.5	Clearing Configuration on CacheCade Pro 2.0 Virtual Drives	10-16
10.5.6	Removing Blocked Access	10-17
10.5.7	Deleting a Virtual Drive With SSD Caching Enabled	10-18
10.6	FastPath Advanced Software	10-19
10.6.1	Setting Fast Path Options	10-19

Appendix A

Events and Messages

Appendix B

SNMP Extension Agent Trap Definitions

B.1	SNMP Traps for RAID Controllers	B-20
B.2	SNMP Traps for Virtual Drives	B-23
B.3	SNMP Traps for Physical Drives	B-24

Appendix C
Glossary

Appendix D
Battery Glossary

Appendix E
Notices

E.1	Trademarks	E-2
E.2	Important Notes	E-3

Figures

1.1	Example of a SAS Direct-Connect Application	1-11
1.2	Example of a SAS RAID Controller Configured with an Expander	1-12
2.1	Example of Disk Striping (RAID 0)	2-8
2.2	Example of Disk Mirroring (RAID 1)	2-9
2.3	Example of Distributed Parity (RAID 5)	2-10
2.4	Example of Disk Spanning	2-11
2.5	RAID 0 Drive Group Example with Two Drives	2-19
2.6	RAID 1 Drive Group	2-20
2.7	RAID 5 Drive Group with Six Drives	2-21
2.8	Example of Distributed Parity across Two Blocks in a Stripe (RAID 6)	2-23
2.9	RAID 10 Level Virtual Drive	2-24
2.10	RAID 50 Level Virtual Drive	2-25
2.11	RAID 60 Level Virtual Drive	2-27
4.1	WebBIOS CU Main Screen	4-3
4.2	WebBIOS Configuration Wizard Screen	4-7
4.3	WebBIOS Configuration Method Screen	4-8
4.4	WebBIOS Disk Group Definition Screen	4-11
4.5	WebBIOS Virtual Drive Definition Screen	4-12
4.6	RAID 0 Configuration Preview	4-15
4.7	WebBIOS Disk Group Definition Screen	4-17
4.8	WebBIOS Virtual Drive Definition Screen	4-18
4.9	RAID 1 Configuration Preview	4-21
4.10	WebBIOS Disk Group Definition Screen	4-23
4.11	WebBIOS Virtual Drive Definition Screen	4-24
4.12	RAID 5 Configuration Preview	4-27
4.13	WebBIOS Disk Group Definition Screen	4-29
4.14	WebBIOS Virtual Drive Definition Screen	4-30
4.15	RAID 6 Configuration Preview	4-33
4.16	WebBIOS Drive Group Definition Screen	4-35
4.17	WebBIOS Span Definition Screen	4-36
4.18	WebBIOS Virtual Drive Definition Screen	4-37
4.19	RAID 10 Configuration Preview	4-40
4.20	WebBIOS Disk Group Definition Screen	4-42
4.21	WebBIOS Span Definition Screen	4-43

4.22	WebBIOS Virtual Drive Definition Screen	4-44
4.23	RAID 50 Configuration Preview	4-47
4.24	WebBIOS Disk Group Definition Screen	4-49
4.25	WebBIOS Span Definition Screen	4-50
4.26	WebBIOS Virtual Drive Definition Screen	4-51
4.27	RAID 60 Configuration Preview	4-54
4.28	Encryption Settings Screen	4-56
4.29	Enable Drive Security - Introduction Screen	4-57
4.30	Enable Drive Security – Enter Security Key ID Screen	4-58
4.31	Enable Drive Security – Enter Security Key	4-59
4.32	Enable Drive Security – Enter Pass Phrase	4-60
4.33	Confirm Enable Drive Security Screen	4-61
4.34	Encryption Settings Screen	4-62
4.35	Change Security Settings – Introduction	4-63
4.36	Change Security Settings – Security Key ID	4-64
4.37	Change Security Settings – Security Key	4-65
4.38	Authenticate Drive Security Key	4-66
4.39	Change Security Settings – Pass Phrase	4-67
4.40	Confirm Change Drive Security Settings	4-68
4.41	Encryption Settings	4-69
4.42	Confirm Disable Drive Security Settings	4-69
4.43	First Controller Properties Screen	4-71
4.44	Second Controller Properties Screen	4-71
4.45	Third Controller Properties Screen	4-72
4.46	Fourth Controller Properties Screen	4-72
4.47	Virtual Drive Screen	4-75
4.48	Physical Drive Screen	4-78
4.49	Physical View Shield State Dialog	4-80
4.50	Logical View Shield State	4-81
4.51	Physical Drive Properties of a Drive in Shield State	4-82
4.52	Shield State Support	4-82
4.53	First Controller Properties Screen	4-83
4.54	Third Controller Properties Screen	4-84
4.55	Battery Module Screen: iBBU08 Battery	4-85
4.56	Fourth Controller Properties Screen	4-88
4.57	Manage Link Speed Screen	4-89
4.58	System Restart Required Message Box	4-90
4.59	Enclosure Properties	4-90

4.60	Additional Enclosure Properties	4-91
4.61	Enclosure More Information - Temperature Sensors	4-91
4.62	Enclosure More Information - Number of Fans	4-92
4.63	Enclosure More Information - Number of Power Supplies	4-92
4.64	SSD Disk Cache Setting in Controller Properties Dialog	4-93
4.65	Virtual Drive Properties	4-94
4.66	Expand Virtual Drive Dialog	4-95
4.67	Virtual Drives Dialog	4-97
4.68	Physical Drive Dialog	4-98
4.69	Mode Selection - Drive Erase	4-99
4.70	Drive Erase Confirm Page	4-99
4.71	Drive Erase Progress	4-100
4.72	Virtual Drive Dialog	4-101
4.73	Virtual Drive Erase Dialog	4-102
4.74	Virtual Drive Dialog	4-103
4.75	Event Information Screen	4-104
4.76	Foreign Configuration Import Screen	4-106
4.77	Foreign Configuration Preview Screen	4-107
4.78	Advanced Operations Screen	4-112
5.1	License Agreement	5-5
5.2	Customer Information Screen	5-6
5.3	Setup Type Screen	5-7
5.4	LDAP Logon Information	5-8
5.5	Setup Type Screen	5-9
5.6	Custom Setup Screen	5-11
5.7	MegaRAID Storage Manager - Host View Window	5-12
6.1	Select Server	6-3
6.2	Configure Host Window	6-4
6.3	Server Login Window	6-6
6.4	LDAP Login	6-8
6.5	Configuring Host LDAP	6-10
6.6	MegaRAID Storage Manager Main Screen	6-11
6.7	MSM Dashboard View	6-12
6.8	MSM Physical View	6-13
6.9	Logical View	6-14
6.10	Chip and Controller Temperature	6-16
6.11	Physical Drive Temperature	6-17
6.12	Physical View Shield State	6-18

6.13	Logical View Shield State	6-19
6.14	Physical Drive Properties of a Drive in Shield State	6-20
6.15	Server Profile View of a Drive in Shield State	6-21
6.16	Parity Size	6-22
6.17	Mirror Data Size	6-23
6.18	Metadata Size	6-24
6.19	Controller Properties – SSD Disk Cache Policy	6-25
6.20	Virtual Drive Settings	6-26
6.21	Virtual Drive Properties	6-27
6.22	Mode Selection - Drive Erase Window	6-28
6.23	Drive Erase Message	6-29
6.24	Group Show Progress	6-30
6.25	Mode Selection – Virtual Drive Erase Dialog	6-31
6.26	Warning Message for Virtual Drive Erase	6-32
6.27	Warning Message for Virtual Drive Erase Without Virtual Drive Delete	6-32
6.28	Group Show Progress – Virtual Drive	6-33
6.29	Group Show Progress Dialog	6-35
6.30	Enclosure Properties	6-36
6.31	Battery Backup Unit Properties for iBBU Batteries	6-38
6.32	Battery Backup Unit Properties for TMM-C Batteries	6-38
6.33	Properties Tab and Graphical View Tab	6-41
7.1	Virtual Drive Creation Menu	7-5
7.2	Virtual Drive Creation Mode	7-6
7.3	Using the Free Capacity of an Existing Drive Group	7-7
7.4	Create Virtual Drive Screen	7-8
7.5	Create Virtual Drive - Summary Window	7-9
7.6	Virtual Drive Creation Menu	7-11
7.7	Virtual Drive Advanced Configuration Mode	7-11
7.8	Create Drive Group Settings Screen	7-12
7.9	Span 0 of Drive Group 0	7-13
7.10	Span 0 and Span 1 of Drive Group 0	7-14
7.11	Virtual Drive Settings Window	7-15
7.12	New Virtual Drive 0	7-16
7.13	Create Virtual Drive Summary Window	7-17
7.14	Enable SSD Caching on New Virtual Drives	7-18
7.15	Drive Security Settings Menu	7-20
7.16	Enable Drive Security - Introduction Screen	7-21

7.17	Enter Security Key ID Screen	7-22
7.18	Enter Security Key Screen	7-23
7.19	Enable Drive Security - Enter Pass Phrase Screen	7-24
7.20	Confirm Create Security Key Screen	7-25
7.21	Change Drive Security Menu	7-26
7.22	Change Security Settings - Introduction Screen	7-27
7.23	Change Security Settings - Security Key ID Screen	7-28
7.24	Change Security Settings - Security Key Screen	7-29
7.25	Authenticate Drive Security Settings Screen	7-30
7.26	Change Security Settings - Pass Phrase Screen	7-30
7.27	Change Drive Security Menu	7-32
7.28	Confirm Disable Drive Security Screen	7-33
7.29	Foreign Configuration Detected Screen	7-34
7.30	Set Adjustable Task Rates Menu	7-38
7.31	Set Adjustable Task Rates	7-39
7.32	Set Virtual Drive Properties Menu	7-41
7.33	Set Virtual Drive Properties Screen	7-42
7.34	Reboot Warning Message	7-44
7.35	Modify Drive Group Wizard	7-44
7.36	Modify Drive Group Wizard	7-46
7.37	Add Drive(s) to the Current Configuration Screen	7-47
7.38	Modify Drive Group Summary Screen	7-48
7.39	Drive Replacement Window	7-50
7.40	Modify Drive Group Wizard	7-51
7.41	Add Drive(s) to the Current Configuration Screen	7-52
7.42	Modify Drive Group Summary Screen	7-53
8.1	Pop-Up Notification	8-3
8.2	E-Mail Notification	8-4
8.3	Event Notification Configuration Menu	8-5
8.4	Configure Alerts Screen	8-5
8.5	Alert Notification Delivery Methods Dialog Box	8-7
8.6	Change Individual Events Dialog Box	8-8
8.7	Change Individual Events Severity Level Menu	8-9
8.8	Mail Server Options	8-11
8.9	Email Settings	8-14
8.10	Controller Properties	8-16
8.11	Drive Information	8-17
8.12	Patrol Read Configuration	8-19

8.13	Virtual Drive Properties	8-22
8.14	Enclosure Properties	8-24
8.15	Battery Backup Unit Properties for iBBU Battery	8-25
8.16	Battery Backup Unit Properties for TMM-C Battery	8-26
8.17	Group Show Progress Window	8-28
9.1	Group Initialization Dialog Box	9-3
9.2	Set Consistency Check Properties Option	9-5
9.3	Set Consistency Check Properties Dialog Box	9-5
9.4	Schedule Consistency Check Dialog	9-6
9.5	Group Consistency Check Dialog Box	9-8
9.6	Update Controller Firmware Dialog	9-13
10.1	Virtual Drive Properties Menu	10-2
10.2	WebBIOS Configuration Wizard Screen	10-4
10.3	CacheCade Array Selection Screen	10-5
10.4	CacheCade Disk Screen	10-6
10.5	CacheCade Configuration Preview	10-7
10.6	WebBIOS Main Menu with a CacheCade Virtual Drive	10-8
10.7	CacheCade SSD Caching Wizard - First Screen	10-9
10.8	Parameters for CacheCade SSD Caching Virtual Drive	10-10
10.9	Create CacheCade - SSD Caching - Summary	10-11
10.10	Set Virtual Drive Properties Window	10-12
10.11	Enable SSD Caching	10-13
10.12	Disable SSD Caching	10-14
10.13	Manage SSD Caching	10-15
10.14	Confirm Clear Configuration	10-16
10.15	Confirm Remove Blocked Access	10-17
10.16	Confirm Delete Virtual Disk	10-18
10.17	Set Virtual Drive Properties Menu	10-20

Tables

1.1	List of Feature on Demand Upgrades	1-6
1.2	Transportable Memory Modules	1-8
1.3	ServeRAID M5100 Series Performance Accelerator for IBM System x Upgrade	1-9
1.4	ServeRAID M5100 Series SSD Caching Enabler for IBM System x Upgrade	1-10
2.1	Types of Parity	2-9
2.2	Spanning for RAID 10, RAID 50, and RAID 60	2-11
2.3	Drive States	2-15
2.4	Virtual Drive States	2-16
2.5	RAID 0 Overview	2-19
2.6	RAID 1 Overview	2-20
2.7	RAID 5 Overview	2-21
2.8	RAID 6 Overview	2-22
2.9	RAID 10 Overview	2-24
2.10	RAID 50 Overview	2-25
2.11	RAID 60 Overview	2-26
2.12	RAID Levels and Fault Tolerance	2-28
2.13	RAID Levels and Performance	2-29
2.14	RAID Levels and Capacity	2-31
2.15	Factors to Consider for Drive Group Configuration	2-34
3.1	SED Terminology	3-2
4.1	WebBIOS CU Toolbar Icons	4-4
4.2	Controller Properties Menu Options	4-73
4.3	BBU Modes	4-85
4.4	Additional Drives Required for RAID-Level Migration	4-111
5.1	Source Port and Destination Port for MSM Components	5-12
6.1	Device Icons	6-39
8.1	Event Severity Levels	8-2
A.1	Event Error Levels	A-1
A.2	Event Messages	A-2
B.1	Supported SNMP Traps and Definitions for RAID Controllers	B-20
B.2	Supported SNMP Traps and Definitions for Virtual Drives	B-23
B.3	MegaRAID Trap Definitions	B-24

Chapter 1

Overview

This guide documents the utilities used to configure, monitor, and maintain IBM ServeRAID-M[®] Serial-attached SCSI (SAS)/Serial-ATA (SATA) controllers with RAID control capabilities and the storage-related devices connected to them. This guide explains how to use the MegaRAID Storage Manager[™] software and the WebBIOS utility. In addition, it documents self-encrypting disks (SED), SAS technology, SATA technology, Solid State Drive (SSD) technology, configuration scenarios, and drive types.

This chapter consists of the following sections:

- [Section 1.1, “SAS Technology”](#)
- [Section 1.2, “Serial-attached SCSI Device Interface”](#)
- [Section 1.3, “Serial ATA Features”](#)
- [Section 1.4, “Solid State Drive Features”](#)
- [Section 1.5, “Integrated MegaRAID Mode and MegaRAID Mode”](#)
- [Section 1.6, “Feature on Demand \(FoD\) Upgrades”](#)
- [Section 1.7, “UEFI 2.0 Support”](#)
- [Section 1.8, “Configuration Scenarios”](#)

1.1 SAS Technology

The ServeRAID-M SAS/SATA RAID controllers are high-performance ServeRAID Serial-attached-SCSI/Serial ATA controllers with RAID control capabilities. ServeRAID-M SAS/SATA RAID controllers provide reliability, high performance, and fault-tolerant disk subsystem management. They are an ideal RAID solution for the internal storage of workgroup, departmental, and enterprise systems. ServeRAID-M

SAS/SATA RAID controllers offer a cost-effective way to implement RAID in a server.

SAS technology brings a wealth of options and flexibility with the use of SAS devices, Serial ATA (SATA) devices, and SSD devices within the same storage infrastructure. These devices bring individual characteristics that make each one a more suitable choice depending on your storage needs. ServeRAID-M gives you the flexibility to combine these two similar technologies on the same controller, within the same enclosure, and in the same virtual drive.

Note: Carefully assess any decision to mix SAS drives and SATA drives within the same *virtual drives*. Although you can mix drives, IBM strongly discourages the practice. This recommendation applies to both HDDs and SSDs.

The ServeRAID-M SAS/SATA RAID controllers are based on the SAS IC technology and proven RAID technology. As second-generation PCI Express SAS RAID controllers, the ServeRAID-M SAS/SATA RAID controllers address the growing demand for increased data throughput and scalability requirements across midrange and enterprise-class server platforms. IBM offers a family of ServeRAID-M SAS/SATA RAID controllers addressing the needs for both internal and external solutions.

The controllers support the SAS protocol as described in the *Serial Attached SCSI Standard, Version 2.0*. In addition, the controller supports the SATA II protocol defined by the *Serial ATA specification, Version 2.0*, and the SATA III protocol defined by the *Serial ATA Specification, Version 3.0*. SATA III is an extension to SATA II.

Each port on the SAS RAID controller supports SAS devices, SATA devices, or SSD devices using the following protocols:

- SAS Serial SCSI Protocol (SSP), which enables communication with other SAS devices
- SATA, which enables communication with other SATA devices
- Serial Management Protocol (SMP), which communicates topology management information directly with an attached SAS expander device
- Serial Tunneling Protocol (STP), which enables communication with a SATA device through an attached expander

1.2 Serial-attached SCSI Device Interface

SAS is a serial, point-to-point, enterprise-level device interface that leverages the proven SCSI protocol set. SAS is a convergence of the advantages of SATA, SCSI, and Fibre Channel, and is the future mainstay of the enterprise and high-end workstation storage markets. SAS offers a higher bandwidth per pin than parallel SCSI, and it improves signal and data integrity.

The SAS interface uses the proven SCSI command set to ensure reliable data transfers, while providing the connectivity and flexibility of point-to-point serial data transfers. The serial transmission of SCSI commands eliminates clock-skew challenges. The SAS interface provides improved performance, simplified cabling, smaller connectors, lower pin count, and lower power requirements when compared to parallel SCSI.

SAS controllers leverage a common electrical and physical connection interface that is compatible with Serial ATA technology. The SAS protocols and the SATA protocols use a thin, 7-wire connector instead of the 68-wire SCSI cable or 26-wire ATA cable. The SAS/SATA connector and cable are easier to manipulate, allow connections to smaller devices, and do not inhibit airflow. The point-to-point SATA architecture eliminates inherent difficulties created by the legacy ATA master-slave architecture, while maintaining compatibility with existing ATA firmware.

1.3 Serial ATA Features

The SATA bus is a high-speed, internal bus that provides a low pin count, low voltage level bus for device connections between a host controller and a SATA device.

Note: The ServeRAID SAS/SATA controllers support SATA, SATA II, and SATA III technologies.

The following list describes the SATA features of the RAID controllers:

- Supports SATA III data transfers of 6.0 Gbits/s
- Supports STP data transfers of 6.0 Gbits/s

- Provides a serial, point-to-point storage interface
- Simplifies cabling between devices
- Eliminates the master-slave construction used in parallel ATA
- Allows addressing of multiple SATA targets through an expander
- Allows multiple initiators to address a single target (in a fail-over configuration) through an expander

1.4 Solid State Drive Features

The firmware supports Solid State Drives attached to ServeRAID-M SAS controllers. The features and operations for SSDs are the same as for hard disk drives (HDDs), and these drives are expected to behave like SATA HDDs or SAS HDDs. The major advantages of SSDs include:

- High random read speed (because there is no read-write head to move)
- High performance-to-power ratio, as these drives have very low power consumption compared to HDDs
- Low latency
- High mechanical reliability
- Lower weight and size (for low-capacity SSD drives)

The WebBIOS Configuration Utility and the MegaRAID Storage Manager utility display the SSDs by the type, either SAS or SATA. For example, a SATA SSD drive displays as “SSD (SATA)”. HDDs are identified simply as “SAS” or “SATA”.

Note: ServeRAID-M implements support for only those SATA SSD drives which support ATA-8 ACS compliance.

You can choose whether to allow a virtual drive to consist of both SSD devices and HDDs. For a virtual drive that consists of SSDs only, you can choose whether to allow SAS SSD drives and SATA SSD drives in that virtual drive. For virtual drives that have both SSDs and HDDs, you can choose whether to mix SAS and SATA HDD drives with SAS and SATA SSD devices in various combinations.

Note: Support for SATA SSD drives applies only to those drives that support ATA-8 ACS compliance.

1.4.1 Solid State Drive Guard

SSDs are known for their reliability and performance. SSD Guard™ increases the reliability of SSDs by automatically copying data from a drive with potential to fail to a designated hot spare or newly inserted drive. Because SSDs are very reliable, non-redundant RAID 0 configurations are much more common than in the past. SSD Guard offers added data protection for RAID 0 configurations.

SSD Guard works by looking for a predictive failure while monitoring the SSD S.M.A.R.T. (Self-Monitoring, Analysis and Reporting Technology) error log. If errors indicate that an SSD failure is imminent, ServeRAID-M starts a rebuild to preserve the data on the SSD and sends appropriate warning event notifications.

1.5 Integrated MegaRAID Mode and MegaRAID Mode

Some ServeRAID SAS/SATA controllers function in either integrated MegaRAID (iMR) mode or in MegaRAID (MR) mode.

Integrated MegaRAID is a highly integrated, low-cost RAID solution made possible by Fusion-MPT™ architecture. Integrated MegaRAID is a processor-based, hardware RAID solution designed for system environments requiring redundancy and high availability where a full-featured RAID implementation is not desired or might be cost prohibitive.

The major advantage of Integrated MegaRAID is that iMR provides RAID at the processor level, so it does not burden the CPU, which allows for more efficient operation.

The major advantage of MegaRAID mode is that the MR mode supports more RAID levels than iMR mode. iMR mode supports RAID levels 0, 1, 5, 10, and 50. MR mode supports RAID levels 0, 1, 5, 6, 10, 50, and 60.

Note: For the ServeRAID M1100 SAS/SATA controllers and the ServeRAID M5100 SAS/SATA controllers, iMR RAID 5 requires purchase of the Feature on Demand (FoD) upgrade

Note: For the ServeRAID M1100 SAS/SATA controllers and the ServeRAID M5100 SAS/SATA controllers, MegaRAID RAID 5/50 requires a transportable memory module (3 options) or the Feature on Demand upgrade.

Note: For the ServeRAID M1100 SAS/SATA controllers and the ServeRAID M5100 SAS/SATA controllers, MegaRAID RAID 6/60 requires a transportable memory module (3 options) *and* the Feature on Demand upgrade.

See [Section 1.6, “Feature on Demand \(FoD\) Upgrades”](#) for information about these upgrades.

See [Section 2.5.1, “Summary of RAID Levels”](#) for information about the supported RAID levels.

1.6 Feature on Demand (FoD) Upgrades

To use RAID levels 5, 6, 50, or 60, the FastPath feature, or the CacheCade 2.0 feature with selected controllers, you must install a Feature on Demand (FoD) upgrade and/or a transportable memory module (TMM). The following sections describe these upgrades, required installations, and supported controllers.

The following table lists the FoD upgrades available.

Table 1.1 List of Feature on Demand Upgrades

IBM Option	Official Name	Functionality
81Y4542	ServeRAID M1100 Series Zero Cache/RAID 5 Upgrade for IBM System x	iMR RAID 5 + SED
81Y4544	ServeRAID M5100 Series Zero Cache/RAID 5 Upgrade for IBM System x	iMR RAID 5 + SED
81Y4546	ServeRAID M5100 Series RAID 6 Upgrade for IBM System x	MegaRAID RAID 6
90Y4273	ServeRAID M5100 Series Performance Accelerator for IBM System x	FastPath
90Y4318	ServeRAID M5100 Series SSD Caching Enabler for IBM System x	CacheCade 2.0

1.6.1 Feature on Demand: iMR RAID 5 + Self-Encrypting Disk Upgrade

The ServeRAID M5100 Series Zero Cache/RAID 5 Upgrade for IBM System x supports iMR RAID levels 5 and 50, and self-encrypting disks for the following ServeRAID SAS/SATA controllers:

- ServeRAID M5110 SAS/SATA controller for IBM System x
- ServeRAID M5110e SAS/SATA controller for IBM System x
- ServeRAID M5120 SAS/SATA controller for IBM System x

The ServeRAID M1100 Series Zero Cache/RAID 5 Upgrade for IBM System x supports iMR RAID levels 5 and 50, and self-encrypting disks for the following ServeRAID SAS/SATA controller:

- The ServeRAID M1115 SAS/SATA controller for IBM System x

The SED feature offers the ability to encrypt data on disks and use disk-based key management to provide data security. With the SED feature, data is encrypted by the drives. You can designate which data to encrypt at the individual virtual drive (VD) level.

This solution provides data protection in the event of theft or loss of physical drives. With self-encrypting disks, if you remove a drive from its storage system or the server it is housed in, the data on that drive is encrypted and useless to anyone who attempts to access without the the appropriate security authorization.

See [Chapter 3, “Self-Encrypting Disk Feature”](#) for more information about the self-encrypting disk feature.

See [Section 2.5.1, “Summary of RAID Levels”](#) for information about the supported RAID levels.

1.6.2 Feature on Demand: MegaRAID RAID 6/60 Upgrade

The ServeRAID M5100 Series RAID 6 Upgrade for System x is a Feature on Demand that supports MegaRAID RAID levels 6 and 60 for the following ServeRAID SAS/SATA controllers:

- ServeRAID M5110 SAS/SATA controller for IBM System x
- ServeRAID M5110e SAS/SATA controller for IBM System x
- ServeRAID M5120 SAS/SATA controller for IBM System x

Install the ServeRAID M5100 Series RAID 6 Upgrade for System FoD *and* any of the transportable memory modules in the following table to upgrade to support RAID 6 and 60.

Table 1.2 Transportable Memory Modules

IBM Option	Official Name	Functionality
81Y4584	ServeRAID M5100 Series 512MB Cache/RAID 5 Upgrade for IBM System x	MegaRAID RAID 5 + SED
81Y4587	ServeRAID M5100 Series 512MB Flash/RAID 5 Upgrade for IBM System x	MegaRAID RAID 5 + SED
81Y4559	ServeRAID M5100 Series 1GB Flash/RAID 5 Upgrade for IBM System x	MegaRAID RAID 5 + SED

1.6.3 Feature on Demand: FastPath Upgrade

ServeRAID M5100 Series Performance Accelerator for IBM System x is a Feature on Demand that supports the FastPath feature for the following ServeRAID SAS/SATA controllers:

- ServeRAID M5110 SAS/SATA controller for IBM System x
- ServeRAID M5110e SAS/SATA controller for IBM System x
- ServeRAID M5120 SAS/SATA controller for IBM System x

The FastPath feature is a high-performance IO accelerator for SSD drive groups connected to a ServeRAID controller card. FastPath software combined with SSDs delivers a performance advantage over HDD installations and consumes less power. This feature dramatically boosts storage subsystem bandwidth and overall transactional application performance when used with a 6Gb/s ServeRAID SAS/SATA controller.

The FastPath feature delivers optimization of SSD virtual disk groups to enable read and write IOPS three times greater than ServeRAID controllers not using FastPath technology. In addition, the FastPath advanced software is faster and more cost-effective than current flash-based adapter card solutions. This feature is faster and more cost-effective than current Flash-based adapter card solutions.

To use the FastPath feature, you must install the ServeRAID M5100 Series Performance Accelerator for IBM System x FoD and a transportable memory module. See [Section 1.6.2, “Feature on Demand:](#)

[MegaRAID RAID 6/60 Upgrade](#)” for the list of transportable memory modules.

The following table lists the option number and functionality for the ServeRAID M5100 Series Performance Accelerator for IBM System x.

Table 1.3 ServeRAID M5100 Series Performance Accelerator for IBM System x Upgrade

IBM Option	Official Name	Functionality
81Y4544	ServeRAID M5100 Series Performance Accelerator for IBM System x	FastPath

1.6.4 Feature on Demand: CacheCade 2 Upgrade

The ServeRAID M5100 Series SSD Caching Enabler for IBM System x is a Feature on Demand that supports the CacheCade 2 feature for the following ServeRAID SAS/SATA controllers:

- ServeRAID M5110 SAS/SATA controller for IBM System x
- ServeRAID M5110e SAS/SATA controller for IBM System x
- ServeRAID M5120 SAS/SATA controller for IBM System x

The CacheCade 2 Software feature improves I/O performance to meet the needs of high-performing Solid State Drives (SSDs). In addition, this feature benefits hard disk drives (HDDs). To support full-throughput for multiple direct-attached SSDs, the CacheCade 2 Software feature reduces I/O-processing overhead for the ServeRAID controllers. CacheCade 2 Software offers performance equivalent to Flash-based controllers.

To use the CacheCade 2.0 feature, you must install the ServeRAID M5100 Series SSD Caching Enabler for IBM System x FoD and a transportable memory module. See [Section 1.6.2, “Feature on Demand: MegaRAID RAID 6/60 Upgrade”](#) for the list of transportable memory modules.

The following table lists the option number and functionality for the ServeRAID M5100 Series SSD Caching Enabler for IBM System x.

Table 1.4 ServeRAID M5100 Series SSD Caching Enabler for IBM System x Upgrade

IBM Option	Official Name
90Y4318	ServeRAID M5100 Series SSD Caching Enabler for IBM System x

1.7 UEFI 2.0 Support

Significant challenges face operating system and platform developers to innovate using the legacy PC-AT BIOS boot environment. These include memory constraints, maintenance challenges, and increased complexities due to a lack of industry-wide standards.

To handle these challenges, the Unified Extensible Firmware Interface (UEFI) was developed to do the following:

- Define a clean interface between operating systems and the hardware platform at boot time
- Support an architecture-independent mechanism for initializing add-in cards.

UEFI 2.0 provides ServeRAID-M customers with expanded platform support. The UEFI 2.0 driver, a boot service device driver, handles block IO requests and SCSI pass-through commands (SPTs), and offers the ability to launch pre-boot ServeRAID-M management applications through a driver configuration protocol (DCP). The UEFI driver also supports driver diagnostic protocol, which allows administrators to access pre-boot diagnostics.

1.8 Configuration Scenarios

You can use the SAS RAID controllers in the following three main scenarios:

- **Low-end, internal SATA configurations:** In this configuration, use the RAID controller as a high-end SATA compatible controller that


connects up to eight disks either directly or through a port expander. This configuration is mostly for low-end or entry servers. Enclosure management is provided through out-of-band I²C bus. Side bands of both types of internal SAS connectors support the SFF-8485 (SGPIO) interface.

- **Midrange internal SAS configurations:** This configuration is like the internal SATA configurations, but with high-end disks. This configuration is more suitable for low-range to midrange servers.
- **High-end external SAS/SATA configurations:** This configuration is for both internal connectivity and external connectivity, using SATA drives, SAS drives, or both. External enclosure management is supported through in-band, SCSI-enclosed storage. The configuration must support STP and SMP.

The following figure shows a direct-connect configuration. The Inter-IC (I²C) interface communicates with peripherals. The external memory bus provides a 32-bit memory bus, parity checking, and chip select signals for pipelined synchronous burst static random access memory (PSBRAM), nonvolatile static random access memory (NVSRAM), and Flash ROM.

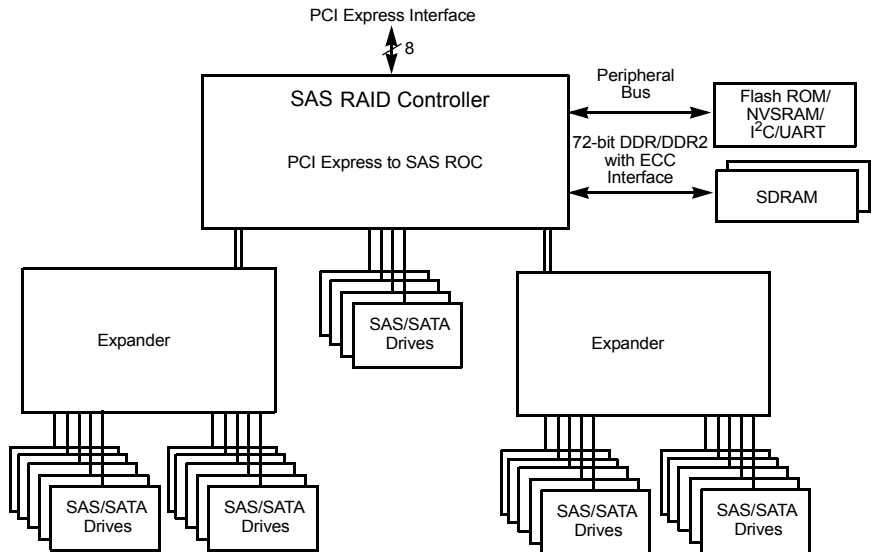
Note: The external memory bus is 64-bit for the Server RAID-MR10il RAID controller and the Server RAID-MR10M RAID controller.

Figure 1.1 Example of a SAS Direct-Connect Application



The following figure shows an example of a SAS RAID controller configured with an expander that is connected to SAS disks, SATA disks, or both.

Figure 1.2 Example of a SAS RAID Controller Configured with an Expander



1.8.1 Valid Drive Mix Configurations

You *cannot* have both SSDs and HDDs in a virtual drive. For a virtual drive that consists of SSDs only, you can choose whether to allow both SAS SSD drives and SATA SSD drives in that virtual drive.

Note: The valid drive mix applies to hot spares, also. For hot spare information, see [Section 2.4.13, “Hot Spares”](#).

Chapter 2

Introduction to RAID

This chapter describes RAID (Redundant Array of Independent Disks), RAID functions and benefits, RAID components, RAID levels, and configuration strategies. In addition, it defines the RAID availability concept, and offers tips for configuration planning.

This chapter consists of the following sections:

- [Section 2.1, “RAID Description”](#)
 - [Section 2.2, “RAID Benefits”](#)
 - [Section 2.3, “RAID Functions”](#)
 - [Section 2.4, “Components and Features”](#)
 - [Section 2.5, “RAID Levels”](#)
 - [Section 2.6, “RAID Configuration Strategies”](#)
 - [Section 2.7, “RAID Availability”](#)
 - [Section 2.8, “Configuration Planning”](#)
-

2.1 RAID Description

RAID is an array, or group of multiple independent physical drives that provide high performance and fault tolerance. A RAID drive group improves I/O (input/output) performance and reliability. The RAID drive group appears to the host computer as a single storage unit or as multiple virtual units. I/O is expedited because several drives can be accessed simultaneously.

2.2 RAID Benefits

RAID drive groups improve data storage reliability and fault tolerance compared to single-drive storage systems. Data loss resulting from a drive failure can be prevented by reconstructing missing data from the remaining drives. RAID has gained popularity because it improves I/O performance and increases storage subsystem reliability.

2.3 RAID Functions

Virtual drives are drive groups or spanned drive groups that are available to the operating system. The storage space in a virtual drive is spread across all of the drives in the drive group.

Your drives must be organized into virtual drives in a drive group and they must be able to support the RAID level that you select. Below are some common RAID functions:

- Creating hot spare drives
- Configuring drive groups and virtual drives
- Initializing one or more virtual drives
- Accessing controllers, virtual drives, and drives individually
- Rebuilding failed drives
- Verifying that the redundancy data in virtual drives using RAID level 1, 5, 6, 10, 50, or 60 is correct
- Reconstructing virtual drives after changing RAID levels or adding a drive to a drive group
- Selecting a host controller to work on

2.4 Components and Features

RAID levels describe a system for ensuring the availability and redundancy of data stored on large disk subsystems. See [Section 2.5, “RAID Levels,”](#) for detailed information about RAID levels. The following

subsections describes the components of RAID drive groups and RAID levels.

2.4.1 Physical Array

A physical array is a group of drives. The drives are managed in partitions known as virtual drives.

2.4.2 Virtual Drive

A virtual drive is a partition in a drive group that is made up of contiguous data segments on the drives. A virtual drive can consist of an entire drive group, more than one entire drive group, a part of a drive group, parts of more than one drive group, or a combination of any two of these conditions.

2.4.3 RAID Drive Group

A RAID drive group is one or more drives controlled by the RAID controller.

2.4.4 Fault Tolerance

Fault tolerance is the capability of the subsystem to undergo a drive failure or failures without compromising data integrity, and processing capability. The RAID controller provides this support through redundant drive groups in RAID levels 1, 5, 6, 10, 50, and 60. The system can still work properly even with drive failure in a drive group, though performance can be degraded to some extent.

In a span of RAID 1 drive groups, each RAID 1 drive group has two drives and can tolerate one drive failure. The span of RAID 1 drive groups can contain up to 32 drives, and tolerate up to 16 drive failures - one in each drive group. A RAID 5 drive group can tolerate one drive failure in each RAID 5 drive group. A RAID 6 drive group can tolerate up to two drive failures.

Each spanned RAID 10 virtual drive can tolerate multiple drive failures, as long as each failure is in a separate drive group. A RAID 50 virtual drive can tolerate two drive failures, as long as each failure is in a separate drive group. RAID 60 drive groups can tolerate up to two drive failures in each drive group.

Note: RAID level 0 is not fault tolerant. If a drive in a RAID 0 drive group fails, the whole virtual drive (all drives associated with the virtual drive) will fail.

Fault tolerance is often associated with system availability because it allows the system to be available during the failures. However, this means that it is also important for the system to be available during the repair of the problem.

A hot spare is an unused drive that, in case of a drive failure in a redundant RAID drive group, can be used to rebuild the data and re-establish redundancy. After the hot spare is automatically moved into the RAID drive group, the data is automatically rebuilt on the hot spare drive. The RAID drive group continues to handle requests while the rebuild occurs.

Auto-rebuild allows a failed drive to be replaced and the data automatically rebuilt by “hot-swapping” the drive in the same drive bay. The RAID drive group continues to handle requests while the rebuild occurs.

2.4.4.1 Multipathing

The firmware provides support for detecting and using multiple paths from the RAID controllers to the SAS devices that are in enclosures. Devices connected to enclosures have multiple paths to them. With redundant paths to the same port of a device, if one path fails, another path can be used to communicate between the controller and the device. Using multiple paths with load balancing, instead of a single path, can increase reliability through redundancy.

Applications show the enclosures and the drives connected to the enclosures. The firmware dynamically recognizes new enclosures added to a configuration along with their contents (new drives). In addition, the firmware dynamically adds the enclosure and its contents to the management entity currently in-use.

Multipathing provides the following features:

- Support for failover, in the event of path failure
- Auto-discovery of new or restored paths while the system is online, and reversion to system load balancing policy

- Measurable bandwidth improvement to the multi-path device
- Support for changing the load balancing path while the system is online

The firmware determines whether enclosure modules (ESMs) are part of the same enclosure. When a new enclosure module is added (allowing multi-path) or removed (going single path), an Asynchronous Event Notification (AEN) is generated. Alerts about drives contain correct information about the "enclosure", when the drives are connected by multiple paths. The enclosure module detects partner ESMs and issue events appropriately.

In a system with two ESMs, you can replace one of the ESMs without affecting the virtual drive availability. For example, the controller can run heavy I/Os, and when you replace one of the ESM modules, I/Os should not stop. The controller uses different paths to balance the load on the entire system.

In the MegaRAID Storage Manager utility, when multiple paths are available to a drive, the drive information will show only one enclosure. The utility shows that a redundant path is available to a drive. All drives with a redundant path display this information. The firmware supports online replacement of enclosure modules.

2.4.5 Consistency Check

The Consistency Check operation verifies correctness of the data in virtual drives that use RAID levels 1, 5, 6, 10, 50, and 60. (RAID 0 does not provide data redundancy). For example, in a system with parity, checking consistency means computing the data on one drive and comparing the results to the contents of the parity drive.

Note: It is recommended that you perform a consistency check at least once a month.

2.4.6 Copyback

The copyback feature allows you to copy data from a source drive of a virtual drive to a destination drive that is not a part of the virtual drive. Copyback is often used to create or restore a specific physical configuration for a drive group (for example, a specific arrangement of

drive group members on the device I/O buses). Copyback can be run automatically or manually.

Typically, when a drive fails or is expected to fail, the data is rebuilt on a hot spare. The failed drive is replaced with a new drive. Then the data is copied from the hot spare to the new drive, and the hot spare reverts from a rebuild drive to its original hot spare status. The copyback operation runs as a background activity, and the virtual drive is still available online to the host.

Copyback is also initiated when the first Self-Monitoring Analysis and Reporting Technology (SMART) error occurs on a drive that is part of a virtual drive. The destination drive is a hot spare that qualifies as a rebuild drive. The drive with the SMART error is marked as "failed" only after the successful completion of the copyback. This avoids putting the drive group in degraded status.

Note: During a copyback operation, if the drive group involved in the copyback is deleted because of a virtual drive deletion, the destination drive reverts to an Unconfigured Good state or hot spare state.

Order of Precedence –

In the following scenarios, rebuild takes precedence over the copyback operation:

1. If a copyback operation is already taking place to a hot spare drive, and any virtual drive on the controller degrades, the copyback operation aborts, and a rebuild starts. The rebuild changes the virtual drive to the optimal state.
2. The rebuild operation takes precedence over the copyback operation when the conditions exist to start both operations. For example:
 - a. Where the hot spare is not configured (or unavailable) in the system.
 - b. There are two drives (both members of virtual drives), with one drive exceeding the SMART error threshold, and the other failed.
 - c. If you add a hot spare (assume a global hot spare) during a copyback operation, the copyback is aborted, and the rebuild operation starts on the hot spare.

2.4.7 Background Initialization

Background initialization is a consistency check that is forced when you create a virtual drive. The difference between a background initialization and a consistency check is that a background initialization is forced on new virtual drives. This is an automatic operation that starts 5 minutes after you create the virtual drive.

Background initialization is a check for media errors on the drives. It ensures that striped data segments are the same on all drives in a drive group. The default and recommended background initialization rate is 30 percent. Before you change the rebuild rate, you must stop the background initialization or the rate change will not affect the background initialization rate. After you stop background initialization and change the rebuild rate, the rate change takes effect when you restart background initialization.

2.4.8 Patrol Read

Patrol read involves the review of your system for possible drive errors that could lead to drive failure and then action to correct errors. The goal is to protect data integrity by detecting drive failure before the failure can damage data. The corrective actions depend on the drive group configuration and the type of errors.

Patrol read starts only when the controller is idle for a defined period of time and no other background tasks are active, though it can continue to run during heavy I/O processes.

You can use the MegaRAID Command Tool or the MegaRAID Storage Manager to select the patrol read options, which you can use to set automatic or manual operation, or disable patrol read. See [Section 8.6, “Running a Patrol Read”](#).


2.4.9 Disk Striping

Disk striping allows you to write data across multiple drives instead of just one drive. Disk striping involves partitioning each drive storage space into stripes that can vary in size from 8 KB to 1024 KB. These stripes are interleaved in a repeated sequential manner. The combined storage space is composed of stripes from each drive. It is

recommended that you keep stripe sizes the same across RAID drive groups.

For example, in a four-disk system using only disk striping (used in RAID level 0), segment 1 is written to disk 1, segment 2 is written to disk 2, and so on. Disk striping enhances performance because multiple drives are accessed simultaneously, but disk striping does not provide data redundancy.

Figure 2.1 Example of Disk Striping (RAID 0)



2.4.9.1 Stripe Width

Stripe width is the number of drives involved in a drive group where striping is implemented. For example, a four-disk drive group with disk striping has a stripe width of four.

2.4.9.2 Stripe Size

The stripe size is the length of the interleaved data segments that the RAID controller writes across multiple drives, not including parity drives. For example, consider a stripe that contains 64 KB of disk space and has 16 KB of data residing on each disk in the stripe. In this case, the stripe size is 64 KB and the strip size is 16 KB.

2.4.9.3 Strip Size


The strip size is the portion of a stripe that resides on a single drive.

2.4.10 Disk Mirroring

With mirroring (used in RAID 1 and RAID 10), data written to one drive is simultaneously written to another drive. The primary advantage of disk mirroring is that it provides 100 percent data redundancy. Because the contents of the disk are completely written to a second disk, data is not lost if one disk fails. In addition, both drives contain the same data at all times, so either disk can act as the operational disk. If one disk fails, the contents of the other disk can be used to run the system and reconstruct the failed disk.

Disk mirroring provides 100 percent redundancy, but is expensive because each drive in the system must be duplicated. [Figure 2.2](#) shows an example of disk mirroring.

Figure 2.2 Example of Disk Mirroring (RAID 1)



2.4.11 Parity


Parity generates a set of redundancy data from two or more parent data sets. The redundancy data can be used to reconstruct one of the parent data sets in the event of a drive failure. Parity data does not fully duplicate the parent data sets, but parity generation can slow the write process. In RAID, this method is applied to entire drives or stripes across all of the drives in a drive group. [Table 2.1](#) describes the types of parity.

Table 2.1 Types of Parity

Parity Type	Description
Dedicated	The parity data on two or more drives is stored on an additional drive.
Distributed	The parity data is distributed across more than one drive in the system.

RAID 5 combines distributed parity with disk striping. If a single drive fails, it can be rebuilt from the parity and the data on the remaining drives. An example of a RAID 5 drive group is shown in [Figure 2.3](#). RAID 5 uses parity to provide redundancy for one drive failure without duplicating the contents of entire drives. RAID 6 uses distributed parity and disk striping, also, but adds a second set of parity data so that it can survive up to two drive failures.

Figure 2.3 Example of Distributed Parity (RAID 5)



Note: Parity is distributed across all drives in the drive group.


2.4.12 Disk Spanning

Disk spanning allows multiple drives to function like one big drive. Spanning overcomes lack of disk space and simplifies storage management by combining existing resources or adding relatively inexpensive resources. For example, four 20 GB drives can be combined to appear to the operating system as a single 80 GB drive.

Spanning alone does not provide reliability or performance enhancements. Spanned virtual drives must have the same stripe size and must be contiguous. In [Figure 2.4](#), RAID 1 drive groups are turned into a RAID 10 drive group.

Note: Make sure that the spans are in different backplanes, so that if one span fails, you do not lose the whole drive group.

Figure 2.4 Example of Disk Spanning



Note: Spanning two contiguous RAID 0 virtual drives does not produce a new RAID level or add fault tolerance. It does increase the capacity of the virtual drive and improves performance by doubling the number of spindles.

2.4.12.1 Spanning for RAID 10, RAID 50, and RAID 60

[Table 2.2](#) describes how to use spanning to configure RAID 10, 50, and 60 virtual drives. The virtual drives must have the same stripe size and the maximum number of spans is eight. The full drive capacity is used when you span virtual drives; you cannot specify a smaller drive capacity.

See [Chapter 7, “Configuration,”](#) for detailed procedures for configuring drive groups and virtual drives, and spanning the drives.

Table 2.2 Spanning for RAID 10, RAID 50, and RAID 60

Level	Description
10	Configure RAID 10 by spanning two contiguous RAID 1 virtual drives, up to the maximum number of supported devices for the controller. RAID 10 supports a maximum of eight spans. You must use an even number of drives in each RAID virtual drive in the span. The RAID 1 virtual drives must have the same stripe size.
50	Configure RAID 50 by spanning two contiguous RAID 5 virtual drives. The RAID 5 virtual drives must have the same stripe size.
60	Configure RAID 60 by spanning two contiguous RAID 6 virtual drives. The RAID 6 virtual drives must have the same stripe size.

2.4.13 Hot Spares

A hot spare is an extra, unused drive that is part of the disk subsystem. It is usually in standby mode, ready for service if a drive fails. Hot spares permit you to replace failed drives without system shutdown or user intervention. MegaRAID SAS RAID controllers can implement automatic and transparent rebuilds of failed drives using hot spare drives, providing a high degree of fault tolerance and zero downtime.

Note: When running RAID 0 and RAID 5 virtual drives on the same set of drives (a sliced configuration), a rebuild to a hot spare will not occur after a drive failure until the RAID 0 virtual drive is deleted.

The RAID management software allows you to specify drives as hot spares. When a hot spare is needed, the RAID controller assigns the hot spare that has a capacity closest to and at least as great as that of the failed drive to take the place of the failed drive. The failed drive is removed from the virtual drive and marked ready awaiting removal once the rebuild to a hot spare begins. You can make hot spares of the drives that are not in a RAID virtual drive.

You can use the RAID management software to designate the hotspare to have enclosure affinity, meaning that if there are drive failures present on a split backplane configuration, the hotspare will be used first on the backplane side that it resides in.

If the hotspare is designated as having enclosure affinity, it will attempt to rebuild any failed drives on the backplane that it resides in before rebuilding any other drives on other backplanes.

Note: If a rebuild to a hot spare fails for any reason, the hot spare drive is marked as "failed". If the source drive fails, both the source drive and the hot spare drive is marked as "failed".

There are two types of hot spares:

- Global hot spare
- Dedicated hot spare

2.4.13.1 Global Hot Spare

A global hot spare drive can be used to replace any failed drive in a redundant drive group as long as its capacity is equal to or larger than the coerced capacity of the failed drive. A global hot spare defined on any channel should be available to replace a failed drive on both channels.

2.4.13.2 Dedicated Hot Spare

A dedicated hot spare can be used to replace a failed drive only in a selected drive group. One or more drives can be designated as a member of a spare drive pool. The most suitable drive from the pool is selected for fail over. A dedicated hot spare is used before one from the global hot spare pool.

Hot spare drives can be located on any RAID channel. Standby hot spares (not being used in RAID drive group) are polled every 60 seconds at a minimum, and their status made available in the drive group management software. RAID controllers offer the ability to rebuild with a disk that is in a system, but not initially set to be a hot spare.

Observe the following parameters when using hot spares:

- Hot spares are used only in drive groups with redundancy: RAID levels 1, 5, 10, and 50.
- A hot spare connected to a specific RAID controller can be used to rebuild a drive that is connected to the same controller only.
- You must assign the hot spare to one or more drives through the controller BIOS or use drive group management software to place it in the hot spare pool.
- A hot spare must have free space equal to or greater than the drive it replaces. For example, to replace an 18 GB drive, the hot spare must be 18 GB or larger.

2.4.14 Disk Rebuilds

When a drive in a RAID drive group fails, you can rebuild the drive by recreating the data that was stored on the drive before it failed. The RAID controller recreates the data using the data stored on the other drives in the drive group. Rebuilding can be done only in drive groups with data redundancy, which includes RAID 1, 5, 6, 10, 50, and 60 drive groups.

The RAID controller uses hot spares to rebuild failed drives automatically and transparently, at user-defined rebuild rates. If a hot spare is available, the rebuild can start automatically when a drive fails. If a hot spare is not available, the failed drive must be replaced with a new drive so that the data on the failed drive can be rebuilt.

The failed drive is removed from the virtual drive and marked ready awaiting removal when the rebuild to a hot spare begins. If the system goes down during a rebuild, the RAID controller automatically restarts the rebuild after the system reboots.

Note: When the rebuild to a hot spare begins, the failed drive is often removed from the virtual drive before management applications detect the failed drive. When this occurs, the events logs show the drive rebuilding to the hot spare without showing the failed drive. The formerly failed drive will be marked as "ready" after a rebuild begins to a hot spare.

Note: If a source drive fails during a rebuild to a hot spare, the rebuild fails, and the failed source drive is marked as offline. In addition, the rebuilding hot spare drive is changed back to a hot spare. After a rebuild fails because of a source drive failure, the dedicated hot spare is still dedicated and assigned to the correct drive group, and the global hot spare is still global.

An automatic drive rebuild will not start if you replace a drive during a RAID-level migration. The rebuild must be started manually after the expansion or migration procedure is complete.

2.4.15 Rebuild Rate

The rebuild rate is the percentage of the compute cycles dedicated to rebuilding failed drives. A rebuild rate of 100 percent means that the system gives priority to rebuilding the failed drives.

The rebuild rate can be configured between 0 percent and 100 percent. At 0 percent, the rebuild is done only if the system is not doing anything else. At 100 percent, the rebuild has a higher priority than any other system activity. Using 0 or 100 percent is not recommended. The default rebuild rate is 30 percent.

2.4.16 Hot Swap

A hot swap is the manual replacement of a defective drive unit while the computer is still running. When a new drive has been installed, a rebuild will occur automatically if:

- The newly inserted drive is the same capacity as or larger than the failed drive
- It is placed in the same drive bay as the failed drive it is replacing

The RAID controller can be configured to detect the new drives and rebuild the contents of the drive automatically.

2.4.17 Drive States

A drive state is a property indicating the status of the drive. The drive states are described in [Table 2.3](#).

Table 2.3 Drive States

State	Description
Online	A drive that can be accessed by the RAID controller and is part of the virtual drive.
Unconfigured Good	A drive that is functioning normally but is not configured as a part of a virtual drive or as a hot spare.
Hot Spare	A drive that is powered up and ready for use as a spare in case an online drive fails.

Table 2.3 Drive States (Cont.)

State	Description
Failed	A drive that was originally configured as Online or Hot Spare, but on which the firmware detects an unrecoverable error.
Rebuild	A drive to which data is being written to restore full redundancy for a virtual drive.
Unconfigured Bad	A drive on which the firmware detects an unrecoverable error; the drive was Unconfigured Good or the drive could not be initialized.
Missing	A drive that was Online but which has been removed from its location.
Offline	A drive that is part of a virtual drive but which has invalid data as far as the RAID configuration is concerned. Note: When a virtual drive with cached data goes offline, the cache for the virtual drive is discarded. Because the virtual drive is offline, the cache cannot be saved.

2.4.18 Virtual Drive States

The virtual drive states are described in [Table 2.4](#).

Table 2.4 Virtual Drive States

State	Description
Optimal	The virtual drive operating condition is good. All configured drives are online.
Degraded	The virtual drive operating condition is not optimal. One of the configured drives has failed or is offline.
Partially Degraded	The operating condition in a RAID 6 virtual drive is not optimal. One of the configured drives has failed or is offline. RAID 6 can tolerate up to two drive failures.
Failed	The virtual drive has failed.
Offline	The virtual drive is not available to the RAID controller.

2.4.19 Enclosure Management

Enclosure management is the intelligent monitoring of the disk subsystem by software and/or hardware. The disk subsystem can be part of the host computer or can reside in an external disk enclosure. Enclosure management helps you stay informed of events in the disk subsystem, such as a drive or power supply failure. Enclosure management increases the fault tolerance of the disk subsystem.

2.5 RAID Levels

The RAID controller supports RAID levels 0, 1, 5, 6, 10, 50, and 60. The supported RAID levels are summarized in the following section. In addition, it supports independent drives (configured as RAID 0). The following sections describe the RAID levels in detail.

2.5.1 Summary of RAID Levels

RAID 0 uses striping to provide high data throughput, especially for large files in an environment that does not require fault tolerance.

RAID 1 uses mirroring so that data written to one drive is simultaneously written to another drive. This is good for small databases or other applications that require small capacity but complete data redundancy.

RAID 5 uses disk striping and parity data across all drives (distributed parity) to provide high data throughput, especially for small random access.

RAID 6 uses distributed parity, with two independent parity blocks per stripe, and disk striping. A RAID 6 virtual drive can survive the loss of two drives without losing data. A RAID 6 drive group, which requires a minimum of three drives, is similar to a RAID 5 drive group. Blocks of data and parity information are written across all drives. The parity information is used to recover the data if one or two drives fail in the drive group.

RAID 10, a combination of RAID 0 and RAID 1, consists of striped data across mirrored spans. A RAID 10 drive group is a spanned drive group that creates a striped set from a series of mirrored drives. RAID 10 allows a maximum of eight spans. You must use an even number of

drives in each RAID virtual drive in the span. The RAID 1 virtual drives must have the same stripe size. RAID 10 provides high data throughput and complete data redundancy but uses a larger number of spans.

RAID 50, a combination of RAID 0 and RAID 5, uses distributed parity and disk striping. A RAID 50 drive group is a spanned drive group in which data is striped across multiple RAID 5 drive groups. RAID 50 works best with data that requires high reliability, high request rates, high data transfers, and medium-to-large capacity.

RAID 60, a combination of RAID 0 and RAID 6, uses distributed parity, with two independent parity blocks per stripe in each RAID set, and disk striping. A RAID 60 virtual drive can survive the loss of two drives in each of the RAID 6 sets without losing data. It works best with data that requires high reliability, high request rates, high data transfers, and medium-to-large capacity.

Note: Having virtual drives of different RAID levels, such as RAID 0 and RAID 5, in the same drive group is not allowed. For example, if an existing RAID 5 virtual drive is created out of partial space in an array, the next virtual drive in the array has to be R5 only.

2.5.2 Selecting a RAID Level

To ensure the best performance, you should select the optimal RAID level when you create a system drive. The optimal RAID level for your drive group depends on a number of factors:

- The number of drives in the drive group
- The capacity of the drives in the drive group
- The need for data redundancy
- The disk performance requirements

2.5.3 RAID 0

RAID 0 provides disk striping across all drives in the RAID drive group. RAID 0 does not provide any data redundancy, but, along with RAID 0, does offer the best performance of any RAID level. RAID 0 breaks up data into smaller segments, and then stripes the data segments across each drive in the drive group. The size of each data segment is determined by the stripe size. RAID 0 offers high bandwidth.


Note: RAID level 0 is not fault tolerant. If a drive in a RAID 0 drive group fails, the whole virtual drive (all drives associated with the virtual drive) will fail.

By breaking up a large file into smaller segments, the RAID controller can use both SAS drives and SATA drives to read or write the file faster. RAID 0 involves no parity calculations to complicate the write operation. This makes RAID 0 ideal for applications that require high bandwidth but do not require fault tolerance. [Table 2.5](#) provides an overview of RAID 0. The following figure provides a graphic example of a RAID 0 drive group.

Table 2.5 RAID 0 Overview

Uses	Provides high data throughput, especially for large files. Any environment that does not require fault tolerance.
Strong Points	Provides increased data throughput for large files. No capacity loss penalty for parity.
Weak Points	Does not provide fault tolerance or high bandwidth. All data lost if any drive fails.
Drives	1 to 32

Figure 2.5 RAID 0 Drive Group Example with Two Drives




2.5.4 RAID 1

In RAID 1, the RAID controller duplicates all data from one drive to a second drive in the drive group. RAID 1 supports an even number of drives from 2 to 32 in a single span. RAID 1 provides complete data redundancy, but at the cost of doubling the required data storage capacity. [Table 2.6](#) provides an overview of RAID 1. The following figure provides a graphic example of a RAID 1 drive group.

Table 2.6 RAID 1 Overview

Uses	Use RAID 1 for small databases or any other environment that requires fault tolerance but small capacity.
Strong Points	Provides complete data redundancy. RAID 1 is ideal for any application that requires fault tolerance and minimal capacity.
Weak Points	Requires twice as many drives. Performance is impaired during drive rebuilds.
Drives	2 - 32 (must be an even number of drives)

Figure 2.6 RAID 1 Drive Group



2.5.5 RAID 5

RAID 5 includes disk striping at the block level and parity. Parity is the data's property of being odd or even, and parity checking is used to detect errors in the data. In RAID 5, the parity information is written to all drives. RAID 5 is best suited for networks that perform a lot of small input/output (I/O) transactions simultaneously.


RAID 5 addresses the bottleneck issue for random I/O operations. Because each drive contains both data and parity, numerous writes can take place concurrently.

Table 2.7 provides an overview of RAID 5. The following figure provides a graphic example of a RAID 5 drive group.

Table 2.7 RAID 5 Overview

Uses	Provides high data throughput, especially for large files. Use RAID 5 for transaction processing applications because each drive can read and write independently. If a drive fails, the RAID controller uses the parity drive to recreate all missing information. Use also for office automation and online customer service that requires fault tolerance. Use for any application that has high read request rates but low write request rates.
Strong Points	Provides data redundancy, high read rates, and good performance in most environments. Provides redundancy with lowest loss of capacity.
Weak Points	Not well-suited to tasks requiring lot of writes. Suffers more impact if no cache is used (clustering). Drive performance will be reduced if a drive is being rebuilt. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes.
Drives	3 to 32

Figure 2.7 RAID 5 Drive Group with Six Drives



Note: Parity is distributed across all drives in the drive group.

2.5.6 RAID 6

RAID 6 is similar to RAID 5 (disk striping and parity), except that instead of one parity block per stripe, there are two. With two independent parity blocks, RAID 6 can survive the loss of two drives in a virtual drive without losing data. Provides a high level of data protection through the use of a second parity block in each stripe. Use RAID 6 for data that requires a very high level of protection from loss.

In the case of a failure of one drive or two drives in a virtual drive, the RAID controller uses the parity blocks to recreate all of the missing information. If two drives in a RAID 6 virtual drive fail, two drive rebuilds are required, one for each drive. These rebuilds do not occur at the same time. The controller rebuilds one failed drive, and then the other failed drive.


Table 2.6 provides a graphic example of a RAID 6 drive group.

Table 2.8 RAID 6 Overview

Uses	Use for office automation and online customer service that requires fault tolerance. Use for any application that has high read request rates but low write request rates.
Strong Points	Provides data redundancy, high read rates, and good performance in most environments. Can survive the loss of two drives or the loss of a drive while another drive is being rebuilt. Provides the highest level of protection against drive failures of all of the RAID levels. Read performance is similar to that of RAID 5.
Weak Points	Not well-suited to tasks requiring a lot of writes. A RAID 6 virtual drive has to generate two sets of parity data for each write operation, which results in a significant decrease in performance during writes. Drive performance is reduced during a drive rebuild. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes. RAID 6 costs more because of the extra capacity required by using two parity blocks per stripe.
Drives	3 to 32

The following figure shows a RAID 6 data layout. The second set of parity drives are denoted by *Q*. The *P* drives follow the RAID 5 parity scheme.

Figure 2.8 Example of Distributed Parity across Two Blocks in a Stripe (RAID 6)



Parity is distributed across all drives in the drive group.

2.5.7 RAID 10

RAID 10 is a combination of RAID 0 and RAID 1, and consists of stripes across mirrored drives. RAID 10 breaks up data into smaller blocks and then mirrors the blocks of data to each RAID 1 drive group. The first RAID 1 drive in each drive group then duplicates its data to the second drive. The size of each block is determined by the stripe size parameter, which is set during the creation of the RAID set. The RAID 1 virtual drives must have the same stripe size.

Spanning is used because one virtual drive is defined across more than one drive group. Virtual drives defined across multiple RAID 1 level drive groups are referred to as RAID level 10, (1+0). Data is striped across drive groups to increase performance by enabling access to multiple drive groups simultaneously.

Each spanned RAID 10 virtual drive can tolerate multiple drive failures, as long as each failure is in a separate drive group. If there are drive failures, less than total drive capacity is available.

Configure RAID 10 by spanning two contiguous RAID 1 virtual drives, up to the maximum number of supported devices for the controller. RAID 10 supports a maximum of eight spans, with a maximum of 32 drives per span. You must use an even number of drives in each RAID 10 virtual drive in the span.

Note: Other factors, such as the type of controller, can restrict the number of drives supported by RAID 10 virtual drives.


Table 2.9 provides an overview of RAID 10.

Table 2.9 RAID 10 Overview

Uses	Appropriate when used with data storage that needs 100 percent redundancy of mirrored drive groups and that also needs the enhanced I/O performance of RAID 0 (striped drive groups.) RAID 10 works well for medium-sized databases or any environment that requires a higher degree of fault tolerance and moderate to medium capacity.
Strong Points	Provides both high data transfer rates and complete data redundancy.
Weak Points	Requires twice as many drives as all other RAID levels except RAID 1.
Drives	4 - the maximum number of drives supported by the controller (with a maximum of eight spans)

In the following figure, virtual drive 0 is created by distributing data across four drive groups (drive groups 0 through 3).

Figure 2.9 RAID 10 Level Virtual Drive



2.5.8 RAID 50

RAID 50 provides the features of both RAID 0 and RAID 5. RAID 50 includes both parity and disk striping across multiple drive groups. RAID 50 is best implemented on two RAID 5 drive groups with data striped across both drive groups.

RAID 50 breaks up data into smaller blocks and then stripes the blocks of data to each RAID 5 disk set. RAID 5 breaks up data into smaller blocks, calculates parity by performing an exclusive-or on the blocks and

then writes the blocks of data and parity to each drive in the drive group. The size of each block is determined by the stripe size parameter, which is set during the creation of the RAID set.

Table 2.10 provides an overview of RAID 50. The following figure shows an example of a RAID 50 array.

Uses	Appropriate when used with data that requires high reliability, high request rates, high data transfer, and medium to large capacity.
Strong Points	Provides high data throughput, data redundancy, and very good performance.
Weak Points	Requires 2 to 8 times as many parity drives as RAID 5.
Drives	Eight spans of RAID 5 drive groups containing 3-32 drives each (limited by the maximum number of devices supported by the controller)

Figure 2.10 RAID 50 Level Virtual Drive

2.5.9 RAID 60

RAID 60 provides the features of both RAID 0 and RAID 6, and includes both parity and disk striping across multiple drive groups. RAID 6 supports two independent parity blocks per stripe. A RAID 60 virtual drive can survive the loss of two drives in each of the RAID 6 sets without losing data. RAID 60 is best implemented on two RAID 6 drive groups with data striped across both drive groups.

RAID 60 breaks up data into smaller blocks, and then stripes the blocks of data to each RAID 6 disk set. RAID 6 breaks up data into smaller blocks, calculates parity by performing an exclusive-or on the blocks and then writes the blocks of data and parity to each drive in the drive group. The size of each block is determined by the stripe size parameter, which is set during the creation of the RAID set.

RAID 60 can support up to 8 spans and tolerate up to 16 drive failures, though less than total drive capacity is available. Two drive failures can be tolerated in each RAID 6 level drive group.

Table 2.11 RAID 60 Overview


Uses	<p>Provides a high level of data protection through the use of a second parity block in each stripe. Use RAID 60 for data that requires a very high level of protection from loss.</p> <p>In the case of a failure of one drive or two drives in a RAID set in a virtual drive, the RAID controller uses the parity blocks to recreate all of the missing information. If two drives in a RAID 6 set in a RAID 60 virtual drive fail, two drive rebuilds are required, one for each drive. These rebuilds can occur at the same time.</p> <p>Use for office automation and online customer service that requires fault tolerance. Use for any application that has high read request rates but low write request rates.</p>
Strong Points	<p>Provides data redundancy, high read rates, and good performance in most environments. Each RAID 6 set can survive the loss of two drives or the loss of a drive while another drive is being rebuilt. Provides the highest level of protection against drive failures of all of the RAID levels. Read performance is similar to that of RAID 50, though random reads in RAID 60 might be slightly faster because data is spread across at least one more disk in each RAID 6 set.</p>

Table 2.11 RAID 60 Overview (Cont.)

Weak Points	Not well suited to tasks requiring lot of writes. A RAID 60 virtual drive has to generate two sets of parity data for each write operation, which results in a significant decrease in performance during writes. Drive performance is reduced during a drive rebuild. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes. RAID 6 costs more because of the extra capacity required by using two parity blocks per stripe.
Drives	A minimum of 8

The following figure shows a RAID 6 data layout. The second set of parity drives are denoted by *Q*. The *P* drives follow the RAID 5 parity scheme.

Figure 2.11 RAID 60 Level Virtual Drive



Parity is distributed across all drives in the drive group.

2.6 RAID Configuration Strategies

The most important factors in RAID drive group configuration are:

- Virtual drive availability (fault tolerance)
- Virtual drive performance
- Virtual drive capacity

You cannot configure a virtual drive that optimizes all three factors, but it is easy to choose a virtual drive configuration that maximizes one factor at the expense of another factor. For example, RAID 1 (mirroring) provides excellent fault tolerance, but requires a redundant drive. The

following subsections describe how to use the RAID levels to maximize virtual drive availability (fault tolerance), virtual drive performance, and virtual drive capacity.

2.6.1 Maximizing Fault Tolerance

Fault tolerance is achieved through the ability to perform automatic and transparent rebuilds using hot spare drives and hot swaps. A hot spare drive is an unused online available drive that the RAID controller instantly plugs into the system when an active drive fails. After the hot spare is automatically moved into the RAID drive group, the failed drive is automatically rebuilt on the spare drive. The RAID drive group continues to handle requests while the rebuild occurs.

A hot swap is the manual substitution of a replacement unit in a disk subsystem for a defective one, where the substitution can be performed while the subsystem is running hot swap drives. Auto-Rebuild in the WebBIOS Configuration Utility allows a failed drive to be replaced and automatically rebuilt by “hot-swapping” the drive in the same drive bay. The RAID drive group continues to handle requests while the rebuild occurs, providing a high degree of fault tolerance and zero downtime.

[Table 2.12](#) describes the fault tolerance for each RAID level.

Table 2.12 RAID Levels and Fault Tolerance

RAID Level	Fault Tolerance
0	Does not provide fault tolerance. All data lost is if any drive fails. Disk striping writes data across multiple drives instead of just one drive. It involves partitioning each drive storage space into stripes that can vary in size. RAID 0 is ideal for applications that require high bandwidth but do not require fault tolerance.
1	Provides complete data redundancy. If one drive fails, the contents of the other drive in the drive group can be used to run the system and reconstruct the failed drive. The primary advantage of disk mirroring is that it provides 100 percent data redundancy. Since the contents of the drive are completely written to a second drive, no data is lost if one of the drives fails. Both drives contain the same data at all times. RAID 1 is ideal for any application that requires fault tolerance and minimal capacity.
5	Combines distributed parity with disk striping. Parity provides redundancy for one drive failure without duplicating the contents of entire drives. If a drive fails, the RAID controller uses the parity data to reconstruct all missing information. In RAID 5, this method is applied to entire drives or stripes across all drives in a drive group. Using distributed parity, RAID 5 offers fault tolerance with limited overhead.

Table 2.12 RAID Levels and Fault Tolerance (Cont.)

RAID Level	Fault Tolerance
6	Combines distributed parity with disk striping. RAID 6 can sustain two drive failures and still maintain data integrity. Parity provides redundancy for two drive failures without duplicating the contents of entire drives. If a drive fails, the RAID controller uses the parity data to reconstruct all missing information. In RAID 6, this method is applied to entire drives or stripes across all of the drives in a drive group. Using distributed parity, RAID 6 offers fault tolerance with limited overhead.
10	Provides complete data redundancy using striping across spanned RAID 1 drive groups. RAID 10 works well for any environment that requires the 100 percent redundancy offered by mirrored drive groups. RAID 10 can sustain a drive failure in each mirrored drive group and maintain drive integrity.
50	Provides data redundancy using distributed parity across spanned RAID 5 drive groups. RAID 50 includes both parity and disk striping across multiple drives. If a drive fails, the RAID controller uses the parity data to recreate all missing information. RAID 50 can sustain one drive failure per RAID 5 drive group and still maintain data integrity.
60	Provides data redundancy using distributed parity across spanned RAID 6 drive groups. RAID 60 can sustain two drive failures per RAID 6 drive group and still maintain data integrity. It provides the highest level of protection against drive failures of all of the RAID levels. RAID 60 includes both parity and disk striping across multiple drives. If a drive fails, the RAID controller uses the parity data to recreate all missing information.

2.6.2 Maximizing Performance

A RAID disk subsystem improves I/O performance. The RAID drive group appears to the host computer as a single storage unit or as multiple virtual units. I/O is faster because drives can be accessed simultaneously. [Table 2.13](#) describes the performance for each RAID level.

Table 2.13 RAID Levels and Performance

RAID Level	Performance
0	RAID 0 (striping) offers excellent performance. RAID 0 breaks up data into smaller blocks, then writes a block to each drive in the drive group. Disk striping writes data across multiple drives instead of just one drive. It involves partitioning each drive storage space into stripes that can vary in size from 8 KB to 128 KB. These stripes are interleaved in a repeated sequential manner. Disk striping enhances performance because multiple drives are accessed simultaneously.
1	With RAID 1 (mirroring), each drive in the system must be duplicated, which requires more time and resources than striping. Performance is impaired during drive rebuilds.

Table 2.13 RAID Levels and Performance (Cont.)

RAID Level	Performance
5	<p>RAID 5 provides high data throughput, especially for large files. Use this RAID level for any application that requires high read request rates, but low write request rates, such as transaction processing applications, because each drive can read and write independently. Since each drive contains both data and parity, numerous writes can take place concurrently. In addition, robust caching algorithms and hardware based exclusive-or assist make RAID 5 performance exceptional in many different environments.</p> <p>Parity generation can slow the write process, making write performance significantly lower for RAID 5 than for RAID 0 or RAID 1. Drive performance is reduced when a drive is being rebuilt. Clustering can also reduce drive performance. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes.</p>
6	<p>RAID 6 works best when used with data that requires high reliability, high request rates, and high data transfer. It provides high data throughput, data redundancy, and very good performance. However, RAID 6 is not well suited to tasks requiring a lot of writes. A RAID 6 virtual drive has to generate two sets of parity data for each write operation, which results in a significant decrease in performance during writes. Drive performance is reduced during a drive rebuild. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes.</p>
10	<p>RAID 10 works best for data storage that need the enhanced I/O performance of RAID 0 (striped drive groups), which provides high data transfer rates. Spanning increases the capacity of the virtual drive and improves performance by doubling the number of spindles. The system performance improves as the number of spans increases. (The maximum number of spans is eight.) As the storage space in the spans is filled, the system stripes data over fewer and fewer spans and RAID performance degrades to that of a RAID 1 or RAID 5 drive group.</p>
50	<p>RAID 50 works best when used with data that requires high reliability, high request rates, and high data transfer. It provides high data throughput, data redundancy, and very good performance. Spanning increases the capacity of the virtual drive and improves performance by doubling the number of spindles. The system performance improves as the number of spans increases. (The maximum number of spans is eight.) As the storage space in the spans is filled, the system stripes data over fewer and fewer spans and RAID performance degrades to that of a RAID 1 or RAID 5 drive group.</p>
60	<p>RAID 60 works best when used with data that requires high reliability, high request rates, and high data transfer. It provides high data throughput, data redundancy, and very good performance. Spanning increases the capacity of the virtual drive and improves performance by doubling the number of spindles. The system performance improves as the number of spans increases. (The maximum number of spans is eight.) As the storage space in the spans is filled, the system stripes data over fewer and fewer spans and RAID performance degrades to that of a RAID 1 or RAID 6 drive group.</p> <p>RAID 60 is not well suited to tasks requiring a lot of writes. A RAID 60 virtual drive has to generate two sets of parity data for each write operation, which results in a significant decrease in performance during writes. Drive performance is reduced during a drive rebuild. Environments with few processes do not perform as well because the RAID overhead is not offset by the performance gains in handling simultaneous processes.</p>

2.6.3 Maximizing Storage Capacity

Storage capacity is an important factor when selecting a RAID level. There are several variables to consider. Striping alone (RAID 0) requires less storage space than mirrored data (RAID 1) or distributed parity (RAID 5 and RAID 6). RAID 5, which provides redundancy for one drive failure without duplicating the contents of entire drives, requires less space than RAID 1. [Table 2.14](#) explains the effects of the RAID levels on storage capacity.

Table 2.14 RAID Levels and Capacity

RAID Level	Capacity
0	RAID 0 (striping) involves partitioning each drive storage space into stripes that can vary in size. The combined storage space is composed of stripes from each drive. RAID 0 provides maximum storage capacity for a given set of drives.
1	With RAID 1 (mirroring), data written to one drive is simultaneously written to another drive, which doubles the required data storage capacity. This is expensive because each drive in the system must be duplicated.
5	RAID 5 provides redundancy for one drive failure without duplicating the contents of entire drives. RAID 5 breaks up data into smaller blocks, calculates parity by performing an exclusive-or on the blocks, then writes the blocks of data and parity to each drive in the drive group. The size of each block is determined by the stripe size parameter, which is set during the creation of the RAID set.
6	RAID 6 provides redundancy for two drive failures without duplicating the contents of entire drives. However, it requires extra capacity because it uses two parity blocks per stripe. This makes RAID 60 more expensive to implement.
10	RAID 10 requires twice as many drives as all other RAID levels except RAID 1. RAID 10 works well for medium-sized databases or any environment that requires a higher degree of fault tolerance and moderate to medium capacity. Disk spanning allows multiple drives to function like one big drive. Spanning overcomes lack of disk space and simplifies storage management by combining existing resources or adding relatively inexpensive resources.
50	RAID 50 requires two to four times as many parity drives as RAID 5. This RAID level works best when used with data that requires medium to large capacity.
60	RAID 60 provides redundancy for two drive failures in each RAID set without duplicating the contents of entire drives. However, it requires extra capacity because a RAID 60 virtual drive has to generate two sets of parity data for each write operation. This makes RAID 60 more expensive to implement.

2.7 RAID Availability

2.7.1 RAID Availability Concept

Data availability without downtime is essential for many types of data processing and storage systems. Businesses want to avoid the financial costs and customer frustration associated with failed servers. RAID helps you maintain data availability and avoid downtime for the servers that provide that data. RAID offers several features, such as spare drives and rebuilds, that you can use to fix any drive problems, while keeping the servers running and data available. The following subsections describe these features.

2.7.1.1 Spare Drives

You can use spare drives to replace failed or defective drives in a drive group. A replacement drive must be at least as large as the drive it replaces. Spare drives include hot swaps, hot spares, and cold swaps.

A hot swap is the manual substitution of a replacement unit in a disk subsystem for a defective one, where the substitution can be performed while the subsystem is running (performing its normal functions). The backplane and enclosure must support hot swap in order for the functionality to work.

Hot spare drives are drives that power up along with the RAID drives and operate in a standby state. If a drive used in a RAID virtual drive fails, a hot spare automatically takes its place and the data on the failed drive is rebuilt on the hot spare. Hot spares can be used for RAID levels 1, 5, 6, 10, 50, and 60.

Note: If a rebuild to a hot spare fails for any reason, the hot spare drive will be marked as "failed." If the source drive fails, both the source drive and the hot spare drive will be marked as "failed."

A cold swap requires that you power down the system before replacing a defective drive in a disk subsystem.

2.7.1.2 Rebuilding

If a drive fails in a drive group that is configured as a RAID 1, 5, 6, 10, 50, or 60 virtual drive, you can recover the lost data by rebuilding the drive. If you have configured hot spares, the RAID controller automatically tries to use them to rebuild failed drives. Manual rebuild is necessary if no hot spares with enough capacity to rebuild the failed drives are available. You must insert a drive with enough storage into the subsystem before rebuilding the failed drive.

2.8 Configuration Planning

Factors to consider when planning a configuration are the number of drives the RAID controller can support, the purpose of the drive group, and the availability of spare drives.

Each type of data stored in the disk subsystem has a different frequency of read and write activity. If you know the data access requirements, you can more successfully determine a strategy for optimizing the disk subsystem capacity, availability, and performance.

Servers that support video on demand typically read the data often, but write data infrequently. Both the read and write operations tend to be long. Data stored on a general-purpose file server involves relatively short read and write operations with relatively small files.

2.8.1 Number of Drives

Your configuration planning for the SAS RAID controller depends in part on the number of drives that you want to use in a RAID drive group. The number of drives in a drive group determines the RAID levels that can be supported. Only one RAID level can be assigned to each virtual drive.

2.8.2 Drive Group Purpose

Important factors to consider when creating RAID drive groups include availability, performance, and capacity. Define the major purpose of the drive group by answering questions related to these factors, such as the following, which are followed by suggested RAID levels for each situation:

- Will this drive group increase the system storage capacity for general-purpose file and print servers? Use RAID 5, 6, 10, 50, or 60.
- Does this drive group support any software system that must be available 24 hours per day? Use RAID 1, 5, 6, 10, 50, or 60.
- Will the information stored in this drive group contain large audio or video files that must be available on demand? Use RAID 0.
- Will this drive group contain data from an imaging system? Use RAID 0, or 10.

Fill out [Table 2.15](#) to help you plan the drive group configuration. Rank the requirements for your drive group, such as storage space and data redundancy, in order of importance, and then review the suggested RAID levels.

Table 2.15 Factors to Consider for Drive Group Configuration

Requirement	Rank	Suggested RAID Level(s)
Storage space		RAID 0, RAID 5
Data redundancy		RAID 5, RAID 6, RAID 10, RAID 50, RAID 60
Drive performance and throughput		RAID 0, RAID 10
Hot spares (extra drives required)		RAID 1, RAID 5, RAID 5, RAID 10, RAID 50, RAID 60

Chapter 3

Self-Encrypting Disk Feature

This chapter describes the self-encrypting disk (SED) feature and consists of the following sections:

- [Section 3.1, “Overview”](#)
- [Section 3.2, “Purpose”](#)
- [Section 3.3, “Terminology”](#)
- [Section 3.4, “Workflow”](#)

3.1 Overview

The SED feature offers the ability to encrypt data on disks and use disk-based key management to provide data security. With the SED feature, data is encrypted by the drives. You can designate which data to encrypt at the individual virtual drive (VD) level.

This solution provides data protection in the event of theft or loss of physical drives. With self-encrypting disks, if you remove a drive from its storage system or the server it is housed in, the data on that drive is encrypted and useless to anyone who attempts to access without the appropriate security authorization.

Any encryption solution requires management of the encryption keys. The security feature provides a way to manage these keys. You can change the encryption key for all ServeRAID controllers that are connected to SED drives. All SED drives, whether locked or unlocked, always have an encryption key. This key is set by the drive and is always active. When the drive is unlocked, the data to host from the drive (on reads) and from the host to the drive cache (on writes) is always provided. However, when resting on the drive platters, the data is always encrypted by the drive.

You might not want to lock your drives because you have to manage a password if they are locked. Even if you do not lock the drives, there is still a benefit to using SED drives.

The WebBIOS Configuration Utility ([Section Figure 4.5, “WebBIOS Virtual Drive Definition Screen”](#)) and MegaRAID Storage Manager ([Chapter 8, “Monitoring System Events and Storage Devices,”](#)) offer procedures that you can use to manage the security settings for the drives.

3.2 Purpose

Security is a growing market concern and requirement. ServeRAID customers are looking for a comprehensive storage encryption solution to protect data. You can use the SED feature to help protect your data.

3.3 Terminology

[Table 3.1](#) describes the terminology related to the SED feature.

Table 3.1 SED Terminology

Option	Description
Authenticated Mode	The RAID configuration is keyed to a user passphrase. The passphrase must be provided on system boot to authenticate the user and facilitate unlocking the configuration for user access to the encrypted data.
Blob	A blob is created by encrypting a key(s) using another key. There are two types of blob in the system – encryption key blob and security key blob.
Key backup	You need to provide the controller with a lock key if the controller is replaced or if you choose to migrate secure virtual drives. To do this, you must back up the security key.
Passphrase	An optional authenticated mode is supported in which you must provide a passphrase on each boot to make sure the system boots only if the user is authenticated. Firmware uses the user passphrase to encrypt the security key in the security key blob stored on the controller.

Table 3.1 SED Terminology (Cont.)

Option	Description
Re-provisioning	Re-provisioning disables the security system of a device. For a controller, it involves destroying the security key. For self-encrypting drives, when the drive lock key is deleted, the drive is unlocked and any user data on the drive is securely deleted. This does not apply to controller-encrypted drives, because deleting the virtual drive destroys the encryption keys and causes a secure erase. See Section 3.5, “Instant Secure Erase” for information about the instant secure erase feature.
Security Key	A key based on a user-provided string. The controller uses the security key to lock and unlock access to the secure user data. This key is encrypted into the security key blob and stored on the controller. If the security key is unavailable, user data is irretrievably lost. You must take all precautions to never lose the security key.
Un-Authenticated Mode	This mode allows controller to boot and unlock access to user configuration without user intervention. In this mode, the security key is encrypted into a security key blob, stored on the controller, but instead of a user passphrase, an internal key specific to the controller is used to create the security key blob.
Volume Encryption Keys (VEK)	The controller uses the Volume Encryption Keys to encrypt data when a controller-encrypted virtual drive is created. These keys are not available to the user. The firmware (FW) uses a unique 512-bit key for each virtual drive. The VEK for the VDs are stored on the physical drives in a VEK blob.

3.4 Workflow

3.4.1 Enable Security

You can enable security on the controller. After you enable security, you have the option to create secure virtual drives using a security key.

There are three procedures you can perform to create secure virtual drives using a security key:

- Create the security key identifier
- Create the security key
- Create a pass phrase (optional)

See [Section 4.5, “Selecting Self-Encrypting Disk Security Options”](#) for the procedures used to enable security in WebBIOS or [Section 8.1, “Monitoring System Events,”](#) for the procedures used to enable security in MegaRAID Storage Manager.

3.4.1.1 Create the Security Key Identifier

The security key identifier appears whenever you enter the security key. If you have multiple security keys, the identifier helps you determine which security key to enter. The controller provides a default identifier for you. You can use the default or enter your own identifier.

3.4.1.2 Create the Security Key

You need to enter the security key to perform certain operations. You can choose a strong security key that the controller suggests.

Attention: If you forget the security key, you will lose access to your data.

3.4.1.3 Create a Passphrase (Optional)

The pass phrase provides additional security. The pass phrase should be different from the security key. If you choose this option, you must enter it whenever you boot your server.

Attention: If you forget the pass phrase, you will lose access to your data.

When you use the specified security key identifier, security key, and pass phrase, security will be enabled on the controller.

3.4.2 Change Security

You can change the security settings on the controller, and you have the option to change the security key identifier, security key, and pass phrase. If you have previously removed any secured drives, you still need to supply the old security key to import them.

There are three procedures you can perform to change the security settings on the controller:

- Change the security key identifier
- Change the security key
- Change a pass phrase

See [Section 4.5, “Selecting Self-Encrypting Disk Security Options”](#) for the procedures used to change security options in WebBIOS or [Section 7.2, “Selecting Self-Encrypting Disk Security Options”](#) for the procedures used to change security options in MegaRAID Storage Manager.

3.4.2.1 Change the Security Key Identifier

You have the option to edit the security key identifier. If you plan to change the security key, it is highly recommended that you change the security key identifier. Otherwise, you will not be able to differentiate between the security keys.

You can select whether you want to keep the current security key identifier or enter a new one. To change the security key identifier, enter a new security key identifier.

3.4.2.2 Change the Security Key

You can choose to keep the current security key or enter a new one. To change the security key, you can either enter the new security key or accept the security key that the controller suggests.

3.4.2.3 Add or Change the Pass Phrase

You have the option to add a pass phrase or change the existing one. To change the pass phrase, enter the new pass phrase. To keep the existing pass phrase, enter the current pass phrase. If you choose this option, you must enter the pass phrase whenever you boot your server.

This procedure updates the existing configuration on the controller to use the new security settings.

3.4.3 Create Secure Virtual Drives

You can create a secure virtual drive and set their parameters as desired. To create a secure virtual drive, select a configuration method. You can select either simple configuration or advanced configuration.

3.4.3.1 Simple Configuration

If you select simple configuration, select the redundancy type and the drive security method to use for the drive group.

See [Section 7.1.2, “Creating a Virtual Drive Using Simple Configuration”](#) for the procedures used to select the redundancy type and drive security method for a configuration.

3.4.3.2 Advanced Configuration

If you select advanced configuration, select the drive security method, and add the drives to the drive group.

See [Section 7.1.3, “Creating a Virtual Drive Using Advanced Configuration”](#) for the procedures used to import a foreign configuration.

After the drive group is secured, you cannot remove the security without deleting the virtual drives.

3.4.4 Import a Foreign Configuration

After you create a security key, you can run a scan for a foreign configuration and import a locked configuration. (You can import unsecured or unlocked configurations when security is disabled.) A foreign configuration is a RAID configuration that already exists on a replacement set of drives that you install in a computer system. WebBIOS Configuration Utility and MSM allows you to import the existing configuration to the RAID controller or clear the configuration so you can create a new one.

See [Section 4.5.4, “Importing Foreign Configurations”](#) for the procedure used to import a foreign configuration in WebBIOS or [Section 7.2.4, “Importing or Clearing a Foreign Configuration”](#) for the procedure in MegaRAID Storage Manager.

To import a foreign configuration, you must first enable security to allow importation of locked foreign drives. If the drives are locked and the controller security is disabled, you cannot import the foreign drives. Only unlocked drives can be imported when security is disabled.

After you enable the security, you can import the locked drives. To import the locked drives, you must provide the security key used to secure them. Verify whether any drives are left to import as the locked drives can use different security keys. If there are any drives left, repeat the import process for the remaining drives. After all of the drives are imported, there is no configuration to import.

3.5 Instant Secure Erase

The Instant Secure Erase feature offers a way to erase data that you can use with SED drives. After the initial investment into a SED drive, there is no additional cost in dollars or time to erase data using the Instant Secure Erase feature.

If you are concerned about data theft or other security issues, you might already invest in drive disposal costs, and there are benefits to using SED over other technologies that exists today, both in terms of the security provided and time saved.

If the encryption key on the drive changes, the drive cannot decrypt the data on the platters, effectively erasing the data on the drives. The National Institute of Standards and Technology (<http://www.nist.gov>) values this type of data erasure above secure erase and below physical destruction of the device.

There are four major reasons for using instant secure erase.

If there is a need to repurpose the hard drive for a different application – You might need to move the drive to another server to expand storage elsewhere, but the drive is in use. The data on the drive might contain sensitive data including customer information that, if lost or divulged, could cause an embarrassing disclosure of a security hole. You can use the instant secure erase feature to effectively erase the data so the drive can be moved to another server or area without concern that old data could be found.

If there is a need to replace drives – If the amount of data has outgrown the storage system, and there is no room to expand capacity by adding drives, you might choose to purchase upgrade drives. If the older drives support SED, you can erase the data instantly so the new drives can be used.

If there is a need to return a drive for warranty activity – If the drive is beginning to show SMART predictive failure alerts, you might want to return the drive for replacement. If so, the drive needs to be effectively erased if there is sensitive data. Occasionally a drive is in such bad condition that standard erasure applications do not work. If the

drive still allows any access, it might be possible to destroy the encryption key.

Chapter 4

WebBIOS Configuration Utility

This chapter describes the WebBIOS Configuration Utility (CU) and consists of the following sections:

- [Section 4.1, “Overview”](#)
- [Section 4.2, “Starting the WebBIOS CU”](#)
- [Section 4.3, “WebBIOS Configuration Utility Main Screen Options”](#)
- [Section 4.4, “Creating a Storage Configuration”](#)
- [Section 4.5, “Selecting Self-Encrypting Disk Security Options”](#)
- [Section 4.6, “Viewing and Changing Device Properties”](#)
- [Section 4.7, “Viewing and Expanding a Virtual Drive”](#)
- [Section 4.8, “Recovering/Clearing Punctured Block Entries”](#)
- [Section 4.9, “Suspending and Resuming Virtual Drive Operations”](#)
- [Section 4.10, “Non-SED Secure Erase”](#)
- [Section 4.11, “Viewing System Event Information”](#)
- [Section 4.12, “Managing Configurations”](#)

4.1 Overview

The WebBIOS CU enables you to create and manage RAID configurations on IBM ServeRAID SAS/SATA controllers. Unlike the MegaRAID Storage Manager™ software, the WebBIOS CU resides in the controller BIOS and operates independently of the operating system.

You can use the WebBIOS CU to do the following tasks:

- Create drive groups and virtual drives for storage configurations

- Display controller, virtual drive, drive, and battery backup unit (BBU) properties, and change parameters
- Delete virtual drives
- Migrate a storage configuration to a different RAID level
- Detect configuration mismatches
- Import a foreign configuration
- Scan devices connected to the controller
- Initialize virtual drives
- Check configurations for data consistency

The WebBIOS CU provides a configuration wizard to guide you through the configuration of virtual drives and drive groups.

4.2 Starting the WebBIOS CU

Follow these steps to start the WebBIOS CU and access the main screen.

1. When the host computer is booting, hold down the <Ctrl> key and press the <H> key when the following text appears on the screen:

Copyright© LSI Corporation
Press <Ctrl><H> for WebBIOS

The Controller Selection screen appears.

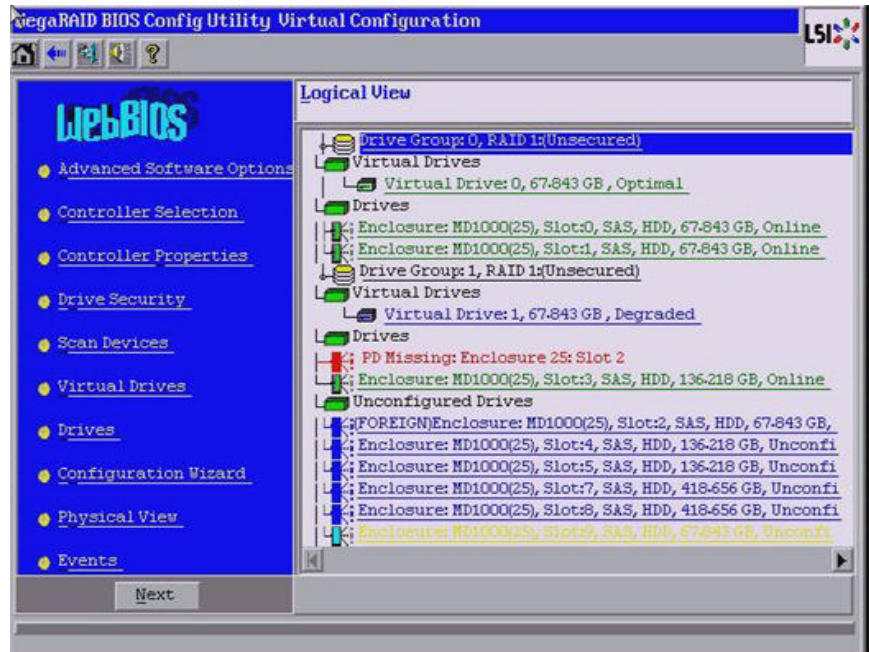
2. If the system has multiple SAS controllers, select a controller.
3. Click **Start** to continue.

The main WebBIOS CU screen appears.

4.3 WebBIOS Configuration Utility Main Screen Options

The following figure shows the screen that appears when you start the WebBIOS CU and select a controller.

Figure 4.1 WebBIOS CU Main Screen



The right frame of the screen shows the virtual drives configured on the controller, and the drives that are connected to the controller. In addition, the screen identifies drives that are foreign or missing.

Note: In the list of virtual drives, the drive nodes are sorted based on the order in which you added the drives to the drive group, rather than the physical slot order that displays in the physical trees.

Note: The minimum screen resolution for WebBIOS is 640x480.

To toggle between the physical view and logical view of the storage devices connected to the controller, click **Physical View** or **Logical View**

in the menu on the left. When the physical view screen appears, it shows the drive groups that are configured on this controller.






Note: Unconfigured Bad drives are only displayed in the Physical View.

For drives in an enclosure, the screen shows the following drive information:

- Enclosure
- Slot
- Interface type (such as SAS or SATA)
- Drive type (HDD or SSD)
- Drive size
- Drive status (such as **Online** or **Unconfigured Good**)

The toolbar at the top of the WebBIOS CU has the following buttons, as listed in [Table 4.1](#).

Table 4.1 WebBIOS CU Toolbar Icons

Icon	Description
	Click this icon to return to the main screen from any other WebBIOS CU screen.
	Click this icon to return to the previous screen that you were viewing.
	Click this icon to exit the WebBIOS CU program.
	Click this icon to turn off the sound on the onboard controller alarm.
	Click this icon to display information about the WebBIOS CU version, browser version, and HTML interface engine.

Here is a description of the options listed on the left of the main WebBIOS CU screen:

- **Controller Selection: (Alt+c)** Select this option to view the Controller Selection screen, where you can select a different SAS controller. You can then view information about the controller and the devices connected to it, or create a new configuration on the controller.
- **Controller Properties: (Alt+p)** Select this option to view the properties of the currently selected SAS controller. For more information, see [Section 4.6.1, “Viewing and Changing Controller Properties.”](#)
- **Drive Security: (Alt+r)** Select this option to encrypt data on the drives and use disk-based key management for the data security solution. This solution protects your data in case of theft or loss of physical drives. For more information, see [Section 4.5, “Selecting Self-Encrypting Disk Security Options.”](#)
- **Scan Devices: (Alt+s)** Select this option to have the WebBIOS CU re-scan the physical and virtual drives for any changes in the drive status or the physical configuration. The WebBIOS CU displays the results of the scan in the physical and virtual drive descriptions.
- **Virtual Drives: (Alt+v)** Select this option to view the Virtual Drives screen, where you can change and view virtual drive properties, delete virtual drives, initialize drives, and perform other tasks. For more information, see [Section 4.6.2, “Viewing and Changing Virtual Drive Properties.”](#)
- **Drives: (Alt+d)** Select this option to view the Drives screen, where you can view drive properties, create hot spares, and perform other tasks. For more information, see [Section 4.6.3, “Viewing Drive Properties.”](#)
- **Configuration Wizard: (Alt+o)** Select this option to start the Configuration Wizard and create a new storage configuration, clear a configuration, or add a configuration. For more information, see [Section 4.4, “Creating a Storage Configuration.”](#)
- **Physical View/Logical View: (Alt+l)** Select this option to toggle between the Physical View and Logical View screens.

- **Events: (Alt+e)** Select this option to view system events in the Event Information screen. For more information, see [Section 4.11, “Viewing System Event Information.”](#)
- **Exit: (Alt+x)** Select this option to exit the WebBIOS CU and continue with system boot.

4.4 Creating a Storage Configuration

This section explains how to use the WebBIOS CU Configuration Wizard to configure RAID drive groups and virtual drives. The following subsections explain how to use the Configuration Wizard to create storage configurations:

- [Section 4.4.1, “Selecting the Configuration with the Configuration Wizard”](#)
- [Section 4.4.2, “Using Automatic Configuration”](#)
- [Section 4.4.3, “Using Manual Configuration”](#)


4.4.1 Selecting the Configuration with the Configuration Wizard

Follow these steps to start the Configuration Wizard, and select a configuration option and mode:

1. Click **Configuration Wizard** on the WebBIOS main screen.

The first Configuration Wizard screen appears, as shown in the following figure.

Figure 4.2 WebBIOS Configuration Wizard Screen



2. Select a configuration option.

Attention: If you choose the first or second option, all existing data in the configuration will be deleted. Make a backup of any data that you want to keep before you choose an option.

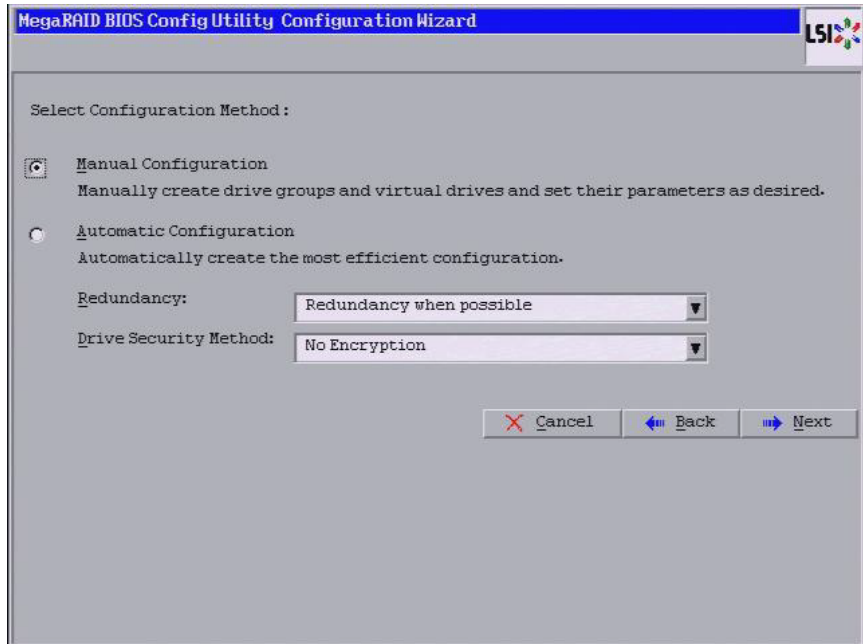
- **Clear Configuration:** Clears the existing configuration.
- **New Configuration:** Clears the existing configuration and lets you create a new configuration.
- **Add Configuration:** Retains the existing storage configuration and adds new drives to it (this does not cause any data loss).

3. Click **Next**.

A dialog box warns that you will lose data if you select Clear Configuration or New Configuration.

The WebBIOS Configuration Method screen appears, as shown in the following figure.

Figure 4.3 WebBIOS Configuration Method Screen



4. On this screen, select a configuration mode:
 - **Manual Configuration:** Allows you to control all attributes of the new storage configuration as you create drive groups and virtual drives, and set their parameters.
 - **Automatic Configuration:** Automatically creates an optimal RAID configuration.
5. If you select Automatic Configuration, you can choose whether to create a redundant RAID drive group or a non-redundant RAID 0 drive group. Select one of the following options in the Redundancy field:

- **Redundancy when possible**
- **No redundancy**

If you select **Automatic Configuration**, you can choose whether to use a drive security method. Select one of the following options in the **Drive Security Method** drop down list:

- **No Encryption**
- **Drive Encryption**

6. Click **Next** to continue.

If you select the Automatic Configuration option, continue with [Section 4.4.2, “Using Automatic Configuration.”](#) If you select Manual Configuration, continue with [Section 4.4.3, “Using Manual Configuration.”](#)

4.4.2 Using Automatic Configuration

Follow these instructions to create a configuration with automatic configuration, either with or without redundancy:

1. When WebBIOS displays the proposed new configuration, review the information on the screen, and click **Accept** to accept it. (Or click **Back** to go back and change the configuration.)
 - **RAID 0:** If you select **Automatic Configuration** and **No Redundancy**, WebBIOS creates a RAID 0 configuration.
 - **RAID 1:** If you select **Automatic Configuration** and **Redundancy when possible**, and only two drives are available, WebBIOS creates a RAID 1 configuration.
 - **RAID 5:** If you select **Automatic Configuration** and **Redundancy when possible**, and three or more drives are available, WebBIOS creates a RAID 5 configuration.
 - **RAID 6:** If you select **Automatic Configuration** and **Redundancy when possible**, and the RAID 6 option is enabled, and three or more drives are available, WebBIOS creates a RAID 6 configuration.
2. Click **Yes** when you are prompted to save the configuration.
3. Click **Yes** when you are prompted to initialize the new virtual drive(s).

The WebBIOS CU begins a background initialization of the virtual drives.

New RAID 5 virtual drives and new RAID 6 virtual drives require a minimum number of drives for a background initialization to start. If there are fewer drives, the background initialization will not start. The following number of drives is required:

- New RAID 5 virtual drives must have at least five drives for a background initialization to start.
- New RAID 6 virtual drives must have at least seven drives for a background initialization to start.

4.4.3 Using Manual Configuration

The following subsections contain the procedures for creating RAID drive groups for RAID levels 0, 1, 5, 6, 10, 50, and 60:

- [Section 4.4.3.1, “Using Manual Configuration: RAID 0”](#)
- [Section 4.4.3.2, “Using Manual Configuration: RAID 1”](#)
- [Section 4.4.3.3, “Using Manual Configuration: RAID 5”](#)
- [Section 4.4.3.4, “Using Manual Configuration: RAID 6”](#)
- [Section 4.4.3.5, “Using Manual Configuration: RAID 10”](#)
- [Section 4.4.3.6, “Using Manual Configuration: RAID 50”](#)
- [Section 4.4.3.7, “Using Manual Configuration: RAID 60”](#)

4.4.3.1 Using Manual Configuration: RAID 0

RAID 0 provides drive striping across all drives in the RAID drive group. RAID 0 does not provide any data redundancy but does offer excellent performance. RAID 0 is ideal for applications that require high bandwidth but do not require fault tolerance. RAID 0 also denotes an independent or single drive.

Note: RAID level 0 is not fault-tolerant. If a drive in a RAID 0 drive group fails, the whole virtual drive (all drives associated with the virtual drive) fails.


When you select **Manual Configuration** and click **Next**, the drive group Definition screen appears. You use this screen to select drives to create drive groups.

1. Hold <Ctrl> while selecting two or more ready drives in the Drives panel on the left until you have selected all desired drives for the drive group.
2. Click **Add To Array** to move the drives to a proposed drive group configuration in the Disk Groups panel on the right, as shown in the following figure.

If you need to undo the changes, click the **Reclaim** button.

3. Choose whether to use power save mode.
4. Choose whether to use drive encryption.

Figure 4.4 WebBIOS Disk Group Definition Screen



5. After you finish selecting drives for the drive group, click **Accept DG**.
6. Click **Next**.

The Virtual Drive Definition screen appears, as shown in the following figure. This screen lists the possible RAID levels for the drive group. Use this screen to select the RAID level, stripe size, read policy, and other attributes for the new virtual drives.

Figure 4.5 WebBIOS Virtual Drive Definition Screen

The screenshot shows the 'MegaRAID BIOS Config Utility Config Wizard - Virtual Drive Definition' window. On the left, there are configuration options for a virtual drive:

- RAID Level:** RAID 0
- Strip Size:** 64 KB
- Access Policy:** RW
- Read Policy:** Ahead
- Write Policy:** Write Back with BBU
- IO Policy:** Direct
- Drive Cache:** Disable
- Disable BGI:** No
- Select Size:** GB

On the right, under 'Virtual Drives', there is a list of drives: R0:836.625 GB, R5:557.750 GB, and R6:278.875 GB. Below this list is a green text label: 'Next LD, Possible RAID Levels'. At the bottom, there are buttons for 'Accept', 'Reclaim', 'Cancel', 'Back', and 'Next'.

7. Change the virtual drive options from the defaults listed on the screen as needed.

Here are brief explanations of the virtual drive options:

- **RAID Level:** The drop-down menu lists the possible RAID levels for the virtual drive. Select RAID 0.
- **Strip Size:** The strip size specifies the length of the data segments that the RAID controller writes across multiple drives, not including parity drives. For example, consider a stripe that contains 64 KB of drive space and has 16 KB of data residing on each drive in the stripe. In this case, the stripe size is 64 KB and the strip size is 16 KB. You can set the strip size to 8, 16, 32, 64, 128, 256, 512, and 1024 Kbytes. A larger strip size produces higher read performance. If your computer regularly performs random read requests, choose a smaller strip size. The default is 64 Kbytes.

Note: The ServeRAID M1015 and M1115 SAS/SATA controllers do not support the cache policies. These policies include access policy, read policy, write policy, IO policy, and drive cache.

- **Access Policy:** Select the type of data access that is allowed for this virtual drive:
 - ◇ *RW:* Allow read/write access. This is the default.
 - ◇ *Read Only:* Allow read-only access.
 - ◇ *Blocked:* Do not allow access.
- **Read Policy:** Specify the read policy for this virtual drive:
 - ◇ *Normal:* This disables the read ahead capability. This is the default.
 - ◇ *Ahead:* This enables read ahead capability, which allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data, but there is little improvement when accessing random data.
- **Write Policy:** Specify the write policy for this virtual drive:
 - ◇ *WBack:* In Writeback mode the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. This setting is recommended in Standard mode.
 - ◇ *WThru:* In Writethrough mode the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction. This is the default.
 - ◇ *Write Back with BBU:* Select this mode if you want the controller to use Writeback mode but the controller has no BBU or the BBU is bad. If you do not choose this option, the controller firmware automatically switches to Writethrough mode if it detects a bad BBU or missing BBU.

Attention: You can use Writeback mode with or without a battery. You should use either a battery to protect the controller cache or an uninterruptible power supply (UPS) to protect the entire system. If you do not use a battery or a UPS, and there is a power failure, you risk losing the data in the controller cache.


Although you can enable or disable the disk cache, you should disable it. If you enable the disk cache, the drive

sends a data transfer completion signal to the controller when the drive cache has received all the data in a transaction. The data has not been actually transferred to the disk media. If a power failure happens, you risk losing the data in the disk cache. This data will be unrecoverable.

- **IO Policy:** The IO Policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.
 - ◇ *Direct:* In direct I/O mode, reads are not buffered in cache memory. Data is transferred to the cache and the host concurrently. If the same data block is read again, it comes from cache memory. This is the default.
 - ◇ *Cached:* In cached I/O mode, all reads are buffered in cache memory.
 - **Drive Cache:** Specify the drive cache policy:
 - ◇ *Enable:* Enable the drive cache.
 - ◇ *Disable:* Disable the drive cache.
 - ◇ *NoChange:* Leave the current drive cache policy as is. This is the default.
 - **Disable BGI:** Specify the background initialization status:
 - ◇ *No:* Leave background initialization enabled. This means that a new configuration can be initialized in the background while you use WebBIOS to do other configuration tasks. This is the default.
 - ◇ *Yes:* Select Yes if you do not want to allow background initializations for configurations on this controller.
 - **Select Size:** Specify the size of the virtual drive in megabytes, gigabytes, or terabytes. Normally, this would be the full size for RAID 0 shown in the Configuration panel on the right. You can specify a smaller size if you want to create other virtual drives on the same drive group.
 - **Update Size:** Click **Update Size** to update the Select size field value for the selected RAID levels
8. Click **Accept** to accept the changes to the virtual drive definition, or click **Reclaim** to return to the previous settings.
 9. Click **Next** when you are finished defining virtual drives.

The Configuration Preview screen appears, as shown in the following figure.

Figure 4.6 RAID 0 Configuration Preview



10. Check the information in the configuration preview.
11. If the virtual drive configuration is acceptable, click **Accept** to save the configuration. Otherwise, click **Back** to return to the previous screens and change the configuration.
12. If you accept the configuration, click **Yes** at the prompt to save the configuration.

The WebBIOS main menu appears.

4.4.3.2 Using Manual Configuration: RAID 1

In RAID 1, the RAID controller duplicates all data from one drive to a second drive. RAID 1 provides complete data redundancy, but at the cost of doubling the required data storage capacity. It is appropriate for small databases or any other environment that requires fault tolerance but small capacity.

When you select **Manual Configuration** and click **Next**, the Disk Group Definition screen appears. You use this screen to select drives to create drive groups.


1. Hold <Ctrl> while you select two ready drives in the Drives panel on the left. You must select an even number of drives.
2. Click **Add To Array** to move the drives to a proposed drive group configuration in the Disk Groups panel on the right, as shown in the following figure.

If you need to undo the changes, click the **Reclaim** button.

3. Choose whether to use power save mode.
4. Choose whether to use drive encryption.

Note: A RAID 1 virtual drive can contain up to 16 drive groups and up to 32 drives in a single span. (Other factors, such as the type of controller, can limit the number of drives.) You must use two drives in each RAID 1 drive group in the span.

Figure 4.7 WebBIOS Disk Group Definition Screen



5. When you have finished selecting drives for the drive group, click **Accept DG**.
6. Click **Next**.

The Virtual Drive Definition screen appears, as shown in the following figure. You use this screen to select the RAID level, stripe size, read policy, and other attributes for the new virtual drives.

Figure 4.8 WebBIOS Virtual Drive Definition Screen

The screenshot shows the 'MegaRAID BIOS Config Utility Config Wizard - Virtual Drive Definition' window. On the left, there are several configuration options, each with a label and a drop-down menu:

- RAID Level:** RAID 1
- Strip Size:** 64 KB
- Access Policy:** RW
- Read Policy:** Ahead
- Write Policy:** Write Back with BBU
- IO Policy:** Direct
- Drive Cache:** Disable
- Disable BGI:** No
- Select Size:** (empty) GB

Below these options are buttons for 'Accept' and 'Reclaim'. On the right side, there is a section titled 'Virtual Drives' which contains a large empty box. Below this box, it says 'Next LD, Possible RAID Levels' followed by 'R0:557.750 GB R1:278.875 GB'. At the bottom right of the window are buttons for 'Cancel', 'Back', and 'Next'.

7. Change the virtual drive options from the defaults listed on the screen as needed.

Here are brief explanations of the virtual drive options:

- **RAID Level:** The drop-down menu lists the possible RAID levels for the virtual drive. Select RAID 1.
- **Stripe Size:** The stripe size specifies the length of the data segments that the RAID controller writes across multiple drives, not including parity drives. For example, consider a stripe that contains 64 KB of drive space and has 16 KB of data residing on each drive in the stripe. In this case, the stripe size is 64 KB and the strip size is 16 KB. You can set the stripe size to 8, 16, 32, 64, 128, 256, 512, and 1024 Kbytes. A larger stripe size produces higher read performance. If your computer regularly performs random read requests, choose a smaller stripe size. The default is 64 Kbytes.

Note: The ServeRAID M1015 and M1115 SAS/SATA controllers do not support the cache policies. These policies include access policy, read policy, write policy, IO policy, and drive cache.

- **Access Policy:** Select the type of data access that is allowed for this virtual drive:
 - ◇ *RW:* Allow read/write access. This is the default.
 - ◇ *Read Only:* Allow read-only access.
 - ◇ *Blocked:* Do not allow access.
- **Read Policy:** Specify the read policy for this virtual drive:
 - ◇ *Normal:* This disables the read ahead capability. This is the default.
 - ◇ *Ahead:* This enables read ahead capability, which allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data, but there is little improvement when accessing random data.
- **Write Policy:** Specify the write policy for this virtual drive:
 - ◇ *WBack:* In Writeback mode the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. This setting is recommended in Standard mode.
 - ◇ *WThru:* In Writethrough mode the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction. This is the default.
 - ◇ *Write Back with BBU:* Select this mode if you want the controller to use Writeback mode but the controller has no BBU or the BBU is bad. If you do not choose this option, the controller firmware automatically switches to Writethrough mode if it detects a bad or missing BBU.

Attention: You can use Writeback mode with or without a battery. You should use either a battery to protect the controller cache or an uninterruptible power supply (UPS) to protect the entire system. If you do not use a battery or a UPS, and there is a power failure, you risk losing the data in the controller cache.

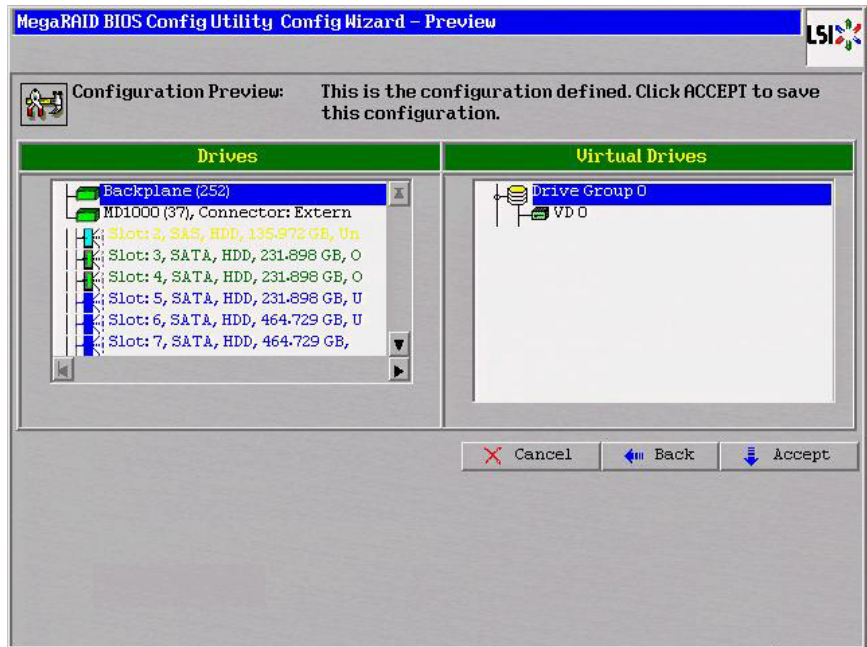
Although you can enable or disable the disk cache, you should disable it. If you enable the disk cache, the drive

sends a data transfer completion signal to the controller when the drive cache has received all the data in a transaction. The data has not been actually transferred to the disk media. If a power failure happens, you risk losing the data in the disk cache. This data will be unrecoverable.

- **IO Policy:** The IO Policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.
 - ◇ *Direct:* In Direct I/O mode, reads are not buffered in cache memory. Data is transferred to the cache and the host concurrently. If the same data block is read again, it comes from cache memory. This is the default.
 - ◇ *Cached:* In Cached I/O mode, all reads are buffered in cache memory.
 - **Drive Policy:** Specify the drive cache policy:
 - ◇ *Enable:* Enable the drive cache.
 - ◇ *Disable:* Disable the drive cache.
 - ◇ *NoChange:* Leave the current drive cache policy as is. This drive policy is the default.
 - **Disable BGI:** Specify the background initialization status:
 - ◇ *No:* Leave background initialization enabled. This means that a new configuration can be initialized in the background while you use WebBIOS to do other configuration tasks. This is the default.
 - ◇ *Yes:* Select Yes if you do not want to allow background initializations for configurations on this controller.
 - **Select Size:** Specify the size of the virtual drive(s) in megabytes, gigabytes, or terabytes. Normally, this would be the full size for RAID 1 shown in the Configuration panel on the right. You may specify a smaller size if you want to create other virtual drives on the same drive group.
 - **Update Size:** Click *Update Size* to update the Select size field value for the selected RAID levels
8. Click **Accept** to accept the changes to the virtual drive definition, or click **Reclaim** to return to the previous settings.
 9. Click **Next** when you are finished defining virtual drives.

The Configuration Preview screen appears, as shown in the following figure.

Figure 4.9 RAID 1 Configuration Preview



10. Check the information in the configuration preview.
11. If the virtual drive configuration is acceptable, click **Accept** to save the configuration. Otherwise, click **Back** to return to the previous screens and change the configuration.
12. If you accept the configuration, click **Yes** at the prompt to save the configuration.

The WebBIOS main menu appears.

4.4.3.3 Using Manual Configuration: RAID 5

RAID 5 uses drive striping at the block level and parity. In RAID 5, the parity information is written to all drives. It is best suited for networks that perform a lot of small input/output (I/O) transactions simultaneously. RAID 5 provides data redundancy, high read rates, and good performance in most environments. It also provides redundancy with lowest loss of capacity.

RAID 5 provides high data throughput. RAID 5 is useful for transaction processing applications because each drive can read and write independently. If a drive fails, the RAID controller uses the parity drive to recreate all missing information. You can use RAID 5 for office automation and online customer service that require fault tolerance. In addition, RAID 5 is good for any application that has high read request rates but low write request rates.


When you select **Manual Configuration** and click **Next**, the Disk Group Definition screen appears. You use this screen to select drives to create drive groups.

1. Hold <Ctrl> while you select at least three ready drives in the Physical Drives panel on the left.
2. Click **Add To Array** to move the drives to a proposed drive group configuration in the Disk Groups panel on the right, as shown in the following figure.

If you need to undo the changes, click the **Reclaim** button.

3. Choose whether to use power save mode.
4. Choose whether to use drive encryption.

Figure 4.10 WebBIOS Disk Group Definition Screen



5. When you have finished selecting drives for the drive group, click **Accept DG**.

6. Click **Next**.

The Virtual Drive Definition screen appears, as shown in the following figure. You use this screen to select the RAID level, stripe size, read policy, and other attributes for the new virtual drives.

Figure 4.11 WebBIOS Virtual Drive Definition Screen

The screenshot displays the 'MegaRAID BIOS Config Utility Config Wizard - Virtual Drive Definition' window. On the left, there is a list of configuration options, each with a dropdown menu:

- RAID Level:** RAID 5
- Strip Size:** 64 KB
- Access Policy:** RW
- Read Policy:** Ahead
- Write Policy:** Write Back with BBU
- IO Policy:** Direct
- Drive Cache:** Disable
- Disable BGI:** No
- Select Size:** (empty) GB

On the right, the 'Virtual Drives' section contains a large empty box. Below it, the text 'Next LD, Possible RAID Levels' is displayed in green, followed by the drive configuration: 'R0:1.089 TB R1:557.750 GB R5:836.625 GB R6: 557.750 GB'. At the bottom of the window, there are navigation buttons: 'Accept' (with a blue arrow icon), 'Reclaim' (with a blue icon), 'Cancel' (with a red X icon), 'Back' (with a blue left arrow icon), and 'Next' (with a blue right arrow icon).

7. Change the virtual drive options from the defaults listed on the screen as needed.

Here are brief explanations of the virtual drive options:

- **RAID Level:** The drop-down menu lists the possible RAID levels for the virtual drive. Select RAID 5.
- **Stripe Size:** The stripe size specifies the length of the data segments that the RAID controller writes across multiple drives, not including parity drives. For example, consider a stripe that contains 64 KB of drive space and has 16 KB of data residing on each drive in the stripe. In this case, the stripe size is 64 KB and the strip size is 16 KB. You can set the stripe size to 8, 16, 32, 64, 128, 256, 512, and 1024 Kbytes. A larger stripe size produces higher read performance. If your computer regularly performs random read requests, choose a smaller stripe size. The default is 64 Kbytes.

Note: The ServeRAID M1015 and M1115 SAS/SATA controllers do not support the cache policies. These policies include access policy, read policy, write policy, IO policy, and drive cache.

- **Access Policy:** Select the type of data access that is allowed for this virtual drive:
 - ◇ *RW:* Allow read/write access. This is the default.
 - ◇ *Read Only:* Allow read-only access.
 - ◇ *Blocked:* Do not allow access.
- **Read Policy:** Specify the read policy for this virtual drive:
 - ◇ *Normal:* This disables the read ahead capability. This is the default.
 - ◇ *Ahead:* This enables read ahead capability, which allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data, but there is little improvement when accessing random data.
- **Write Policy:** Specify the write policy for this virtual drive:
 - ◇ *WBack:* In Writeback mode the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. This setting is recommended in Standard mode.
 - ◇ *WThru:* In Writethrough mode the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction. This is the default.
 - ◇ *Write Back with BBU:* Select this mode if you want the controller to use Writeback mode but the controller has no BBU or the BBU is bad. If you do not choose this option, the controller firmware automatically switches to Writethrough mode if it detects a bad or missing BBU.


Attention: You can use Writeback mode with or without a battery. You should use either a battery to protect the controller cache or an uninterruptible power supply (UPS) to protect the entire system. If you do not use a battery or a UPS, and there is a power failure, you risk losing the data in the controller cache.

Although you can enable or disable the disk cache, you should disable it. If you enable the disk cache, the drive

sends a data transfer completion signal to the controller when the drive cache has received all the data in a transaction. The data has not been actually transferred to the disk media. If a power failure happens, you risk losing the data in the disk cache. This data will be unrecoverable.

- **IO Policy:** The IO Policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.
 - ◇ *Direct:* In Direct I/O mode, reads are not buffered in cache memory. Data is transferred to the cache and the host concurrently. If the same data block is read again, it comes from cache memory. This is the default.
 - ◇ *Cached:* In Cached I/O mode, all reads are buffered in cache memory.
 - **Drive Policy:** Specify the drive cache policy:
 - ◇ *Enable:* Enable the drive cache.
 - ◇ *Disable:* Disable the drive cache.
 - ◇ *NoChange:* Leave the current drive cache policy as is. This drive policy is the default.
 - **Disable BGI:** Specify the background initialization status:
 - ◇ *No:* Leave background initialization enabled. This means that a new configuration can be initialized in the background while you use WebBIOS to do other configuration tasks. This is the default.
 - ◇ *Yes:* Select Yes if you do not want to allow background initializations for configurations on this controller.
 - **Select Size:** Specify the size of the virtual drive in megabytes, gigabytes, or terabytes. Normally, this would be the full size for RAID 5 shown in the Configuration panel on the right. You may specify a smaller size if you want to create other virtual drives on the same drive group.
8. Click **Accept** to accept the changes to the virtual drive definition, or click **Reclaim** to return to the previous settings.
 9. Click **Next** when you are finished defining virtual drives.
- The Configuration Preview screen appears, as shown in the following figure.

Figure 4.12 RAID 5 Configuration Preview



10. Check the information in the configuration preview.
11. If the virtual drive configuration is acceptable, click **Accept** to save the configuration. Otherwise, click **Back** to return to the previous screens and change the configuration.
12. If you accept the configuration, click **Yes** at the prompt to save the configuration.

The WebBIOS main menu appears.

4.4.3.4 Using Manual Configuration: RAID 6

RAID 6 is similar to RAID 5 (drive striping and distributed parity), except that instead of one parity block per stripe, there are two. With two independent parity blocks, RAID 6 can survive the loss of two drives in a virtual drive without losing data. Use RAID 6 for data that requires a very high level of protection from loss.

RAID 6 is best suited for networks that perform a lot of small input/output (I/O) transactions simultaneously. It provides data redundancy, high read rates, and good performance in most environments.

In the case of a failure of one drive or two drives in a virtual drive, the RAID controller uses the parity blocks to recreate all of the missing information. If two drives in a RAID 6 virtual drive fail, two drive rebuilds are required, one for each drive. These rebuilds do not occur at the same time. The controller rebuilds one failed drive, and then the other failed drive.

When you select **Manual Configuration** and click **Next**, the drive Group Definition screen appears. You use this screen to select drives to create drive groups.


1. Hold <Ctrl> while selecting at least three ready drives in the Drives panel on the left.
2. Click **Add To Array** to move the drives to a proposed drive group configuration in the Disk Groups panel on the right, as shown in the following figure.

If you need to undo the changes, click the **Reclaim** button.

3. Choose whether to use power save mode.
4. Choose whether to use drive encryption.

The drop-down list in the **Encryption** field lists the options.

Figure 4.13 WebBIOS Disk Group Definition Screen



5. When you have finished selecting drives for the drive group, click **Accept DG** for each.
6. Click **Next**.

The Virtual Drive Definition screen appears, as shown in the following figure. Use this screen to select the RAID level, stripe size, read policy, and other attributes for the new virtual drives.

Figure 4.14 WebBIOS Virtual Drive Definition Screen

MegaRAID BIOS Config Utility Config Wizard - Virtual Drive Definition

RAID Level: RAID 6

Strip Size: 64 KB

Access Policy: RW

Read Policy: Ahead

Write Policy: Write Back with BBU

IO Policy: Direct

Drive Cache: Disable

Disable BGI: No

Select Size: GB

Update Size

Virtual Drives

Next LD, Possible RAID Levels
R0: 1.361 TB R5: 1.089 TB R6: 836.625 GB

Accept Reclaim

Cancel Back Next

7. Change the virtual drive options from the defaults listed on the screen as needed.

Here are brief explanations of the virtual drive options:

- **RAID Level:** The drop-down menu lists the possible RAID levels for the virtual drive. Select RAID 6.
- **Stripe Size:** The stripe size specifies the length of the data segments that the RAID controller writes across multiple drives, not including parity drives. For example, consider a stripe that contains 64 KB of drive space and has 16 KB of data residing on each drive in the stripe. In this case, the stripe size is 64 KB and the strip size is 16 KB. You can set the stripe size to 8, 16, 32, 64, 128, 256, 512, and 1024 Kbytes. A larger stripe size produces higher read performance. If your computer regularly performs random read requests, choose a smaller stripe size. The default is 64 Kbytes.

Note: WebBIOS does not allow you to select 8 Kbytes as the stripe size when you create a RAID 6 drive group with three drives.

Note: The ServeRAID M1015 and M1115 SAS/SATA controllers do not support the cache policies. These policies include access policy, read policy, write policy, IO policy, and drive cache.

- **Access Policy:** Select the type of data access that is allowed for this virtual drive:
 - ◇ *RW:* Allow read/write access. This is the default.
 - ◇ *Read Only:* Allow read-only access.
 - ◇ *Blocked:* Do not allow access.
- **Read Policy:** Specify the read policy for this virtual drive:
 - ◇ *Normal:* This disables the read ahead capability. This is the default.
 - ◇ *Ahead:* This enables read ahead capability, which allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data, but there is little improvement when accessing random data.
- **Write Policy:** Specify the write policy for this virtual drive:
 - ◇ *WBack:* In Writeback mode the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. This setting is recommended in Standard mode.
 - ◇ *WThru:* In Writethrough mode the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction. This is the default.
 - ◇ *Write Back with BBU:* Select this mode if you want the controller to use Writeback mode but the controller has no BBU or the BBU is bad. If you do not choose this option, the controller firmware automatically switches to Writethrough mode if it detects a bad or missing BBU.

Attention: You can use Writeback mode with or without a battery. You should use either a battery to protect the controller cache or an uninterruptible power supply (UPS) to protect the entire system. If you do not use a battery or a UPS, and

there is a power failure, you risk losing the data in the controller cache.


Although you can enable or disable the disk cache, you should disable it. If you enable the disk cache, the drive sends a data transfer completion signal to the controller when the drive cache has received all the data in a transaction. The data has not been actually transferred to the disk media. If a power failure happens, you risk losing the data in the disk cache. This data will be unrecoverable.

- **IO Policy:** The IO Policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.
 - ◇ *Direct:* In Direct I/O mode, reads are not buffered in cache memory. Data is transferred to the cache and the host concurrently. If the same data block is read again, it comes from cache memory. This is the default.
 - ◇ *Cached:* In Cached I/O mode, all reads are buffered in cache memory.
 - **Drive Policy:** Specify the drive cache policy:
 - ◇ *Enable:* Enable the drive cache.
 - ◇ *Disable:* Disable the drive cache.
 - ◇ *NoChange:* Leave the current drive cache policy as is. This drive policy is the default.
 - **Disable BGI:** Specify the background initialization status:
 - ◇ *No:* Leave background initialization enabled. This means that a new configuration can be initialized in the background while you use WebBIOS to do other configuration tasks. This is the default.
 - ◇ *Yes:* Select Yes if you do not want to allow background initializations for configurations on this controller.
 - **Select Size:** Specify the size of the virtual drive in megabytes, gigabytes, or terabytes. Normally, this would be the full size for RAID 6 shown in the Configuration panel on the right. You may specify a smaller size if you want to create other virtual drives on the same drive group.
8. Click **Accept** to accept the changes to the virtual drive definition, or click **Reclaim** to return to the previous settings.

9. Click **Next** when you are finished defining virtual drives.

The Configuration Preview screen appears, as shown in the following figure.

Figure 4.15 RAID 6 Configuration Preview



10. Check the information in the configuration preview.
11. If the virtual drive configuration is acceptable, click **Accept** to save the configuration. Otherwise, click **Back** to return to the previous screens and change the configuration.
12. If you accept the configuration, click **Yes** at the prompt to save the configuration.

The WebBIOS main menu appears.

4.4.3.5 Using Manual Configuration: RAID 10

RAID 10, a combination of RAID 1 and RAID 0, has mirrored drives. It breaks up data into smaller blocks, then stripes the blocks of data to each RAID 1 drive group. Each RAID 1 drive group then duplicates its data to its other drive. The size of each block is determined by the stripe

size parameter, which is 64 Kbytes. RAID 10 can sustain one drive failure in each drive group while maintaining data integrity.

RAID 10 provides both high data transfer rates and complete data redundancy. It works best for data storage that must have 100 percent redundancy of RAID 1 (mirrored drive groups) and that also needs the enhanced I/O performance of RAID 0 (striped drive groups); it works well for medium-sized databases or any environment that requires a higher degree of fault tolerance and moderate to medium capacity.

When you select **Manual Configuration** and click **Next**, the Disk Group Definition screen appears.

You use the Drive Group Definition screen to select drives to create drive groups.

1. Hold <Ctrl> while selecting two ready drives in the Drives panel on the left.
2. Click **Add To Array** to move the drives to a proposed two-drive drive group configuration in the Drive Groups panel on the right.

If you need to undo the changes, click the **Reclaim** button.

3. Click **Accept DG** to create a RAID 1 drive group.

An icon for the next drive group displays in the right panel.

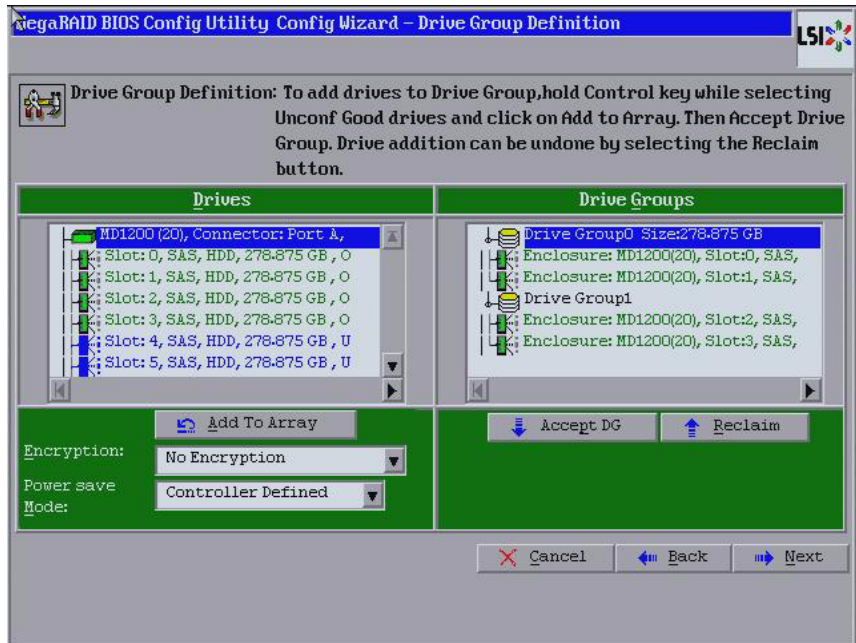
4. Click on the icon for the next drive group to select it.
5. Hold <Ctrl> while selecting two more ready drives in the Drives panel to create a second RAID 1 drive group with two drives.
6. Click **Add To Array** to move the drives to a second two-drive drive group configuration in the Drive Groups panel, as shown in the following figure.

If you need to undo the changes, click the **Reclaim** button.

7. Choose whether to use power saving.
8. Choose whether to use drive encryption.

Note: RAID 10 supports a maximum of eight spans, with a maximum of 32 drives per span. (Other factors, such as the type of controller, can limit the number of drives.) You must use an even number of drives in each RAID 10 drive group in the span.


Figure 4.16 WebBIOS Drive Group Definition Screen



9. Repeat [step 4](#) to [step 6](#) until you have selected all the drives you want for the drive groups.
10. When you have finished selecting drives for the drive groups, select each drive group and click **Accept DG** for each.
11. Click **Next**.

The Span Definition screen appears, as shown in the following figure. This screen displays the drive group holes you can select to add to a span.

Figure 4.17 WebBIOS Span Definition Screen



12. Under the heading Array With Free Space, hold <Ctrl> while you select a drive group with two drives, and click **Add to SPAN**.

The drive group you select displays in the right frame under the heading Span.

13. Hold <Ctrl> while you select a second drive group with two drives, and click **Add to SPAN**.

Both drive groups display in the right frame under Span.

14. If there are additional drive groups with two drives each, you can add them to the virtual drive.

15. Click **Next**.

The Virtual Drive Definition screen appears, as shown in the following figure. You use this screen to select the RAID level, stripe size, read policy, and other attributes for the new virtual drives.

16. Hold <Ctrl> while you select two drive groups with two drives in the Configuration panel on the right.

Figure 4.18 WebBIOS Virtual Drive Definition Screen

MegaRAID BIOS Config Utility Config Wizard - Virtual Drive Definition

RAID Level: RAID 10

Strip Size: 64 KB

Access Policy: RW

Read Policy: Ahead

Write Policy: Write Back with BBU

IO Policy: Direct

Drive Cache: Disable

Disable BGI: No

Select Size: GB

Update Size

Virtual Drives

Next LD, Possible RAID Levels

R00:1.069 TB R10:557.750 GB

Accept Reclaim

Cancel Back Next

Note: The WebBIOS Configuration Utility shows the maximum available capacity while creating the RAID 10 drive group. In version 1.03 of the utility, the maximum capacity of the RAID 10 drive group is the sum total of the two RAID 1 drive groups. In version 1.1, the maximum capacity is the capacity of the smaller drive group multiplied by two.

17. Change the virtual drive options from the defaults listed on the screen as needed.

Here are brief explanations of the virtual drive options:

- **RAID Level:** The drop-down menu lists the possible RAID levels for the virtual drive. Select RAID 10.
- **Stripe Size:** The stripe size specifies the length of the data segments that the RAID controller writes across multiple drives, not including parity drives. For example, consider a stripe that contains 64 KB of drive space and has 16 KB of data residing on each drive in the stripe. In this case, the stripe size is 64 KB and the strip size is 16 KB. You can set the stripe size to 8, 16, 32, 64, 128, 256, 512, and 1024 Kbytes. A larger stripe size produces higher read

performance. If your computer regularly performs random read requests, choose a smaller stripe size. The default is 64 Kbytes.

Note: The ServeRAID M1015 and M1115 SAS/SATA controllers do not support the cache policies. These policies include access policy, read policy, write policy, IO policy, and drive cache.

- **Access Policy:** Select the type of data access that is allowed for this virtual drive:
 - ◇ *RW:* Allow read/write access.
 - ◇ *Read Only:* Allow read-only access. This is the default.
 - ◇ *Blocked:* Do not allow access.
- **Read Policy:** Specify the read policy for this virtual drive:
 - ◇ *Normal:* This disables the read ahead capability. This is the default.
 - ◇ *Ahead:* This enables read ahead capability, which allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data, but there is little improvement when accessing random data.
- **Write Policy:** Specify the write policy for this virtual drive:
 - ◇ *WBack:* In Writeback mode the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. This setting is recommended in Standard mode.
 - ◇ *WThru:* In Writethrough mode the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction. This is the default.
 - ◇ *Write Back with BBU:* Select this mode if you want the controller to use Writeback mode but the controller has no BBU or the BBU is bad. If you do not choose this option, the controller firmware automatically switches to Writethrough mode if it detects a bad or missing BBU.

Attention: You can use Writeback mode with or without a battery. You should use either a battery to protect the controller cache or an uninterruptible power supply (UPS) to protect the entire system. If you do not use a battery or a UPS, and there is a power failure, you risk losing the data in the controller cache.

Although you can enable or disable the disk cache, you should disable it. If you enable the disk cache, the drive sends a data transfer completion signal to the controller when the drive cache has received all the data in a transaction. The data has not been actually transferred to the disk media. If a power failure happens, you risk losing the data in the disk cache. This data will be unrecoverable.

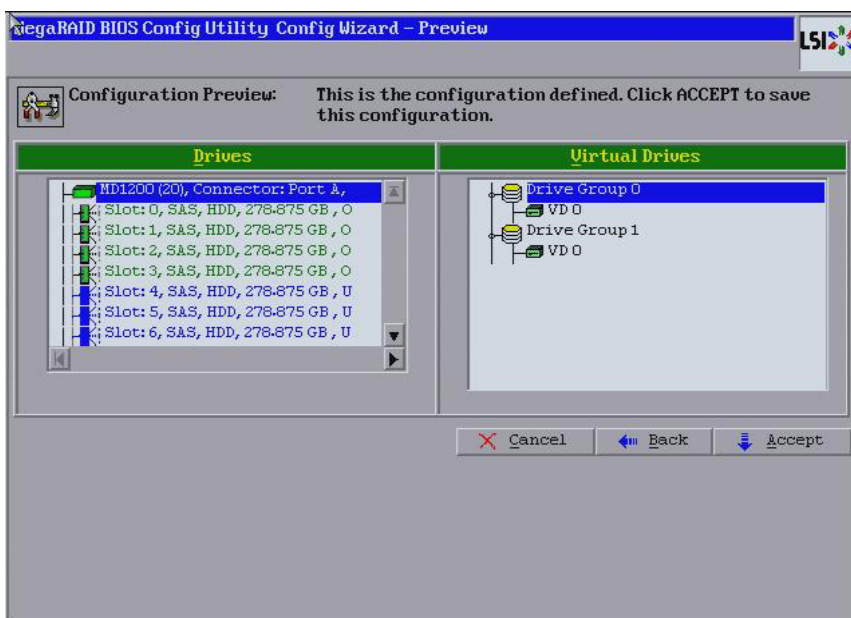
- **IO Policy:** The IO Policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.
 - ◇ *Direct:* In Direct I/O mode, reads are not buffered in cache memory. Data is transferred to the cache and the host concurrently. If the same data block is read again, it comes from cache memory. This is the default.
 - ◇ *Cached:* In Cached I/O mode, all reads are buffered in cache memory.
- **Drive Policy:** Specify the drive cache policy:
 - ◇ *Enable:* Enable the drive cache.
 - ◇ *Disable:* Disable the drive cache.
 - ◇ *NoChange:* Leave the current drive cache policy as is. This drive policy is the default.
- **Disable BGI:** Specify the background initialization status:
 - ◇ *No:* Leave background initialization enabled. This means that a new configuration can be initialized in the background while you use WebBIOS to do other configuration tasks. This is the default.
 - ◇ *Yes:* Select Yes if you do not want to allow background initializations for configurations on this controller.
- **Select Size:** Specify the size of the virtual drive in megabytes, gigabytes, or terabytes. Normally, this would be the full size for RAID 10 shown in the configuration panel on the right. You may specify a

smaller size if you want to create other virtual drives on the same drive group.

- **Update Size:** Click **Update Size** to update the Select size field value for the selected RAID levels.
18. Click **Accept** to accept the changes to the virtual drive definition, or click **Reclaim** to return to the previous settings.
 19. When you are finished defining virtual drives, click **Next**.

The Configuration Preview screen appears, as shown in the following figure.

Figure 4.19 RAID 10 Configuration Preview



20. Check the information in the configuration preview.
21. If the virtual drive configuration is acceptable, click **Accept** to save the configuration. Otherwise, click **Back** to return to the previous screens and change the configuration.
22. If you accept the configuration, click **Yes** at the prompt to save the configuration.

The WebBIOS main menu appears.

4.4.3.6 Using Manual Configuration: RAID 50

RAID 50 provides the features of both RAID 0 and RAID 5. RAID 50 uses both distributed parity and drive striping across multiple drive groups. It provides high data throughput, data redundancy, and very good performance. It is best implemented on two RAID 5 drive groups with data striped across both drive groups. Though multiple drive failures can be tolerated, only one drive failure can be tolerated in each RAID 5 level drive group.

RAID 50 is appropriate when used with data that requires high reliability, high request rates, high data transfer, and medium to large capacity.

When you select **Manual Configuration** and click **Next**, the Disk Group Definition screen appears. You use this screen to select drives to create drive group.

1. Hold <Ctrl> while selecting at least three ready drives in the Drives panel on the left.
2. Click **Add To Array** to move the drives to a proposed drive group configuration in the Disk Groups panel on the right.

If you need to undo the changes, click the **Reclaim** button.

3. Click **Accept DG** to create a RAID 5 drive group.

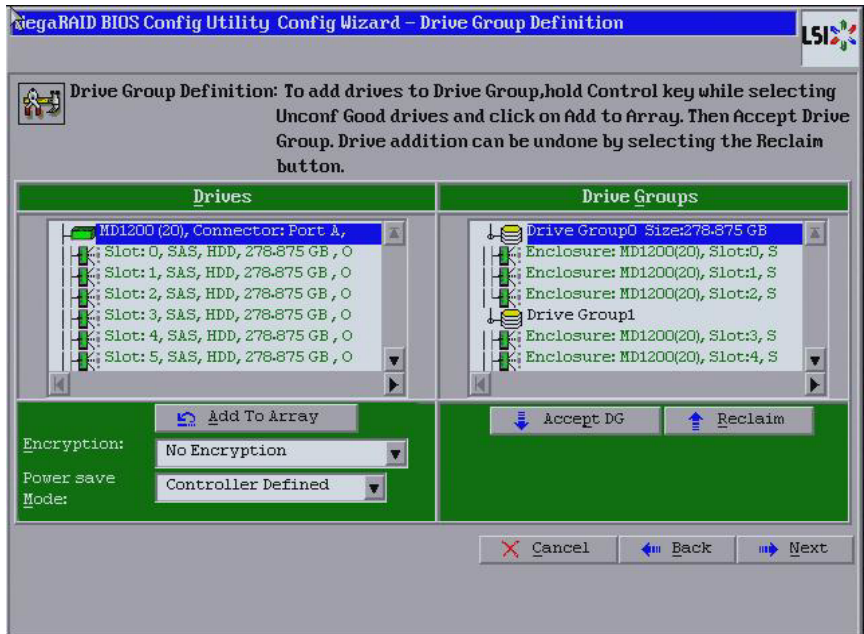
An icon for a second drive group displays in the right panel.

4. Click on the icon for the second drive group to select it.
5. Hold <Ctrl> while selecting at least three more ready drives in the Drives panel to create a second drive group.
6. Click **Add To Array** to move the drives to a proposed drive group configuration in the Disk Groups panel on the right, as shown in the following figure.

If you need to undo the changes, click the **Reclaim** button.

7. Choose whether to use drive encryption.

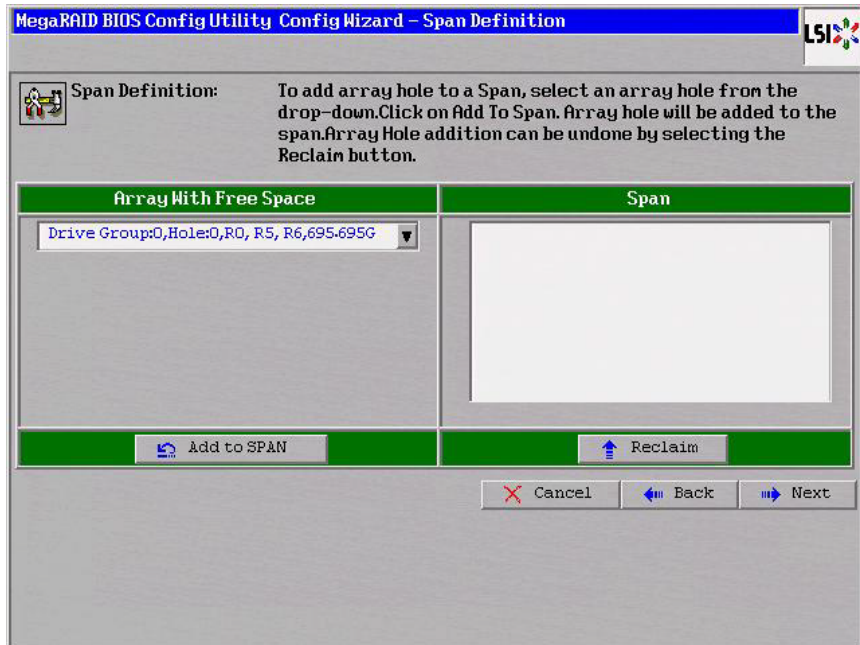
Figure 4.20 WebBIOS Disk Group Definition Screen



8. When you have finished selecting drives for the drive groups, select each drive group and click **Accept DG** for each.
9. Click **Next**.

The Span Definition screen appears, as shown in the following figure. This screen displays the drive group holes you can select to add to a span.

Figure 4.21 WebBIOS Span Definition Screen



10. Under the heading **Array With Free Space**, hold <Ctrl> while you select a drive group of three or more drives, and click **Add to SPAN**.
The drive group you select displays in the right frame under the heading **Span**.
11. Hold <Ctrl> while you select a second drive group of three or more drives, and click **Add to SPAN**.
Both drive groups display in the right frame under **Span**.
12. Click **Next**.
The Virtual Drive Definition screen appears, as shown in the following figure. You use this screen to select the RAID level, stripe size, read policy, and other attributes for the new virtual drive(s).
13. Hold <Ctrl> while you select two 3-drive drive groups in the Configuration panel on the right.

Figure 4.22 WebBIOS Virtual Drive Definition Screen

MegaRAID BIOS Config Utility Config Wizard - Virtual Drive Definition

RAID Level: RAID 50

Strip Size: 64 KB

Access Policy: RW

Read Policy: Ahead

Write Policy: Write Back with BBU

IO Policy: Direct

Drive Cache: Disable

Disable BGI: No

Select Size: GB

Virtual Drives

Next LD, Possible RAID Levels
R00: 1.634 TB R50: 1.089 TB R60: 557.750 GB

Accept Reclaim

Cancel Back Next

14. Change the virtual drive options from the defaults listed on the screen as needed.

Here are brief explanations of the virtual drive options:

- **RAID Level:** The drop-down menu lists the possible RAID levels for the virtual drive. Select RAID 50.
- **Stripe Size:** The stripe size specifies the length of the data segments that the RAID controller writes across multiple drives, not including parity drives. For example, consider a stripe that contains 64 KB of drive space and has 16 KB of data residing on each drive in the stripe. In this case, the stripe size is 64 KB and the strip size is 16 KB. You can set the stripe size to 8, 16, 32, 64, 128, 256, 512, and 1024 Kbytes. A larger stripe size produces higher read performance. If your computer regularly performs random read requests, choose a smaller stripe size. The default is 64 Kbytes.

Note: The ServeRAID M1015 and M1115 SAS/SATA controllers do not support the cache policies. These policies include access policy, read policy, write policy, IO policy, and drive cache.

- **Access Policy:** Select the type of data access that is allowed for this virtual drive:
 - ◇ *RW:* Allow read/write access.
 - ◇ *Read Only:* Allow read-only access. This is the default.
 - ◇ *Blocked:* Do not allow access.
- **Read Policy:** Specify the read policy for this virtual drive:
 - ◇ *Normal:* This disables the read ahead capability. This is the default.
 - ◇ *Ahead:* This enables read ahead capability, which allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data, but there is little improvement when accessing random data.
- **Write Policy:** Specify the write policy for this virtual drive:
 - ◇ *WBack:* In Writeback mode the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. This setting is recommended in Standard mode.
 - ◇ *WThru:* In Writethrough mode the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction. This is the default.
 - ◇ *Write Back with BBU:* Select this mode if you want the controller to use Writeback mode but the controller has no BBU or the BBU is bad. If you do not choose this option, the controller firmware automatically switches to Writethrough mode if it detects a bad or missing BBU.

Attention: You can use Writeback mode with or without a battery. You should use either a battery to protect the controller cache or an uninterruptible power supply (UPS) to protect the entire system. If you do not use a battery or a UPS, and there is a power failure, you risk losing the data in the controller cache.

Although you can enable or disable the disk cache, you should disable it. If you enable the disk cache, the drive


sends a data transfer completion signal to the controller when the drive cache has received all the data in a transaction. The data has not been actually transferred to the disk media. If a power failure happens, you risk losing the data in the disk cache. This data will be unrecoverable.

- **IO Policy:** The IO Policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.
 - ◇ *Direct:* In Direct I/O mode, reads are not buffered in cache memory. Data is transferred to the cache and the host concurrently. If the same data block is read again, it comes from cache memory. This is the default.
 - ◇ *Cached:* In Cached I/O mode, all reads are buffered in cache memory.
 - **Drive Policy: Specify the drive cache policy:**
 - ◇ *Enable:* Enable the drive cache.
 - ◇ *Disable:* Disable the drive cache. This drive policy is the default.
 - ◇ *NoChange:* Leave the current drive cache policy as is. This is the default.
 - **Disable BGI:** Specify the background initialization status:
 - ◇ *No:* Leave background initialization enabled. This means that a new configuration can be initialized in the background while you use WebBIOS to do other configuration tasks. This is the default.
 - ◇ *Yes:* Select **Yes** if you do not want to allow background initializations for configurations on this controller.
 - **Select Size:** Specify the size of the virtual drive in megabytes, gigabytes, or terabytes. Normally, this would be the full size for RAID 50 shown in the Configuration Panel on the right. You may specify a smaller size if you want to create other virtual drives on the same drive group.
 - **Update Size:** Click **Update Size** to update the Select size field value for the selected RAID levels.
15. Click **Accept** to accept the changes to the virtual drive definition or click **Reclaim** to return to the previous settings.

16. Click **Next** when you are finished defining virtual drives.

The Configuration Preview screen appears, as shown in the following figure.

Figure 4.23 RAID 50 Configuration Preview



17. Check the information in the configuration preview.
18. If the virtual drive configuration is acceptable, click **Accept** to save the configuration. Otherwise, click **Back** to return to the previous screens and change the configuration.
19. If you accept the configuration, click **Yes** at the prompt to save the configuration.

The WebBIOS main menu appears.

4.4.3.7 Using Manual Configuration: RAID 60

RAID 60 provides the features of both RAID 0 and RAID 6, and includes both parity and drive striping across multiple drive groups. RAID 6 supports two independent parity blocks per stripe. A RAID 60 virtual drive can survive the loss of two drives in each of the RAID 6 sets without losing data. RAID 60 is best implemented on two RAID 6 drive groups with data striped across both drive groups. Use RAID 60 for data that requires a very high level of protection from loss.

RAID 60 can support up to eight spans and tolerate up to 16 drive failures, though less than total drive capacity is available. Two drive failures can be tolerated in each RAID 6 level drive group.

RAID 60 is appropriate when used with data that requires high reliability, high request rates, high data transfer, and medium to large capacity.

When you select **Manual Configuration** and click **Next**, the Disk Group Definition screen appears. You use this screen to select drives to create drive groups.

1. Hold <Ctrl> while selecting at least three ready drives in the Drives panel on the left.
2. Click **Add To Array** to move the drives to a proposed drive group configuration in the Disk Groups panel on the right.

If you need to undo the changes, click the **Reclaim** button.

3. Click **Accept DG** to create a RAID 6 drive group.

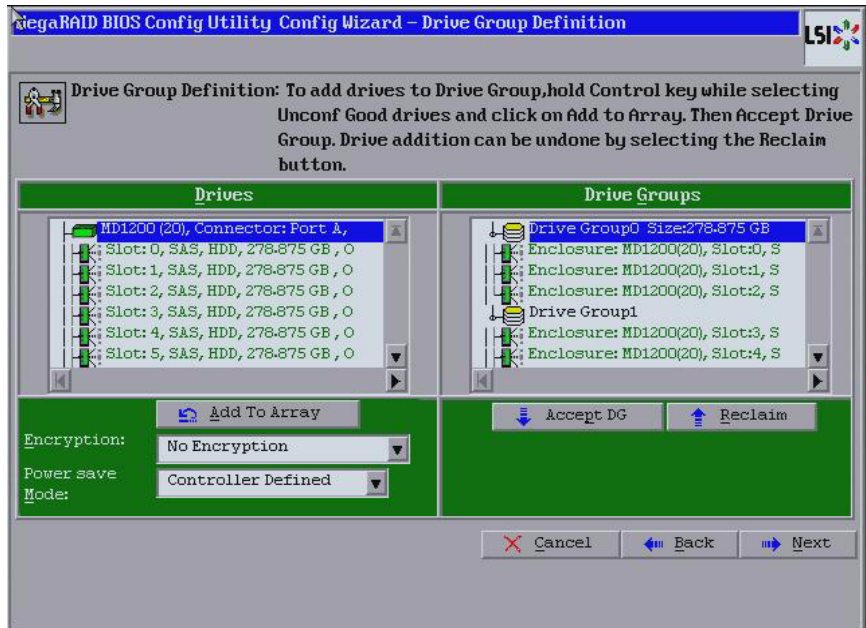
An icon for a second drive group displays in the right panel.

4. Click on the icon for the second drive group to select it.
5. Hold <Ctrl> while selecting at least three more ready drives in the Drives panel to create a second drive group.
6. Click **Add To Array** to move the drives to a proposed drive group configuration in the Disk Groups panel on the right, as shown in the following figure.

If you need to undo the changes, click the **Reclaim** button.

7. Choose whether to use power saving.

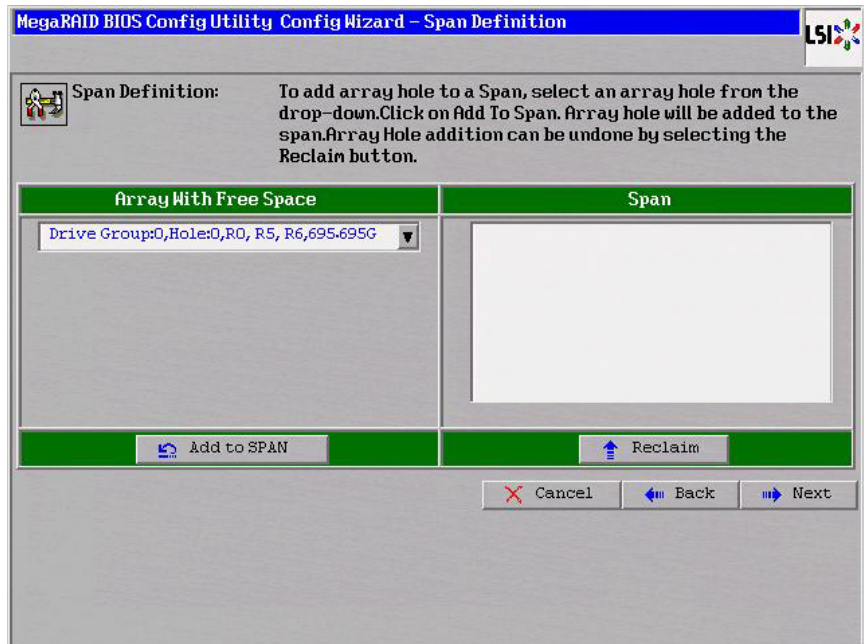
Figure 4.24 WebBIOS Disk Group Definition Screen



8. When you have finished selecting drives for the drive groups, select each drive group and click **Accept DG** for each.
9. Click **Next**.

The Span Definition screen appears, as shown in the following figure. This screen displays the drive group holes you can select to add to a span.

Figure 4.25 WebBIOS Span Definition Screen



10. Under the heading **Array With Free Space**, hold <Ctrl> while you select a drive group of three or more drives, and click **Add to SPAN**.
The drive group you select displays in the right frame under the heading **Span**.
11. Hold <Ctrl> while you select a second drive group of three or more drives, and click **Add to SPAN**.
Both drive groups display in the right frame under **Span**.
12. Click **Next**.
The Virtual Drive Definition screen appears, as shown in the following figure. You use this screen to select the RAID level, stripe size, read policy, and other attributes for the new virtual drive(s).
13. Hold <Ctrl> while you select two 3-drive drive groups in the Configuration window on the right.

Figure 4.26 WebBIOS Virtual Drive Definition Screen

The screenshot shows the 'MegaRAID BIOS Config Utility Config Wizard - Virtual Drive Definition' window. On the left, there are several configuration options with drop-down menus: RAID Level (RAID 60), Strip Size (64 KB), Access Policy (RW), Read Policy (Ahead), Write Policy (Write Back with BBU), IO Policy (Direct), Drive Cache (Disable), Disable BGI (No), and Select Size (with a unit of GB). To the right, under the 'Virtual Drives' heading, there is a large empty box and a green text label 'Next LD, Possible RAID Levels' followed by 'R00:1.634 TB R50:1.089 TB R60: 557.750 GB'. At the bottom, there are buttons for 'Accept', 'Reclaim', 'Cancel', 'Back', and 'Next'.

14. Change the virtual drive options from the defaults listed on the screen as needed.

Here are brief explanations of the virtual drive options:

- **RAID Level:** The drop-down menu lists the possible RAID levels for the virtual drive. Select RAID 60.
- **Stripe Size:** The stripe size specifies the length of the data segments that the RAID controller writes across multiple drives, not including parity drives. For example, consider a stripe that contains 64 KB of drive space and has 16 KB of data residing on each drive in the stripe. In this case, the stripe size is 64 KB and the strip size is 16 KB. You can set the stripe size to 8, 16, 32, 64, 128, 256, 512, and 1024 Kbytes. A larger stripe size produces higher read performance. If your computer regularly performs random read requests, choose a smaller stripe size. The default is 64 Kbytes.

Note: The ServeRAID M1015 and M1115 SAS/SATA controllers do not support the cache policies. These policies include

access policy, read policy, write policy, IO policy, and drive cache.

- **Access Policy:** Select the type of data access that is allowed for this virtual drive:
 - ◇ *RW:* Allow read/write access.
 - ◇ *Read Only:* Allow read-only access. This setting is the default.
 - ◇ *Blocked:* Do not allow access.
- **Read Policy:** Specify the read policy for this virtual drive:
 - ◇ *Normal:* Disables the read ahead capability. This is the default.
 - ◇ *Ahead:* Enables read ahead capability, which allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data, but there is little improvement when accessing random data.
- **Write Policy:** Specify the write policy for this virtual drive:
 - ◇ *WBack:* In Writeback mode the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. This setting is recommended in Standard mode.
 - ◇ *WThru:* In Writethrough mode the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction. This setting is the default.
 - ◇ *Write Back with BBU:* Select this mode if you want the controller to use Writeback mode but the controller has no BBU or the BBU is bad. If you do not choose this option, the controller firmware automatically switches to Writethrough mode if it detects a bad or missing BBU.

Attention: You can use Writeback mode with or without a battery. You should use either a battery to protect the controller cache or an uninterruptible power supply (UPS) to protect the entire system. If you do not use a battery or a UPS, and

there is a power failure, you risk losing the data in the controller cache.

Although you can enable or disable the disk cache, you should disable it. If you enable the disk cache, the drive sends a data transfer completion signal to the controller when the drive cache has received all the data in a transaction. The data has not been actually transferred to the disk media. If a power failure happens, you risk losing the data in the disk cache. This data will be unrecoverable.

- **IO Policy:** The IO Policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.
 - ◇ *Direct:* In Direct I/O mode, reads are not buffered in cache memory. Data is transferred to the cache and the host concurrently. If the same data block is read again, it comes from cache memory. This is the default.
 - ◇ *Cached:* In Cached I/O mode, all reads are buffered in cache memory.
- **Drive Policy:** Specify the drive cache IO Policy.
 - ◇ *Enable:* Enable the drive cache.
 - ◇ *Disable:* Disable the drive cache. This drive policy is the default.
 - ◇ *NoChange:* Leave the current drive cache policy as is. This is the default.
- **Disable BGI:** Specify the background initialization status:
 - ◇ *No:* Leave background initialization enabled. This means that a new configuration can be initialized in the background while you use WebBIOS to do other configuration tasks. This is the default.
 - ◇ *Yes:* Select Yes if you do not want to allow background initializations for configurations on this controller.
- **Select Size:** Specify the size of the virtual drive in megabytes, gigabytes, or terabytes. Normally, this would be the full size for RAID 60 shown in the Configuration panel on the right. You may specify a smaller size if you want to create other virtual drives on the same drive group.


- **Update Size:** Click **Update Size** to update the Select size field value for the selected RAID levels

Note: WebBIOS does not allow you to select 8 Kbytes as the stripe size when you create a RAID 60 drive group with six drives.

15. Click **Accept** to accept the changes to the virtual drive definition, or click **Reclaim** to return to the previous settings.
16. Click **Next** when you are finished defining virtual drives.

The Configuration Preview screen appears, as shown in the following figure.

Figure 4.27 RAID 60 Configuration Preview



17. Check the information in the configuration preview.
18. If the virtual drive configuration is acceptable, click **Accept** to save the configuration. Otherwise, or click **Back** to return to the previous screens and change the configuration.
19. If you accept the configuration, click **Yes** at the prompt to save the configuration.

The WebBIOS main menu appears.

4.5 Selecting Self-Encrypting Disk Security Options

The Self-Encrypting Disk (SED) feature provides the ability to encrypt data and use disk-based key management for the data security solution. This solution protects your data in case of theft or loss of physical drives. This section describes how to enable, change, or disable the drive security settings, and how to import a foreign configuration.


4.5.1 Enabling the Security Key Identifier, Security Key, and Passphrase

Perform the following steps to enable the encryption settings for the security key identifier, security key, and passphrase.

1. Click **Drive Security** on the main WebBIOS screen.

The Encryption Settings screen appears, as shown in the following figure.

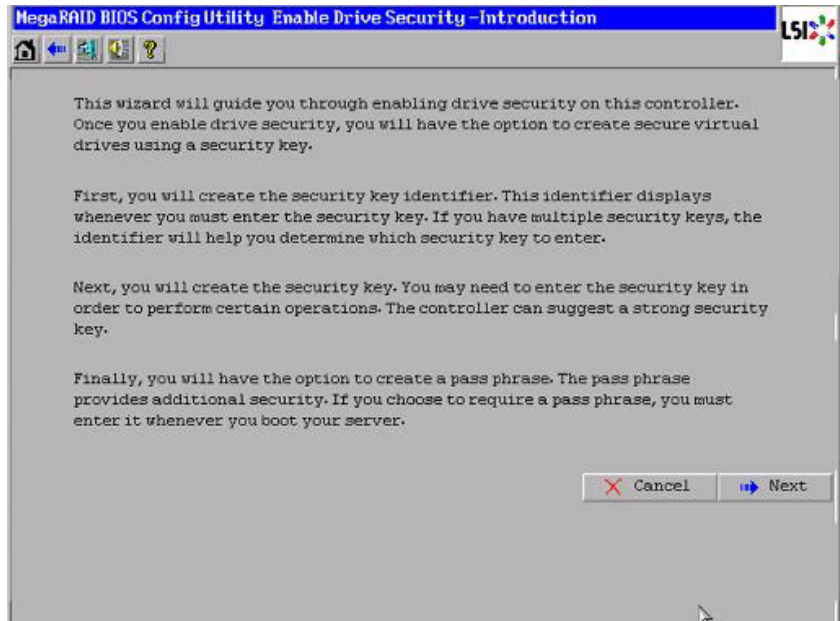
Figure 4.28 Encryption Settings Screen



2. To enable the drive security settings, select **Enable drive security** and click **Accept**.

The Enable Drive Security – Introduction screen appears as shown in the following figure. This screen lists the actions you can perform: editing the security key identifier, editing the security key, and adding or changing the pass phrase (optional).


Figure 4.29 Enable Drive Security - Introduction Screen



3. Click **Next**.

The screen used to create a security key identifier appears, as shown in the following figure.

Figure 4.30 Enable Drive Security – Enter Security Key ID Screen



4. Accept the default security key ID or enter a new security key ID.
5. Click **Next**.

The Enable Drive Security – Enter Security Key screen appears as shown in the following figure.

Figure 4.31 Enable Drive Security – Enter Security Key

MegaRAID BIOS Config Utility Enable Drive Security - Enter Security Key

Next, enter the security key. The security key is case-sensitive. It must be between eight and thirty-two characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (e.g. < > @ +). The space character is not permitted.

Note: For maximum security, use thirty-two varied characters. You may optionally choose for the system to suggest a strong security key.

Be sure to record the security key.

Security key:

Confirm:

6. Enter a new drive security key or click **Suggest** to fill the new security key. Enter the new drive security key again to confirm.

The security key is case-sensitive. It must be between eight and thirty-two characters and must contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (for example, < > @ +). The space character is not permitted.

7. Click **Next**.

The Enable Drive Security – Enter Pass Phrase screen appears as shown in the following figure. You have the option to provide a pass phrase for additional security.

Figure 4.32 Enable Drive Security – Enter Pass Phrase

8. If you want to use a pass phrase, click the checkbox **Use a pass phrase in addition to the security key**.
9. Enter a new pass phrase and then enter the new pass phrase again to confirm.

The pass phrase is case-sensitive. It must be between eight and thirty-two characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (for example, < > @ +). The space character is not permitted.

Non-US keyboard users must be careful not to enter DBCS characters in the pass phrase field or security key field. Firmware works only with the ASCII character set.

10. Click **Accept**.

The Confirm Enable Drive Security screen appears, as shown in the following figure.

Figure 4.33 Confirm Enable Drive Security Screen



11. Click **Yes** on the Confirm Enable Drive Security screen to confirm that you want to enable the drive security settings.

WebBIOS enables the security key ID, security key, and pass phrase (if applicable) that you entered and returns you to the main menu.

Attention: **If you forget the security key, you will lose access to your data.** Be sure to record your security key information. You might need to enter the security key to perform certain operations.


4.5.2 Changing the Security Key Identifier, Security Key, and Pass Phrase

If you selected disk-based encryption when you made the RAID configuration, the drive security will be enabled. Perform the following steps to change the encryption settings for the security key identifier, security key, and pass phrase.

1. Click **Drive Security** on the main WebBIOS screen.

The Encryption Settings screen appears as shown in the following figure.

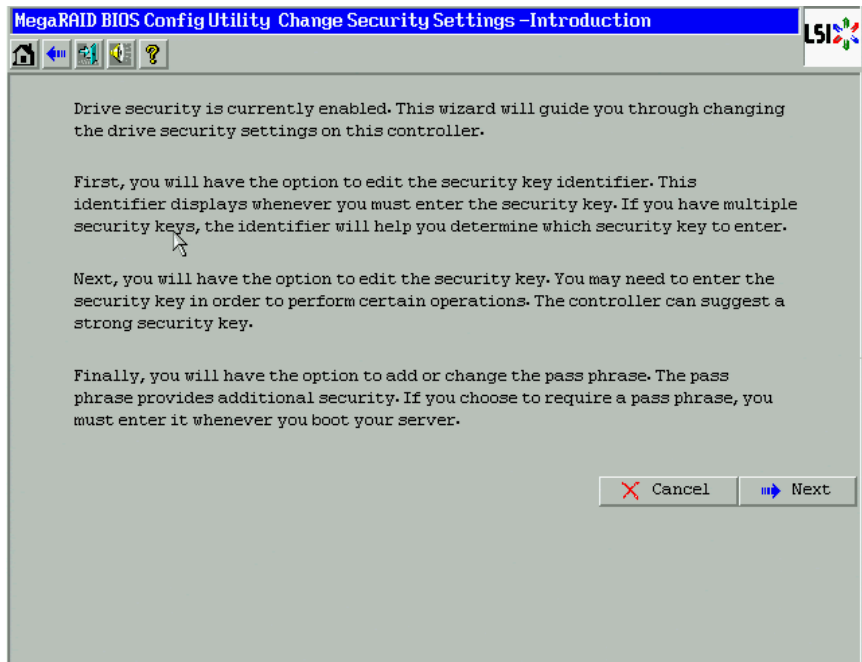
Figure 4.34 Encryption Settings Screen



2. To change the drive security settings, select **Change drive security settings...** and click **Accept**.

The Change Security Settings – Introduction screen appears as shown in the following figure. This screen lists the optional actions you can perform: editing the security key identifier, editing the security key, and adding or changing the pass phrase.

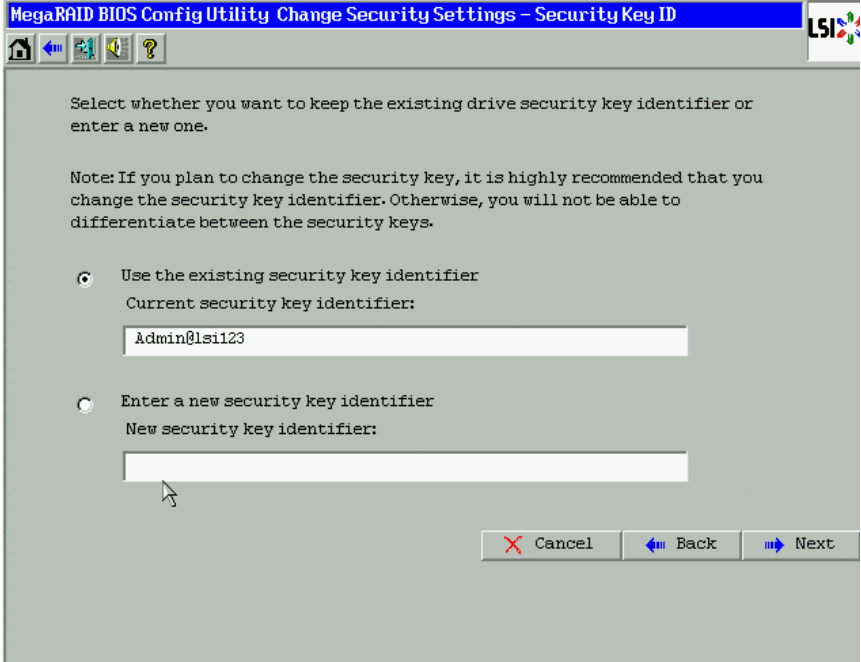
Figure 4.35 Change Security Settings – Introduction



3. To access the option to use the existing security key identifier or enter a new security key identifier, click **Next**.

The Change Security Settings – Security Key ID screen appears as shown in the following figure.

Figure 4.36 Change Security Settings – Security Key ID



MegaRAID BIOS Config Utility Change Security Settings – Security Key ID

Select whether you want to keep the existing drive security key identifier or enter a new one.

Note: If you plan to change the security key, it is highly recommended that you change the security key identifier. Otherwise, you will not be able to differentiate between the security keys.

☒ Use the existing security key identifier
Current security key identifier:
Admin@lsi123

☐ Enter a new security key identifier
New security key identifier:

Cancel Back Next

4. Choose whether you want to use the existing security key ID or enter a new security key ID. The options are:
 - Use the existing security key identifier (Current security key identifier).
 - Enter a new security key identifier (New security key identifier).
5. Click **Next**.

The Change Security Settings – Security Key screen appears as shown in the following figure. You have the option to use the existing security key or enter a new one.

Figure 4.37 Change Security Settings – Security Key

MegaRAID BIOS Config Utility Change Security Settings – Security Key

Select whether you want to keep the existing drive security key or enter a new one.

The security key is case-sensitive. It must be between eight and thirty-two characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (e.g. < > @ +). The space character is not permitted.

☒ Use the existing drive security key

☐ Enter a new drive security key

New security key:

Confirm:

Suggest

Cancel Back Next

6. Choose whether you want to use the existing security key or enter a new security key. The options are:
 - Use the existing drive security key.
 - Enter a new drive security key.
7. If you choose to enter a new drive security key, you can create a new drive security key or click **Suggest** to fill the new security key. Enter the new drive security key again to confirm.

The security key is case-sensitive. It must be between eight and thirty-two characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (for example, < > @ +). The space character is not permitted.

8. Click **Next**.

If you entered a new drive security key, the Authenticate Drive Security Key screen appears as shown in the following figure.

Figure 4.38 Authenticate Drive Security Key



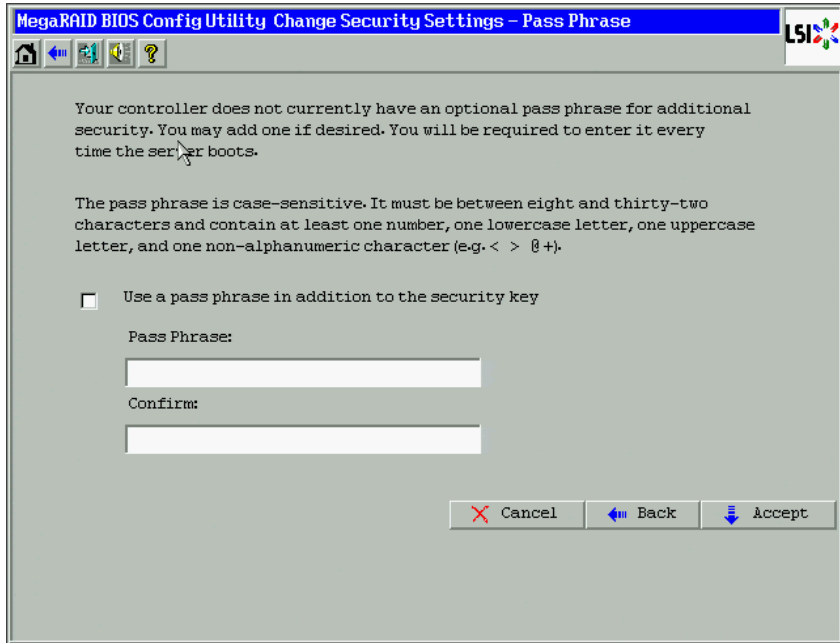
9. Enter the current security key and click **OK**.

The text box for the security key can hold up to 32 characters. The key must be at least eight characters. After you enter the correct security key, the wizard returns to the Change Security Settings – Security Key screen.

10. Click **Next**.

The Change Security Settings – Pass Phrase screen appears as shown in the following figure. You have the option to provide a pass phrase for additional security.

Figure 4.39 Change Security Settings – Pass Phrase



MegaRAID BIOS Config Utility Change Security Settings – Pass Phrase

Your controller does not currently have an optional pass phrase for additional security. You may add one if desired. You will be required to enter it every time the server boots.

The pass phrase is case-sensitive. It must be between eight and thirty-two characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (e.g. < > @ +).

☐ Use a pass phrase in addition to the security key

Pass Phrase:

Confirm:

Cancel Back Accept

11. If you want to use a pass phrase, click the checkbox **Use a pass phrase in addition to the security key**.
12. Enter a new pass phrase and then enter the new pass phrase again to confirm.

The pass phrase is case-sensitive. It must be between eight and thirty-two characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (for example, < > @ +). The space character is not permitted.

Non-US keyboard users must be careful not to enter DBCS characters in the pass phrase field or security key field. Firmware works only with the ASCII character set.

13. Click **Accept**.

If you entered a new a pass phrase, the Authenticate Pass Phrase screen appears.

14. On the Authenticate Pass Phrase screen, enter the pass phrase and click **Finish**.

The Confirm Change Drive Security Settings screen appears as shown in the following figure. This screen lists the changes you made and asks you whether you want to confirm these changes.

Figure 4.40 Confirm Change Drive Security Settings



15. Click **Yes** on the Confirm Change Drive Security Settings screen, confirm that you want to change the drive security settings.

If the current security key is not needed, WebBIOS saves the changes to the security settings and returns you to the main menu. If the current security key is needed, the Authenticate Drive Security Settings screen displays.

4.5.3 Disabling the Drive Security Settings


Perform the following steps to disable the drive security settings.

Note: If you disable the drive security settings, you cannot create any new secure virtual drives. Disabling these settings does not affect the security or data of foreign drives. If you removed any drives that were previously secured, you will still need to enter the security key when you import them.

1. Click **Drive Security** on the main WebBIOS screen.

The Encryption Settings screen appears as shown in the following figure.

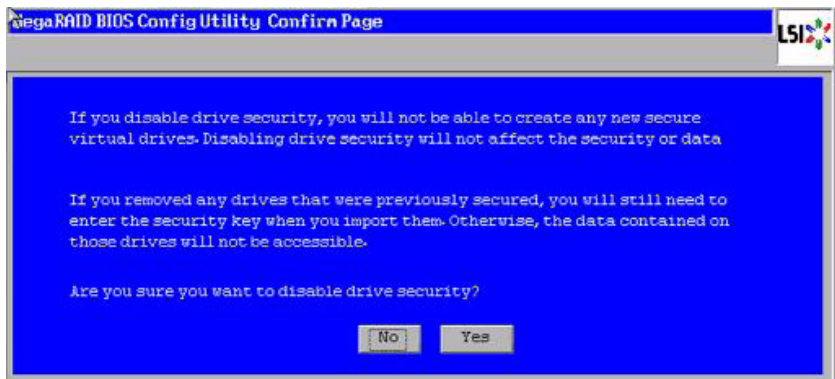
Figure 4.41 Encryption Settings



2. To disable the drive security settings, select **Disable drive security** and click **Accept**.

The Confirm Disable Drive Security screen appears as shown in the following figure.

Figure 4.42 Confirm Disable Drive Security Settings



3. On the Confirm Disable Security Settings screen, click **Yes** to confirm that you want to disable the drive security settings.

WebBIOS returns you to the MSM main menu.

4.5.4 Importing Foreign Configurations

After you create a security key, you can run a scan for a foreign configuration and import a locked configuration. (You can import unsecured or unlocked configurations when security is disabled.) A foreign configuration is a RAID configuration that already exists on a replacement set of drives that you install in a computer system. You can use the WebBIOS utility to import the existing configuration to the RAID controller or clear the configuration so you can create a new one.

See [Section 4.12.3, “Importing or Clearing a Foreign Configuration”](#) for the procedures used to import or clear a foreign configuration.

To import a foreign configuration, you must first enable security to allow importation of locked foreign drives. If the drives are locked and the controller security is disabled, you cannot import the foreign drives. Only unlocked drives can be imported when security is disabled.

After you enable the security, you can import the locked drives. To import the locked drives, you must provide the security key used to secure them. Verify whether any drives are left to import as the locked drives can use different security keys. If there are any drives left, repeat the import process for the remaining drives. After all of the drives are imported, there is no configuration to import.

4.6 Viewing and Changing Device Properties


This section explains how you can use the WebBIOS CU to view and change the properties for controllers, virtual drives, drives, and BBUs.

4.6.1 Viewing and Changing Controller Properties

WebBIOS displays information for one IBM SAS controller at a time. If your computer system has multiple IBM SAS controllers, you can view information for a different controller by clicking **Controller Selection** on the main screen. When the Controller Selection screen appears, select the controller you want from the list.

To view the properties for the currently selected controller, click **Controller Properties** on the main WebBIOS screen. There are three Controller Properties screens. the following figure shows the first screen.

Figure 4.43 First Controller Properties Screen



The information on this screen is read-only and cannot be modified directly. Most of this information is self-explanatory. The screen lists the number of virtual drives that are already defined on this controller, and the number of drives connected to the controller.

Click **Next** to view the second Controller Properties screen, as shown in the following figure.

Figure 4.44 Second Controller Properties Screen



Click **Next** to view the third Controller Properties screen, as shown in the following figure.

Figure 4.45 Third Controller Properties Screen

MegaRAID BIOS Config Utility Controller Properties			
Properties			
Battery Backup	Present	Coercion Mode	None
Set Factory Defaults	No	S.M.A.R.T Polling	300 seconds
Cluster Mode	Disabled	Alarm Control	Disabled
Rebuild Rate	30	Patrol Read Rate	30
BGI Rate	30	Cache Flush Interval	4
CC Rate	30	Spinup Drive Count	4
Reconstruction Rate	30	Spinup Delay	2
NCQ	Enabled		
<input type="button" value="Submit"/> <input type="button" value="Reset"/> <input type="button" value="Next"/>			
<input type="button" value="Home"/>		<input type="button" value="Back"/>	

Click **Next** to view the fourth Controller Properties screen, as shown in the following figure.

Figure 4.46 Fourth Controller Properties Screen

MegaRAID BIOS Config Utility Controller Properties			
Properties			
Stop CC On Error	No	Schedule CC	Supported
Maintain PD Fail History	Enabled	StopOnError	Enabled
Controller BIOS	Enabled	Disk Activity	Disabled
Snapshot	Create	Manage JBOD	Manage
Manage Powersave	Settings	Link Speed	Manage
Global Hotspare for Emergency	Enabled	Unconfigured Good for Emergency	Enabled
Emergency for SMARTer	Disabled		
<input type="button" value="Submit"/> <input type="button" value="Reset"/>			
<input type="button" value="Home"/>		<input type="button" value="Back"/>	

Table 4.2 describes the entries/options listed on the second, third, and fourth Controller Properties screens. Leave these options at their default

settings to achieve the best performance, unless you have a specific reason for changing them.

Table 4.2 Controller Properties Menu Options

Option	Description
Battery Backup	This entry indicates whether the selected controller has a BBU. If present, you can click <i>Present</i> to view information about the BBU. For more information, see Section 4.6.5, “Viewing and Changing Battery Backup Unit Information.”
Set Factory Defaults	Use this option to load the default MegaRAID® WebBIOS CU settings. The default is <i>No</i> .
Cluster Mode	Use this option to enable or disable Cluster mode. The default is <i>Disabled</i> . A cluster is a grouping of independent servers that can access the same data storage and provide services to a common set of clients. When Cluster mode is disabled, the system operates in Standard mode.
Rebuild Rate	Use this option to select the rebuild rate for drives connected to the selected controller. The default is 30 percent. The rebuild rate is the percentage of system resources dedicated to rebuilding a failed drive. The higher the number, the more system resources are devoted to a rebuild.
BGI Rate	Use this option to select the amount of system resources dedicated to background initialization of virtual drives connected to the selected controller. The default is 30 percent.
CC Rate	Use this option to select the amount of system resources dedicated to consistency checks of virtual drives connected to the selected controller. The default is 30 percent.
Reconstruction Rate	Use this option to select the amount of system resources dedicated to reconstruction of drives connected to the selected controller. The default is 30 percent.
Controller BIOS	Use this option to enable or disable the BIOS for the selected controller. The default is <i>Enabled</i> . If the boot device is on the selected controller, the BIOS must be enabled; otherwise, the BIOS should be disabled or it might not be possible to use a boot device elsewhere.
NCQ	Native Command Queuing (NCQ) gives an individual drive the ability to optimize the order in which it executes the read and write commands. The default is <i>Enabled</i> .
Coercion Mode	<p>Drive coercion is a tool for forcing drives of varying capacities to the same size so they can be used in a drive group. The coercion mode options are <i>None</i>, <i>128MB-way</i>, and <i>1GB-way</i>. The default is <i>None</i>.</p> <p>Note: The number you choose depends on how much the drives from various vendors vary in their actual size. Use the 1GB coercion mode option.</p>

Table 4.2 Controller Properties Menu Options (Cont.)

Option	Description
S.M.A.R.T. Polling	Use this option to determine how frequently the controller polls for drives reporting a Predictive Drive Failure (S.M.A.R.T.: Self-Monitoring Analysis and Reporting Technology error). The default is 300 seconds (5 minutes).
Alarm Control	Select this option to enable, disable, or silence the onboard alarm tone generator on the controller. The default is <i>Enabled</i> .
Patrol Read Rate	Use this option to select the rate for patrol reads for drives connected to the selected controller. The default is 30 percent. The patrol read rate is the percentage of system resources dedicated to running a patrol read.
Cache Flush Interval	Use this option to control the interval (in seconds) at which the contents of the onboard data cache are flushed. The default is 4 seconds.
Spinup Drive Count	Use this option to control the number of drives that spin up simultaneously. The default is 4 drives.
Spinup Delay	Use this option to control the interval (in seconds) between spinup of drives connected to this controller. The delay prevents a drain on the system's power supply that would occur if all drives spun up at the same time. The default is 12 seconds.
StopOnError	Enable this option if you want the boot process to stop when the controller BIOS encounters an error during boot-up. The default is <i>Enabled</i> .
Stop CC on Error	Enable this option if you want to stop a consistency check when the controller BIOS encounters an error. The default is <i>No</i> .
Maintain PD Fail History	Enable this option to maintain the history of all drive failures. The default is <i>Enabled</i> .
Schedule CC	Indicates whether the option to schedule the date and time for a consistency check is supported.
Snapshot	Use this option to create a snapshot of a volume. MegaRAID Recovery, also known as Snapshot, offers a simplified way to recover data and provides automatic protection for the boot volume. You can use the Recovery feature to take a snapshot of a volume and to restore a volume or file. Snapshot functionality allows you to capture data changes to the volume, and, if data is deleted accidentally or maliciously, restore the data from the view or roll back to a snapshot at a previous point-in-time (PiT). MegaRAID Recovery supports up to eight snapshots of PiTs for each volume.
Disk Activity	Enable this property if you want to locate a particular disk. This disk can be identified with a continuous blinking of green activity LED. This works only if the disks are installed in an enclosure.

Table 4.2 Controller Properties Menu Options (Cont.)

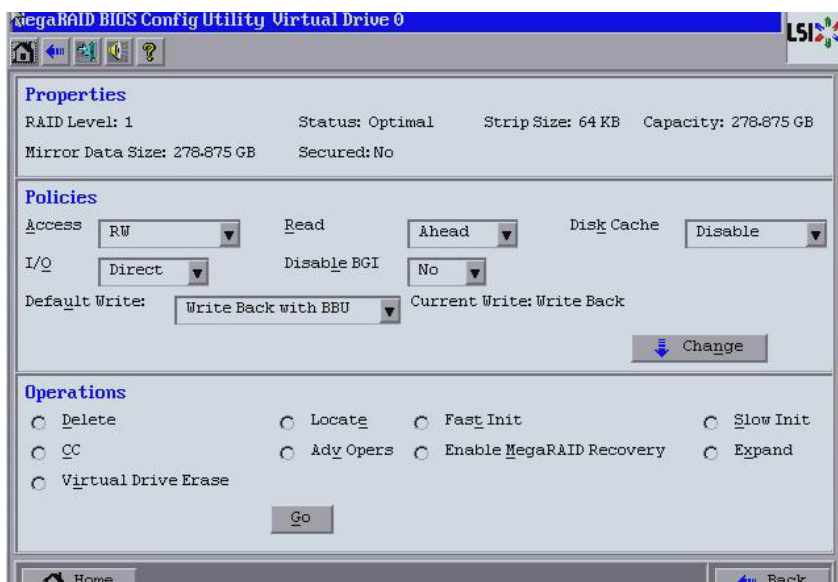
Option	Description
Manage JBOD	Converting the multiple JBOD drives to unconfigured drive at single selection.
Manage Powersave	Use this option to reduce the power consumption of drives that are not in use, by spinning down the unconfigured drives, hot spares, and configured drives.
Link Speed	Use this option to change the link speed between the controller and an expander or between the controller and a drive that is directly connected to the controller.

If you make changes to the options on this screen, click **Submit** to register them. If you change your mind, click **Reset** to return the options to their default values.

4.6.2 Viewing and Changing Virtual Drive Properties

Access the Virtual Drive screen by clicking on a virtual drive in the list of virtual drives in the right panel on the WebBIOS CU main screen. The Virtual Drive screen displays, as shown in the following figure.

Figure 4.47 Virtual Drive Screen



The Properties panel of this screen displays the virtual drive's RAID level, state, size, strip size, and metadata size.

The Policies panel lists the virtual drive policies that were defined when the storage configuration was created. For information about these policies, see [Section 4.4.3, "Using Manual Configuration."](#) To change any of these policies, make a selection from the drop-down menu and click **Change**.

The Operations panel lists operations that can be performed on the virtual drive. To perform an operation, select it and click **Go**. Then choose from the following options:

- Select **Delete** to delete this virtual drive. For more information, see [Section 4.12.2, "Deleting a Virtual Drive."](#)
- Select **Locate** to make the LEDs flash on the drives used by this virtual drive. This works only if the drives are installed in a drive enclosure that supports SAFTE (SCSI-Accessed-Fault-Tolerant-Enclosure).
- Select **Fast Init** or **Slow Init** to initialize this virtual drive. A fast initialization quickly writes zeroes to the first and last 10-Mbyte regions of the new virtual drive and then completes the initialization in the background. A slow initialization is not complete until the entire virtual drive has been initialized with zeroes. It is seldom necessary to use this option, because the virtual drive was already initialized when you created it.

Attention: Before you run an initialization, back up any data on the virtual drive that you want to save. All data on the virtual drive is lost when you initialize it.

- Select **CC** to run a consistency check on this virtual drive. For more information, see [Section 4.12.1, "Running a Consistency Check."](#) (This option is not available for RAID 0 virtual drives.)
- Select **AdvOps** to access screens to remove drives and migrate RAID levels (that is, change the virtual drive configuration by adding a drive and changing the RAID level). For more information, see [Section 4.12.4, "Migrating the RAID Level of a Virtual Drive."](#)
- Select **Expand** to increase the size of a virtual drive to occupy the remaining capacity in the drive group. In addition, you can add drives to the virtual drive in order to increase capacity. For more

information, see [Section 4.7, “Viewing and Expanding a Virtual Drive.”](#)

- Select **Virtual Drive Erase** to erase a virtual drive and over write all users. For more information, see [Section 4.10.2, “Virtual Drive Erase.”](#)

Attention: Before you change a virtual drive configuration, back up any data on the virtual drive that you want to save.

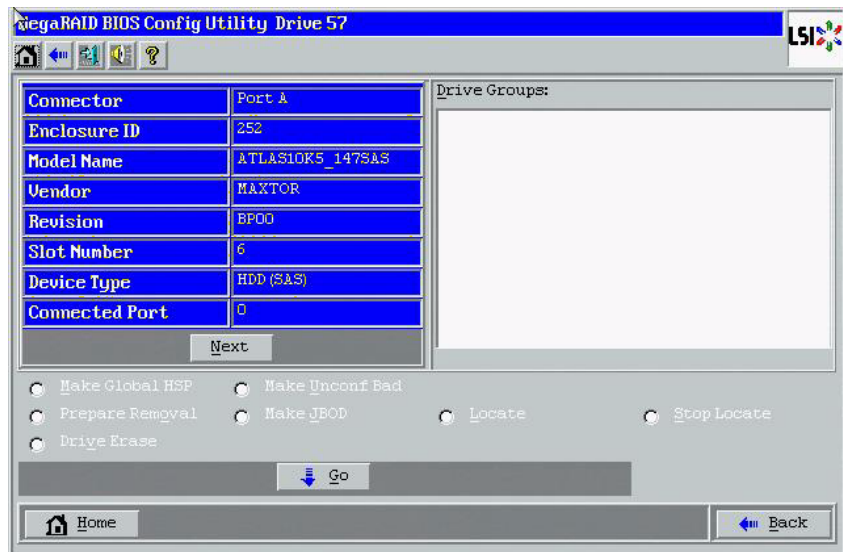
4.6.3 Viewing Drive Properties

The Physical Drive screen displays the properties of a selected drive and enables you to perform operations on the drive. There are two ways to access the Physical Drive screen:

- On the main menu screen, click on a drive in the right panel under the heading **Physical Drives**.
- On the main menu screen, click on **Drives** in the left panel to display the Drive screen. Then click on a drive in the right panel. Click on the **Properties** button, and click **Go**. The properties for the selected drive appear.

The following figure shows the Physical Drive screen.

Figure 4.48 Physical Drive Screen



The drive properties are view-only and are self-explanatory. Note that the properties include the state of the drive.

Operations you can perform are listed at the bottom of the screen. After you select an operation, click **Go** to start the operation. The operations vary depending on the drive state. If the drive state is **Online**, the following operations appear:

- Select **MakeDriveOffline** if you want to force the drive offline.

Note: If you force offline a good drive that is part of a redundant drive group with a hot spare, the drive rebuilds to the hot spare drive. The drive you forced offline goes into the *Unconfigured Bad* state. Access the BIOS utility to set the drive to the *Unconfigured Good* state.

- Select **Locate** to make the LED flash on the drive. This action works only if the drive is installed in a drive enclosure.

If the drive state is *Unconfigured Good*, four additional operations appear on this screen:

- Select **Make Global HSP** to make a global hot spare, available to all of the virtual drives.

- Select **Make Dedicated HSP** to make a hot spare dedicated to a specific virtual drive.

WebBIOS displays the global hot spare as `Global` and the dedicated hot spare as `Ded`. The icon for the dedicated hot spare displays under its associated virtual drive. The drive number, drive state, drive capacity, and drive manufacturer display.

- Select **Enclosure Affinity** so if there are drive failures present on a split backplane configuration, then the hot spare will be used first on the backplane side in which it resides.
- Select **Prepare for Removal** to prepare the drive for removal from the enclosure.

The **Prepare for Removal** feature is different from spinning a drive down into powersave mode because it also involves flagging the drive as ready to remove. If you choose to prepare a drive for removal, **Ready to Remove** displays in the device tree for that drive, instead of **Powersave**.

- Select **Stop Locate** to stop the LED flash on the drive. This works only if the drive is installed in a drive enclosure.
- Select **Drive Erase** to securely erase data on non self-encrypting drives (Non-SED), which are normal HDDs.

4.6.4 Shield State

Physical devices in MegaRAID firmware transit between different states. If the firmware detects a problem or a communication loss for a physical drive, the firmware transitions the drive to a bad (FAILED or UNCONF BAD) state. To avoid transient failures, an interim state called the *shield state* is introduced before marking the drive as being in a bad state.

The shield state is an interim state of a physical drive for diagnostic operations. The results of the diagnostic tests determine if the physical drive is good or bad. If any of the diagnostic tests fail, the physical drive transitions to a bad state (FAILED or UNCONF BAD).

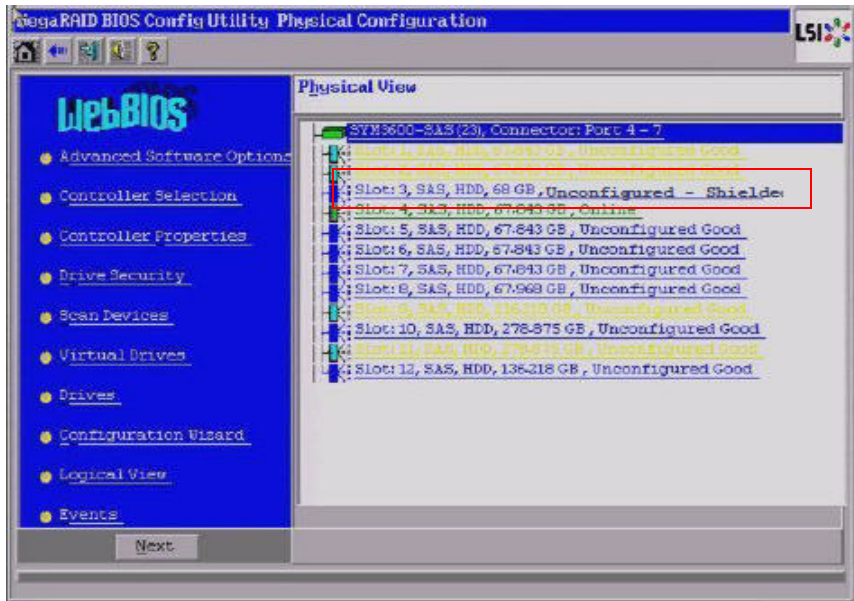
4.6.4.1 Shield State Physical View

Follow these steps to check if a physical drive is in a Shield state in the Physical view.

- Click **Physical View** in the main dialog.

The physical drive that is in a shield state is marked as Shielded. the following figure shows the Physical view shield state.

Figure 4.49 Physical View Shield State Dialog




4.6.4.2 Logical View Shield State

Follow these steps to view the Shield state in the Logical view.

- Click **Logical View** in the main page.

The physical drive that is in a shield state is marked as Shielded. The following figure shows the Logical view shield state.

Figure 4.50 Logical View Shield State




4.6.4.3 Viewing the Physical Drive Properties of a Drive In Shield State

Follow these steps to view the physical properties of the drive in Shield state.

1. Click on the **Physical view** tab or the **Logical view** tab in the device tree.
2. Click the physical drive that is in shield state on the physical or logical view of device tree to view the properties. [Figure 4.51](#) shows the device properties of the drive that is in shield state.

Figure 4.51 Physical Drive Properties of a Drive in Shield State




4.6.4.4 Determining if Shield State is Enabled in a Controller

Follow these steps to determine if the Shield state is enabled in a controller.

1. Click **Controller Properties** on the WebBIOS main menu.
2. The Shield State Supported Column is displayed, as shown in [Figure 4.52](#).

Figure 4.52 Shield State Support



4.6.5 Viewing and Changing Battery Backup Unit Information

If your SAS controller has a battery backup unit (BBU), you can view information about the BBU and change some settings. To do this, follow these steps:

1. Click **Controller Properties** on the WebBIOS CU main menu screen.

The first **Controller Properties** screen appears, as shown in [Figure 4.53](#).

Figure 4.53 First Controller Properties Screen



2. Click **Next**.

The second Controller Properties screen appears.

3. Click **Next**.

The third Controller Properties screen appears, as shown in [Figure 4.54](#). The **Battery Backup** field at the top left of the screen indicates whether the iBBU unit is present.

Figure 4.54 Third Controller Properties Screen

MegaRAID BIOS Config Utility Controller Properties			
Properties			
Battery Backup	Present	Coercion Mode	None
Set Factory Defaults	No	S.M.A.R.T Polling	300 seconds
Cluster Mode	Disabled	Alarm Control	Enabled
Rebuild Rate	30	Patrol Read Rate	30
BGI Rate	30	Cache Flush Interval	4
CC Rate	30	Spinup Drive Count	2
Reconstruction Rate	30	Spinup Delay	12
NCQ	Enabled		
<input type="button" value="Submit"/> <input type="button" value="Reset"/> <input type="button" value="Next"/>			
<input type="button" value="Home"/>		<input type="button" value="Back"/>	

4. Click **Present** in the **Battery Backup** field.

The Battery Module screen appears, as shown in the following screen. The screen contains the following information:

- Battery information
- Design information
- Capacity information
- Properties information

Figure 4.55 Battery Module Screen: iBBU08 Battery

MegaRAID BIOS Config Utility Battery Module	
Battery Type: iBBU08 Voltage: 4068 mV Current: 0 mA Temperature: 30 deg. centigrade Status: gas Gauge Status : Initialized Full Charge Capacity remaining : 98% Design Charge Capacity remaining : 89% expected margin of error : 0%	Design Info Mfg. Name: LS36681 Mfg. Date: 7/9/2010 Serial No.: 37 FRU: None Design Capacity: 1530 mAh Design Voltage: 4100 mV Device Name: bq27541 Device Chemistry: LPMR
Capacity Info FullCharge Capacity: 1416 mAh Remaining Capacity: 1374 mAh	Properties Bbu Mode 4-48 Hour, 3 Years, 45 <input type="button" value="Go"/>

Home Back

Most Battery Module properties are view-only and are self-explanatory.

– Selecting the BBU Mode

You can select the BBU mode to set the length of time that the BBU retains the data in the event of a power failure or outage. The following table describes the BBU modes that display in the drop-down menu in the lower right section of the Battery Module screen.

Table 4.3 BBU Modes

Mode of Operation	Description
1	12 hours retention at 45 °C, transparent learn
3	24 hours retention at 45 °C, transparent learn This mode is the default for the ServeRAID M5100 Battery Kit.
4	48 hours retention at 45 °C, non transparent learn. This mode is the default for the ServeRAID M5000 Battery Kit.

Perform the following steps to select the BBU mode:

1. Open the drop-down menu in the **BBU Mode** field in the Properties section of the screen.
2. Select the mode you want from the drop-down menu.
Mode #4 is the default.
3. Click **Go**.
This action sets the BBU mode.

– **Setting the Learn Delay Interval**

A *learning cycle* is a battery calibration operation performed by the controller periodically to determine the condition of the battery. You can change the learn delay interval, which is the length of time between automatic learning cycles.

Note: Leave the the learn delay interval and the auto learn mode at their default settings.

Perform the following steps to change the interval:

1. Open the drop-down menu in the **Auto Learn Mode** field in the Properties section of the screen.
2. Select the learn mode as *Enabled*.
Use this setting so the controller performs the learning cycle automatically.
3. Change the number of hours in the **Delay Scheduled learn cycle by** field.
You can delay the start of the learn cycles for up to 168 hours (7 days).
4. Click **Go**.
This action sets the length of time between automatic learning cycles.

– Setting the Auto Learn Mode

You can start battery learning cycles manually or automatically. The Auto Learn modes are:

- BBU Auto Learn: Firmware tracks the time since the last learning cycle and performs a learn cycle when it is due.
- BBU Auto Learn Disabled: Firmware does not monitor or initiate a learning cycle. You can schedule learning cycles manually.

If you disable the Auto Learn Mode and re-enable it later, the controller firmware resets the battery module properties to initiate an immediate battery learn cycle. However, in the **Scheduled Learn cycle** field, the value **None** is displayed. The **Scheduled Learn cycle** field is not updated until the battery relearn is completed.

When the relearning cycle is completed, the value in the **Scheduled Learn cycle** field displays the new date and time of the next battery learning cycle.

- BBU Auto Learn Warn: Firmware warns about a pending learning cycle. You can initiate a learning cycle manually. After the learning cycle is complete, firmware resets the counter and warns you when the next learning cycle time is reached.

Perform the following steps to choose an auto learn mode:

1. Open the drop-down menu in the **Auto Learn Mode** field.
2. Select an auto learn mode.
3. Click **Go** to set the auto learn mode.

Note: When you replace the iBBU, the charge cycle counter is reset automatically.

4.6.6 Managing Link Speed

The Managing Link Speed feature allows you to change the link speed between the controller and an expander or between the controller and a drive that is directly connected to the controller.

All phys in a SAS port can have different link speeds or can have the same link speed.

You can select a link speed setting. However, if phys in a SAS port have different link speed settings and if a phy is connected to a drive or an

expander, the firmware overrides the link speed setting you have selected and instead uses the common maximum link speed among all the phys.

You change the link speed on the fourth Controller Properties screen. To access this screen and change the link speed, perform the following steps:

1. Click **Controller Properties** on the WebBIOS main menu.
The first Controller Properties screen appears.
2. Click **Next** to access the second Controller Properties screen.
The second Controller Properties screen appears.
3. Click **Next** to access the third Controller Properties screen.
The third Controller Properties screen appears.
4. Click **Next** to access the fourth Controller Properties screen.
The fourth Controller Properties screen appears, as shown in the following figure.

Figure 4.56 Fourth Controller Properties Screen

Properties			
Stop CC On Error	No ▼	Schedule CC	Supported
Maintain PD Fail History	Enabled ▼	StopOnError	Enabled ▼
Controller BIOS	Enabled ▼	Disk Activity	Disabled ▼
Snapshot	Create	Manage JBOD	Manage
Manage Powersave	Settings	Link Speed	Manage
Global Hotspare for Emergency	Enabled ▼	Unconfigured Good for Emergency	Enabled ▼
Emergency for SMARTer	Disabled ▼		

Submit Reset

Home Back

1. Click **Manage** in the **Link Speed** field.
The Manage Link Speed dialog box appears, as shown in the following figure.

Figure 4.57 Manage Link Speed Screen

MegaRAID BIOS Config Utility Manage Link Speed

View the phy setting for the controller and change them if needed.

Phy Details:

SAS Address	Phy	Select Link Speed
0000000000000000	0	1.5Gbps ▼
0000000000000000	1	3Gbps ▼
0000000000000000	2	1.5Gbps ▼
0000000000000000	3	Auto ▼
500015560006a23f	4	1.5Gbps ▼
500015560006a23f	5	3Gbps ▼
500015560006a23f	6	1.5Gbps ▼
500015560006a23f	7	1.5Gbps ▼

Ok Cancel

- The SAS Address column displays the SAS address that uniquely identifies a device in the SAS domain.
 - The Phy column displays the system-supported phy link values, which are from 0 through 7.
 - The Select Link Speed column displays the phy link speeds.
2. Select the desired link speed from the **Select Link Speed** field using the drop-down selector.


The link speed values are Auto, 1.5, 3.0 or 6.0 Gbps.

Note: By default, the link speed in the controller is *Auto* or the value last saved by the user.

3. Click **OK**.

The link speed value is now reset. The change takes place after you restart the system. The message box appears, as shown in the following figure.

Figure 4.58 System Restart Required Message Box



4.6.7 Viewing Enclosure Properties

Using WebBIOS, you can view the enclosure properties of all of the enclosures connected to the server.

Follow these steps to view enclosure properties.

1. Go to the Physical view of the WebBIOS Utility.
2. Click the enclosure node.

The enclosure properties appear, as shown in the following figure.

Figure 4.59 Enclosure Properties



3. Click **Next** to view additional properties, as shown in the following figure.

Figure 4.60 Additional Enclosure Properties



- Click **More Info** to view additional information on the number of temperature sensors (Figure 4.61), number of fans (Figure 4.62), and the number of power supplies (Figure 4.63).

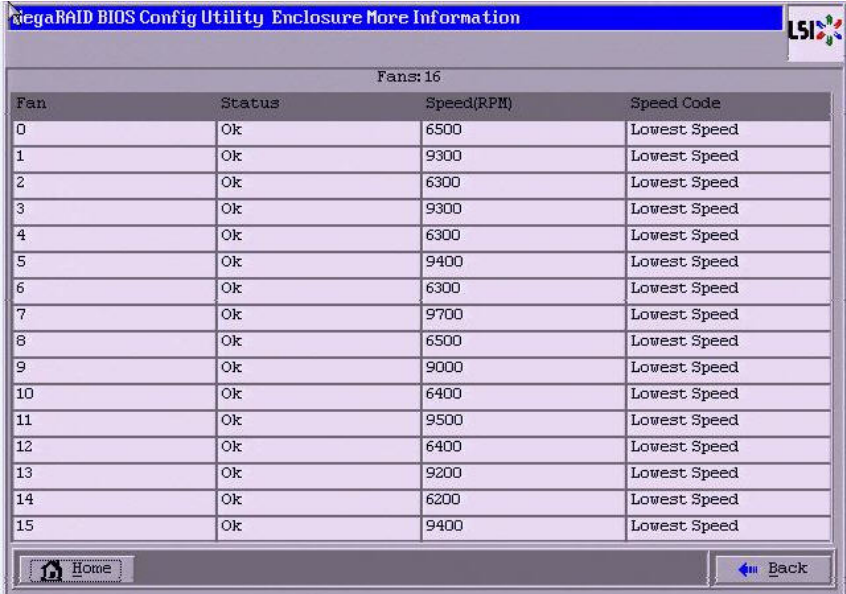
Figure 4.61 Enclosure More Information - Temperature Sensors

The screenshot shows the 'MegaRAID BIOS Config Utility Enclosure More Information' window. It displays a table of 8 temperature sensors with their status and temperature in Celsius. Navigation buttons 'Home' and 'Back' are at the bottom.

Temperature Sensors: 8		
Temperature Sensors	Status	Temperature(Celcius)
0	Ok	41
1	Ok	42
2	Ok	49
3	Ok	39
4	Not Installed	0
5	Ok	42
6	Ok	41
7	Ok	20


Home Back

Figure 4.62 Enclosure More Information - Number of Fans



Fan	Status	Speed(RPM)	Speed Code
0	Ok	6500	Lowest Speed
1	Ok	9300	Lowest Speed
2	Ok	6300	Lowest Speed
3	Ok	9300	Lowest Speed
4	Ok	6300	Lowest Speed
5	Ok	9400	Lowest Speed
6	Ok	6300	Lowest Speed
7	Ok	9700	Lowest Speed
8	Ok	6500	Lowest Speed
9	Ok	9000	Lowest Speed
10	Ok	6400	Lowest Speed
11	Ok	9500	Lowest Speed
12	Ok	6400	Lowest Speed
13	Ok	9200	Lowest Speed
14	Ok	6200	Lowest Speed
15	Ok	9400	Lowest Speed

Figure 4.63 Enclosure More Information - Number of Power Supplies



Power Supply	Status
0	Ok
1	Critical

4.6.8 SSD Disk Cache Policy

MegaRAID supports changes to the write-cache policy for SSD media of individual physical drives.

When SSDs are configured in a mixed disk group with HDDs, the **Physical Device Write-Cache Policy** setting of all of the participating drives is changed to match the SSD cache policy setting.

Follow these steps to view the SSD Disk Cache Setting property.

1. Click the controller properties link in the main menu.
2. Click **Next** to view the controller properties with **SSD Disk Cache Setting** displayed, as shown in the following figure.

Figure 4.64 SSD Disk Cache Setting in Controller Properties Dialog



4.7 Viewing and Expanding a Virtual Drive

Follow these steps to view virtual drive properties.

1. In the Logical view of the device tree, click the **Virtual Drive Node**.
The virtual drive properties appear, as shown in the following figure.

Figure 4.65 Virtual Drive Properties

MegaRAID BIOS Config Utility Virtual Drive 0

Properties

RAID Level: 0 Status: Optimal Strip Size: 64 KB Capacity: 3 GB
SSD Caching: Enabled Cached: No Secured: No

Policies

Access: RW Read: Normal Disk Cache: NoChange
I/O: Cached Disable BGI: No
Default Write: Write Back with BBU Current Write: Write Through
Reason for Difference in Write Policy: BBU is discharged [Change](#)

Operations

☐ Delete ☐ Locate ☐ Fast Init ☐ Slow Init
☐ Adv Opers ☐ Expand
☐ Disable SSD Caching

[Go](#)

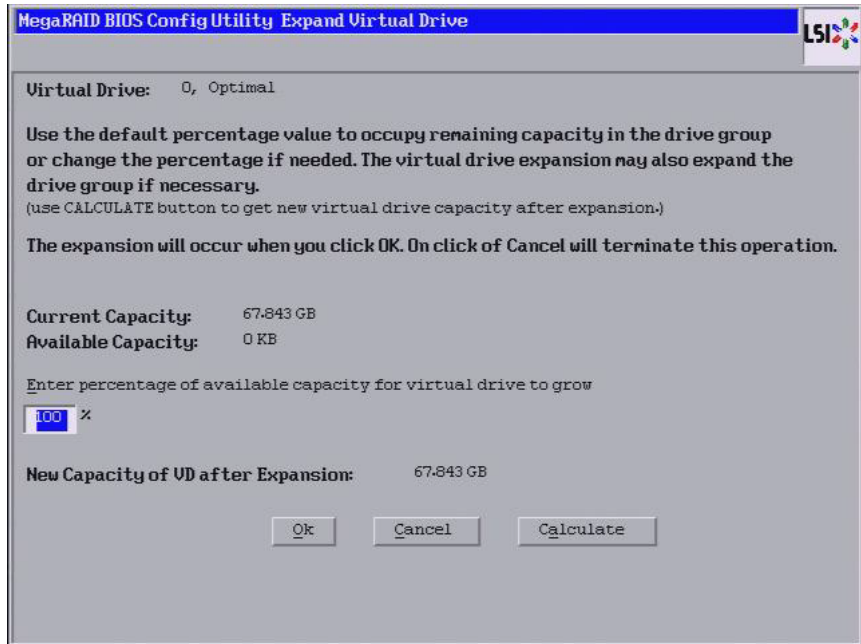
[Home](#) [Back](#)

You can increase the size of a virtual drive to occupy the remaining capacity in a drive group.

2. Select the **Expand** radio button, and click **Go**.

The WebBIOS Config Utility Expand Virtual Drive dialog appears, as shown in the following figure.

Figure 4.66 Expand Virtual Drive Dialog



3. Enter the percentage of the available capacity that you want the virtual drive to use.

For example, if there are 100 GB of capacity available and you want to increase the size of the virtual drive by 30 GB, select 30 percent.
4. Click **Calculate** to determine the capacity of the virtual drive after expansion.
5. Click **Ok**.

The virtual drive expands by the selected percentage of the available capacity.

4.8 Recovering/Clearing Punctured Block Entries

You can recover/clear the punctured block area of a virtual drive.

Caution: This operation removes any data stored on the physical drives. Back up the good data on the drives before making any changes to the configuration.

When a Patrol Read or a Rebuild operation encounters a media error on the source drive, it punctures a block on the target drive to prevent the use of the data with the invalid parity. Any subsequent read operation to the punctured block completes but with an error. Consequently, the puncturing of a block prevents any invalid parity generation later while using this block.

To recover or clear the punctured block area of a virtual drive, run a Slow (or Full) Initialization to zero out and regenerate new parity, causing all bad block entries to be removed from the bad block table. To run a Slow (or Full) Initialization, see [Section 4.6.2, “Viewing and Changing Virtual Drive Properties”](#).

4.9 Suspending and Resuming Virtual Drive Operations

MegaRAID provides background Suspend and Resume features that enhances the functionality. The background operations on a virtual drive can be suspended using the **Suspend** option, and later resumed using the **Resume** option. The suspended operation resumes from the point where the operation was suspended.

If any operation is stopped before completion, it is considered to be aborted. An aborted operation cannot be resumed from the place it was stopped.

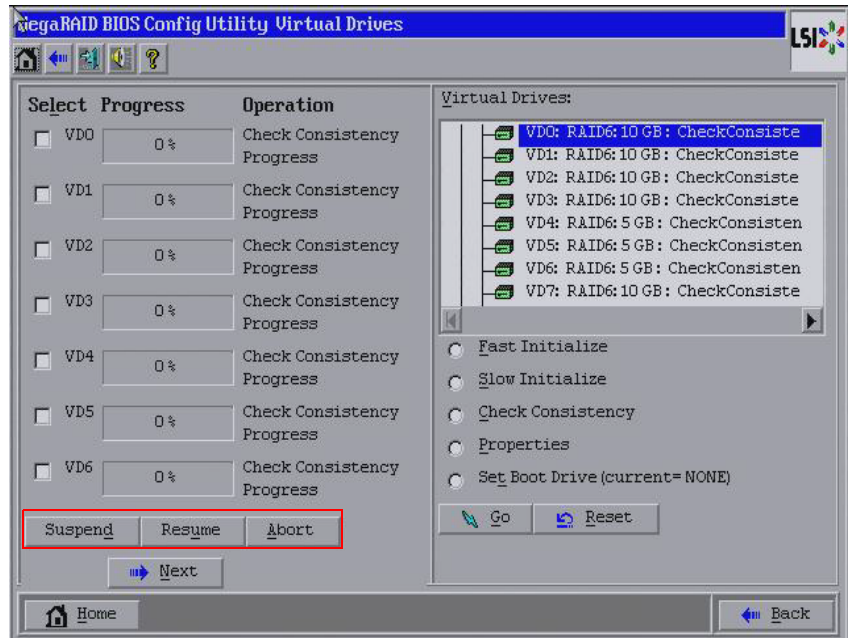
Note: Suspend and resume are applicable for all the background operations, such as background initialization, rebuild and consistency check notes.

Follow these steps to suspend an operation and resume an operation.

1. On the WebBIOS main menu, click **Virtual Drives** link, or click the **VD Progress Info** button on the taskbar.

The Virtual Drives main dialog appears, as shown in the following figure.

Figure 4.67 Virtual Drives Dialog



2. To suspend operations, select the check boxes for the operations that you want to suspend, and click **Suspend** (ALT+D).
3. To abort operations, select the check boxes for the operations and click **Abort** (ALT+A). Aborted operations cannot be resumed and have to be started again.
4. To resume operations, select the check boxes for the suspended operations that you want to resume, and click **Resume** (ALT+U).

4.10 Non-SED Secure Erase

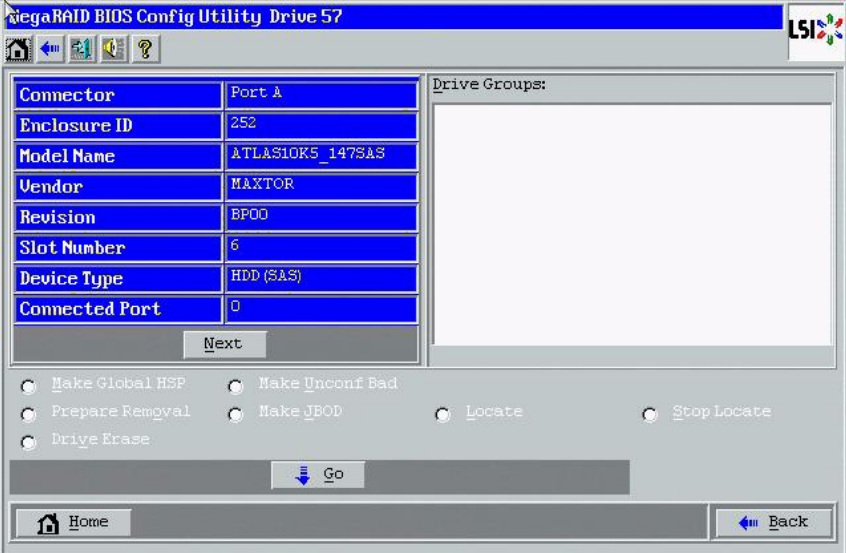
This section describes the procedure used to securely erase data on non self-encrypting drives (Non-SED), which are normal HDDs.

4.10.1 Erasing a Non-SED Physical Drive

Follow these steps for non-SED secure erase.

1. Go to the Physical view in the WebBIOS main menu.
2. Click the physical drive node.
3. Select the **Drive Erase** radio button, as shown in the following figure, and click on **Go**.

Figure 4.68 Physical Drive Dialog



Connector	Port A
Enclosure ID	252
Model Name	ATLAS10K5_147SAS
Vendor	MAXTOR
Revision	BPO0
Slot Number	6
Device Type	HDD (SAS)
Connected Port	0

Next

Drive Groups:

☐ Make Global HSP ☐ Make Unconf Bad ☐ Locate ☐ Stop Locate

☐ Prepare Removal ☐ Make JBOD

☐ Drive Erase


Go

Home

Back

The Mode Selection- Drive Erase dialog appears.

Figure 4.69 Mode Selection - Drive Erase



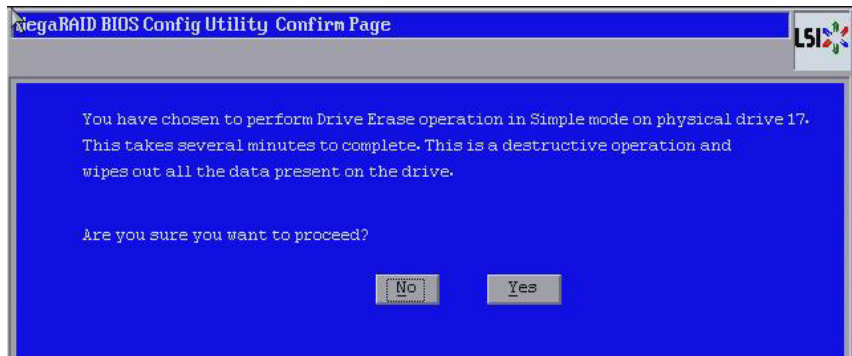
4. Select any of the modes available under the **Select the mode for Drive Erase Operation**

- **Simple** – (Alt + S)
- **Normal** – (Alt + N)
- **Thorough** – (Alt + T)

5. Click **OK**.

A confirmation message dialog appears as shown in the following figure.

Figure 4.70 Drive Erase Confirm Page



4.10.1.1 Drive Erase Progress

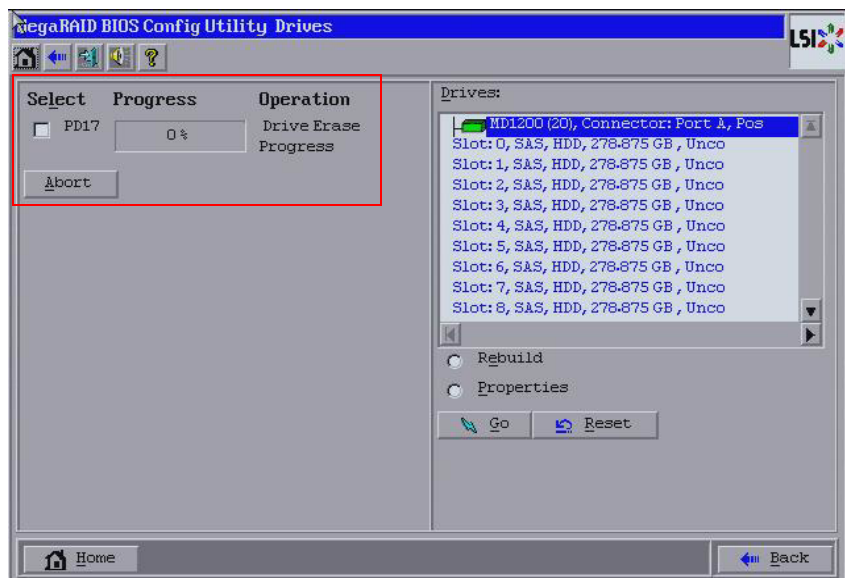
Physical drives, erase operation is generally a time-consuming operation and is performed as a background task.

Follow these steps to check the progress of a physical drive erase operation.

1. Click the **Drives** link in the left panel on the WebBIOS main dialog.

The **Drive Erase Progress** appears, as shown in the following figure.

Figure 4.71 Drive Erase Progress



2. To abort drive erase, select the check box for the operation that you want to abort and click **Abort**.


4.10.2 Virtual Drive Erase

Virtual drive erase is a background operation.

Follow these steps to perform the **Virtual Drive Erase** operation.

1. Go to the **Logical view**.
2. Click on the Virtual Drive node.
3. Select the **Virtual Drive Erase** radio button and click **Go** as shown in the following figure.

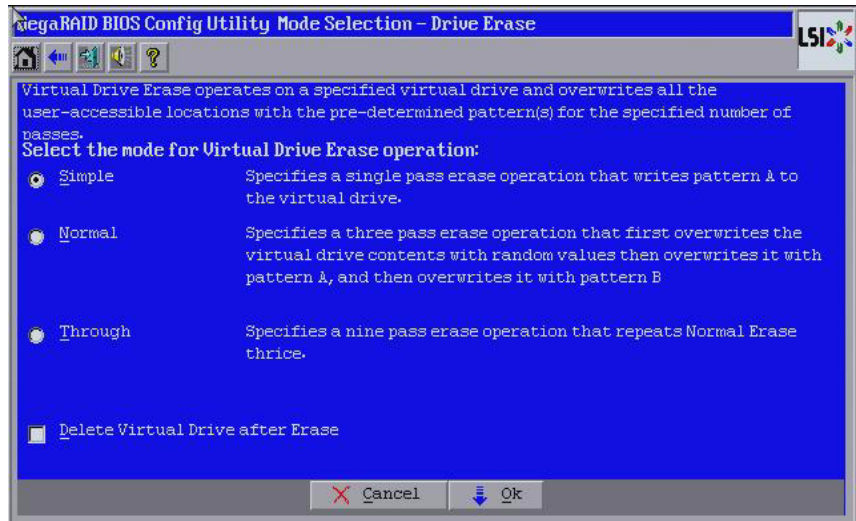
Figure 4.72 Virtual Drive Dialog



Note: The radio button appears for the unsecured virtual drives that are online.

The Mode Selection - Drive Erase Dialog appears, as shown in the following figure.

Figure 4.73 Virtual Drive Erase Dialog



4. Select any of the following options.
 - **Simple** (Alt + S) –After you select this option and click **OK**, if the **Delete Virtual Drive after Erase** check box is selected, a confirmation dialog appears.
 - **Normal** (Alt + N) – After you select this option and click **OK**, if the **Delete Virtual Drive after Erase** check box is selected, a confirmation dialog appears.
 - **Thorough** (Alt + T) –After you select this option and click **OK**, if the **Delete Virtual Drive after Erase** check box is selected, a confirmation dialog appears.
 - **Delete Virtual Drive after Erase** (Alt + D) - if you select this check box, the virtual drive is erased, and a confirmation dialog appears.
 - **OK** (Alt + O) – Click **OK** and, if the **Delete Virtual Drive after Erase** check box is selected a confirmation dialog appears.
 - **Cancel** (Alt + C) – Clicking this option, closes the dialog, and the WebBIOS navigates back to the Virtual Drive dialog.

4.10.2.1 Group Show Progress for Virtual Drive Erase

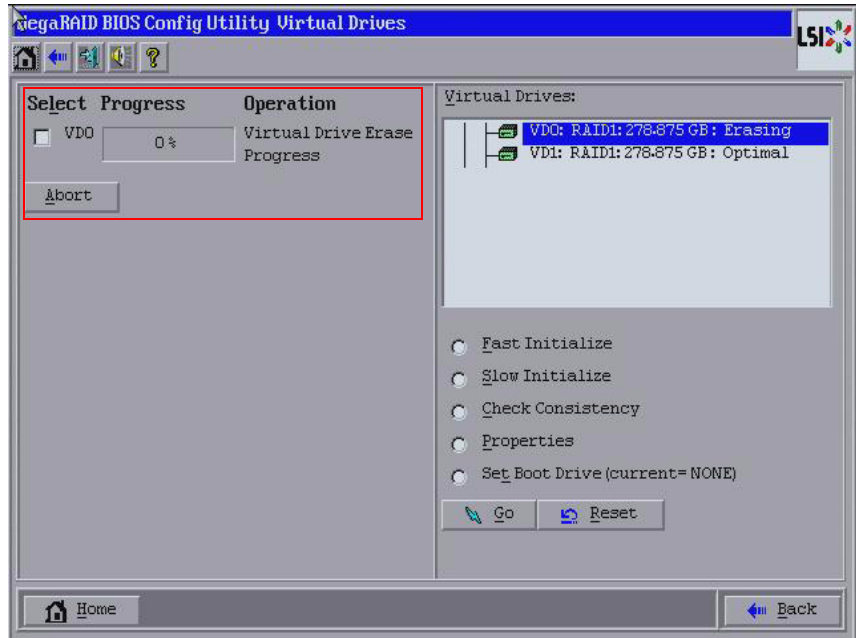
The virtual drive erase operation is a time-consuming operation that is performed as a background task.

Follow these steps to view the progress of **Virtual Drive Erase**.

1. Click on the **Virtual Drives** link on the WebBIOS main menu.

The Virtual Drives dialog appears, as shown in the following figure.

Figure 4.74 Virtual Drive Dialog




2. To abort the virtual drive erase, select the check box of the operation you want to abort, click **Abort**.

4.11 Viewing System Event Information

The SAS controller firmware monitors the activity and performance of all storage configurations and devices in the system. When an event occurs (such as the creation of a new virtual drive or the removal of a drive) an event message is generated and is stored in the controller NVRAM. You can use the WebBIOS CU to view these event messages. To do this, click **Events** on the main WebBIOS CU screen. The Event Information screen appears, as shown in the following figure.

Figure 4.75 Event Information Screen



The right side of the screen is blank until you select an event to view. The First Sequence and Last Sequence fields in the upper left of the screen show you how many event entries are currently stored.

To view event information, follow these steps:

1. Select an Event Locale from the menu. For example, select **Enclosure** to view events relating to the drive enclosure.
2. Select an Event Class: *Information*, *Warning*, *Critical*, *Fatal*, or *Dead*.
3. Enter a Start Sequence number, between the First Sequence number and the Last Sequence number. The higher the number, the more recent the event.
4. Enter the Number of events of this type that you want to view, and click **Go**.

The first event in the sequence appears in the right panel.

5. Click **Next** or **Prev** to page forward or backward through the sequence of events.
6. If you want, select different event criteria in the left panel, and click **Go** again to view a different sequence of events.

Each event entry includes a timestamp and a description to help you determine when the event occurred and what it was.

4.12 Managing Configurations

This section includes information about maintaining and managing storage configurations.

4.12.1 Running a Consistency Check

You should periodically run a consistency check on fault-tolerant virtual drives. A consistency check verifies that the redundancy data is correct and available for RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, and RAID 60 drive groups. To do this, follow these steps:

1. On the main WebBIOS CU screen, select a virtual drive.
2. Click **Virtual Drives**.
3. When the Virtual Drive screen appears, select **CC** in the lower left panel, and click **Go**.

The consistency check begins.

If the WebBIOS CU finds a difference between the data and the parity value on the redundant drive group, it assumes that the data is accurate and automatically corrects the parity value. Be sure to back up the data before running a consistency check if you think the data might be corrupted.

4.12.2 Deleting a Virtual Drive

You can delete any virtual drive on the controller if you want to reuse that space for a new virtual drive. The WebBIOS CU provides a list of configurable drive groups where there is a space to configure. If multiple virtual drives are defined on a single drive group, you can delete a virtual drive without deleting the whole drive group.

To delete a virtual drive, follow these steps:

Attention: Back up any data that you want to keep before you delete the virtual drive.

1. On the main WebBIOS CU screen, click a virtual drive icon to select a virtual drive.

The Virtual Drive screen appears.

2. Click **Virtual Drives**.
3. When the Virtual Drive screen appears, select **Delete** in the lower left panel, and click **Go**.
4. When the message appears, confirm that you want to delete the virtual drive.

4.12.3 Importing or Clearing a Foreign Configuration

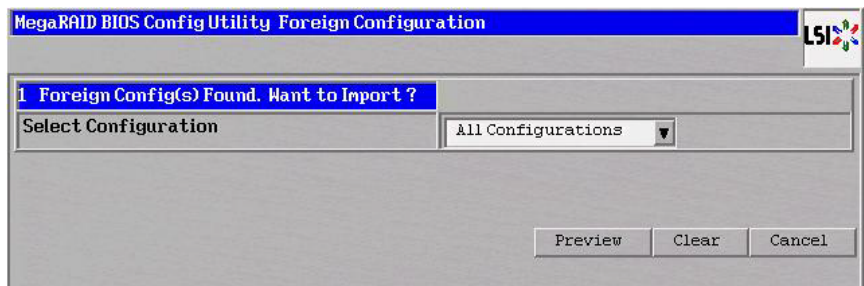
A *foreign configuration* is a storage configuration that already exists on a replacement set of drives that you install in a computer system. In addition, if one or more drives are removed from a configuration, by a cable pull or drive removal, for example, the configuration on those drives is considered a foreign configuration by the RAID controller.

The WebBIOS CU allows you to import the foreign configuration to the RAID controller, or to clear the configuration so you can create a new configuration using these drives.

Note: When you create a new configuration, the WebBIOS CU shows only the unconfigured drives. Drives that have existing configurations, including foreign configurations, do **not** appear. To use drives with existing configurations, you must first clear the configuration on those drives.

If WebBIOS CU detects a foreign configuration, the import screen appears, as shown in the following figure.

Figure 4.76 Foreign Configuration Import Screen



Follow these steps to import or clear a foreign configuration.

1. Click the drop-down list to show the configurations.

The GUID (Global Unique Identifier) entries on the drop-down list are OEM names and will vary from one installation to another.

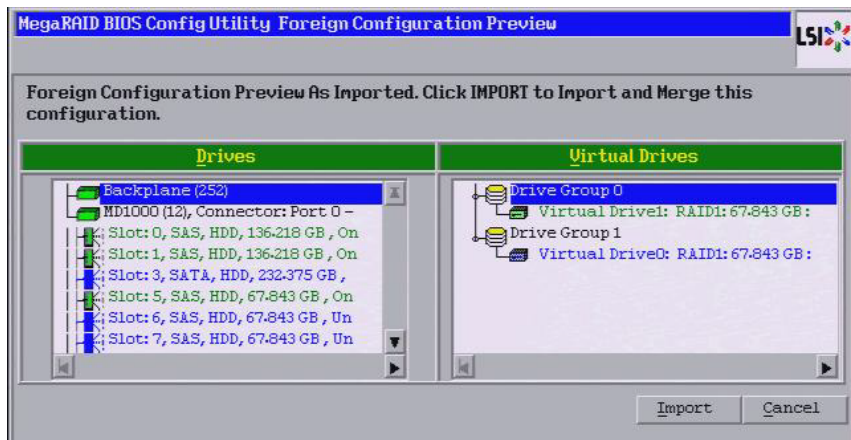
2. Select a configuration or **All Configurations**.

3. Perform one of the following steps:

- a. Click **Preview** if you want to preview the foreign configuration. The preview screen appears, as shown in [Figure 4.77](#).
- b. Click **Clear** if you want to clear the configuration and reuse the drives for another virtual drive.

If you click **Cancel**, it cancels the importation or preview of the configuration.

Figure 4.77 Foreign Configuration Preview Screen



The right panel shows the virtual drive properties of the foreign configuration. In this example, there is a RAID 1 virtual drive with 1,000 Mbytes. The left panel shows the drives that comprise the foreign configuration.

4. Click **Import** to import this foreign configuration and use it on this controller.

If you click **Cancel** to clear the configuration and reuse the drives for another virtual drive, you return to [Figure 4.76](#).

4.12.3.1 Foreign Configurations in Cable Pull and Drive Removal Scenarios

If one or more drives are removed from a configuration, by a cable pull or drive removal, for example, the configuration on those drives is considered a foreign configuration by the RAID controller.

Use the **Foreign Configuration Preview** screen to import or clear the foreign configuration in each case. The import procedure and clear procedure are described in [Section 4.12.3, "Importing or Clearing a Foreign Configuration."](#)

The following scenarios can occur with cable pulls or drive removals.

Note: If you want to import the foreign configuration in any of the following scenarios, you should have all of the drives in the enclosure before you perform the import operation.

1. Scenario #1: If all of the drives in a configuration are removed and re-inserted, the controller considers the drives to have foreign configurations.

Import or clear the foreign configuration. If you select **Import**, automatic rebuilds occur in redundant virtual drives.

Note: Start a consistency check immediately after the rebuild is complete to ensure data integrity for the virtual drives. See [Section 4.12.1, "Running a Consistency Check,"](#) for more information about checking data consistency.

2. Scenario #2: If some of the drives in a configuration are removed and re-inserted, the controller considers the drives to have foreign configurations.

Import or clear the foreign configuration. If you select **Import**, automatic rebuilds occur in redundant virtual drives.

Note: Start a consistency check immediately after the rebuild is complete to ensure data integrity for the virtual drives. See [Section 4.12.1, "Running a Consistency Check,"](#) for more information about checking data consistency.

3. Scenario #3: If all of the drives in a virtual drive are removed, but at different times, and re-inserted, the controller considers the drives to have foreign configurations.

Import or clear the foreign configuration. If you select **Import**, all drives that were pulled *before* the virtual drive became offline will be imported and then automatically rebuilt. Automatic rebuilds will occur in redundant virtual drives.

4. If the drives in a non-redundant virtual drive are removed, the controller considers the drives to have foreign configurations.

Import or clear the foreign configuration. No rebuilds occur after the import operation because there is no redundant data to rebuild the drives with.

4.12.3.2 Importing Foreign Configurations from Integrated RAID to ServeRAID

The IBM Integrated RAID solution simplifies the configuration options and provides firmware support in its host controllers. IBM offers two types of Integrated RAID (IR): Integrated Mirroring (IM) and Integrated Striping (IS).

You can import an IM or IS RAID configuration from an IR system into a ServeRAID system. The ServeRAID system treats the IR configuration as a foreign configuration. You can import or clear the IR configuration.

4.12.3.3 Troubleshooting Information

An IR virtual drive can have either 64 Mbytes or 512 Mbytes available for metadata at the end of the drive. This data is in LSI Data Format (LDF). ServeRAID virtual drives have 512 Mbytes for metadata at the end of the drive in the Disk Data format (DDF).

To import an IR virtual drive into ServeRAID, the IR virtual drive must have 512 Mbytes in the metadata, which is the same amount of megadata as in a ServeRAID virtual drive. If the IR virtual drive has only 64 Mbytes when you attempt to import it into ServeRAID, the import will fail because the last 448 Mbytes of your data will be overwritten and the data lost.

If your IR virtual drive has only 64 Mbytes for metadata at the end of the drive, you cannot import the virtual drive into ServeRAID. You need to use another upgrade method, such as backup/restore to the upgraded virtual drive type.

In order to import an IR virtual drive into a ServeRAID system, use the **Foreign Configuration Preview** screen to import or clear the foreign

configuration. The import procedure and the clear procedure are described in [Section 4.12.3, “Importing or Clearing a Foreign Configuration.”](#)

4.12.4 Migrating the RAID Level of a Virtual Drive

As the amount of data and the number of drives in your system increase, you can use RAID-level migration to change a virtual drive from one RAID level to another. You do not have to power down or reboot the system. When you migrate a virtual drive, you can keep the same number of drives, or you can add drives. You can use the WebBIOS CU to migrate the RAID level of an existing virtual drive.

Note: While you can apply RAID-level migration at any time, you should do so when there are no reboots. Many operating systems issues I/O operations serially (one at a time) during boot. With a RAID-level migration running, a boot can often take more than 15 minutes.

Migrations are allowed for the following RAID levels:

- RAID 0 to RAID 1
- RAID 0 to RAID 5
- RAID 0 to RAID 6
- RAID 1 to RAID 0
- RAID 1 to RAID 5
- RAID 1 to RAID 6
- RAID 5 to RAID 0
- RAID 5 to RAID 6
- RAID 6 to RAID 0
- RAID 6 to RAID 5

[Table 4.4](#) lists the number of additional drives required in some cases when you change the RAID level of a virtual drive.

Table 4.4 Additional Drives Required for RAID-Level Migration

From RAID Level to RAID Level	Original Number of Drives in Drive Group	Additional Drives Required
RAID 0 to RAID 1	RAID 0: 1 drive	1
RAID 0 to RAID 5	RAID 0: 1 drive	2
RAID 0 to RAID 6	RAID 0: 1 drive	3
RAID 1 to RAID 5	RAID 1: 2 drives	1
RAID 1 to RAID 6	RAID 1: 2 drives	1

Follow these steps to migrate the RAID level:

Attention: Back up any data that you want to keep before you change the RAID level of the virtual drive.

1. On the main WebBIOS CU screen, select Virtual Drives.
2. Choose your virtual drive from the list. If only one virtual drive is configured, you will automatically be taken to the **Virtual Drives** menu.
3. From the Virtual Drives menu, select **Properties**.
4. From the **Properties** menu, select **Adv Opers** under the **Advanced Operations** heading.

The Advanced Operations screen appears, as shown in the following figure.

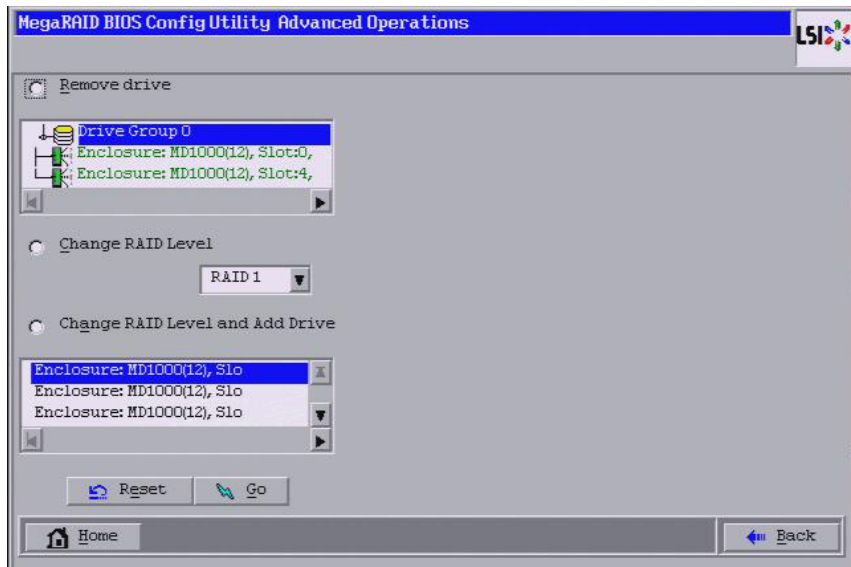


Figure 4.78 Advanced Operations Screen

5. Select either **Change RAID Level** or **Change RAID Level and Add Drive**.
 - If you select **Change RAID Level**, change the RAID level from the drop-down menu.
 - If you select **Change RAID Level and Add Drive**, change the RAID level from the drop-down menu and then select one or more drives to add from the list of drives.

The available RAID levels are limited, based on the current RAID level of the virtual drive plus the number of drives available.

6. Click **Go**.
7. When the message appears, confirm that you want to migrate the RAID level of the virtual drive.

A reconstruction operation begins on the virtual drive. You must wait until the reconstruction is completed before you perform any other tasks in the WebBIOS CU.

4.12.5 New Drives Attached to a MegaRAID Controller

When you insert a new drive on a MegaRAID system, if the inserted drive does not contain valid DDF metadata, the drive displays as JBOD for MegaRAID Entry level controllers. If the drive does contain valid DDF metadata, its drive state is Unconfigured Good.

A new drive in JBOD drive state is exposed to the host operating system as a stand-alone drive. Drives in JBOD drive state are not part of the RAID configuration because they do not have valid DDF records. The operating system can install and run anything on JBOD drives.

Automatic rebuilds always occur when the drive slot status changes, for example, when you insert a drive or remove a drive, so that a hot spare can be used. However, a new drive in JBOD drive state (without a valid DDF record), will not perform an automatic rebuild.

To start an automatic rebuild on the new JBOD drive, you must change the drive state from JBOD to Unconfigured Good. (Rebuilds start only on Unconfigured Good drives.) After you set the drive state to Unconfigured Good, the drive state information always remains on the drive, and you can use the drive for configuration.

See [Section 4.12.3.3, "Troubleshooting Information"](#) for more information about DDF and metadata.

Chapter 5

MegaRAID Storage Manager Overview and Installation

This chapter provides a brief overview of the MegaRAID Storage Manager (MSM) software and explains how to install it on the supported operating systems. This chapter has the following sections:

- [Section 5.1, “Overview”](#)
 - [Section 5.2, “Hardware and Software Requirements”](#)
 - [Section 5.3, “Prerequisites to Running MegaRAID Storage Manager Remote Administration”](#)
 - [Section 5.4, “Installing MegaRAID Storage Manager”](#)
 - [Section 5.5, “MegaRAID Storage Manager Support and Installation on the VMware Operating System”](#)
 - [Section 5.6, “Installing and Configuring an SNMP Agent”](#)
-

5.1 Overview

MegaRAID Storage Manager software enables you to configure, monitor, and maintain storage configurations on ServeRAID-M controllers. The MegaRAID Storage Manager graphical user interface (GUI) makes it easy for you to create and manage storage configurations.

5.1.1 Creating Storage Configurations

MegaRAID Storage Manager software enables you to easily configure the controllers, drives, and virtual drives on your workstation or server. The Configuration Wizard greatly simplifies the process of creating drive groups and virtual drives.

You can use the Configuration Wizard Auto Configuration mode to automatically create the best possible configuration with the available

hardware. You can use the Guided Configuration mode, which asks you a few brief questions about the configuration, and then creates it for you. Or you can use the Manual Configuration mode, which gives you complete control over all aspects of the storage configuration.

The Modify Drive Group Wizard enables you to increase the capacity of a virtual drive and to change the RAID level of a drive group.

Note: The Modify Drive Group Wizard was previously known as the Reconstruction Wizard.

5.1.2 Monitoring Storage Devices

MegaRAID Storage Manager software displays the status of controllers, virtual drives, and drives on the workstation or server that you are monitoring. System errors and events are recorded in an event log file and are displayed on the screen. Special device icons appear on the screen to notify you of drive failures and other events that require immediate attention.

5.1.3 Maintaining Storage Configurations

You can use MegaRAID Storage Manager software to perform system maintenance tasks such as running patrol read operations, updating firmware, and running consistency checks on drive groups that support redundancy.

5.2 Hardware and Software Requirements

The hardware requirements for MegaRAID Storage Manager software are as follows:

- A PC-compatible computer with an IA-32 (32-bit) Intel Architecture processor or an EM64T (64-bit) processor; also compatible with SPARC V9 architecture-based systems
- A minimum of 256 Mbytes of system memory (512 Mbytes recommended)
- A drive with at least 400 Mbytes available free space

The supported operating systems for the MegaRAID Storage Manager software are as follows:

- Microsoft® Windows® 2003, Microsoft Windows 2008, Microsoft Windows 2008 SP2, Microsoft Windows 2008 R2, and Windows R2 SP1
- Red Hat® Enterprise Linux™ versions 4.7, 4.8, 5.5, 5.6, 5.7, 6.0, 6.1, and 6.2
- SUSE™ SLES Server 10 SP3, 10 SP4, 11, and 11 SP1
- VMWare® ESX 4.0, 4.1, and 5.0

Refer to your server documentation and to the operating system documentation for more information on hardware and operating system requirements.

To download the latest operating system drivers, see <http://www.ibm.com/systems/support/>.

5.3 Prerequisites to Running MegaRAID Storage Manager Remote Administration

The MegaRAID Storage Manager software requires ports 3071 and 5571 to be open to function. Follow these steps to prepare to run the MegaRAID Storage Manager Remote Administration.

1. Configure the system with a valid IP address.

Make sure the IP address does not conflict with another in the sub network.

Ports, such as 3071 and 5571, are open and available for the MegaRAID Storage Manager framework communication.

2. Disable all security manager and firewall.
3. Configure the multicasting.

Make sure Class D multicast IP addresses are registered (at least 229.111.112.12 should be registered for the MegaRAID Storage Manager software to work); if not, create a static route using the following command:

```
Route add 229.111.112.12 dev eth1
```

4. Install the MegaRAID Storage Manager software. If the MegaRAID Storage Manager software is already installed, restart the MegaRAID Storage Manager Framework.

5.4 Installing MegaRAID Storage Manager

This section explains how to install (or reinstall) MegaRAID Storage Manager software on your workstation or server for the supported operating systems: Microsoft Windows, Red Hat Linux, and SUSE Linux.

Note: If you want to manage multiple Linux servers, you must configure the network routing properly on a particular subnet. Please refer to the MegaRAID Storage Management installation `readme.txt` for details.

5.4.1 Prerequisite for MegaRAID Storage Manager Installation

The MegaRAID Storage Manager software installation script also installs the LSI SNMP agent, Red Hat Package Manager (RPM). The LSI SNMP agent application depends upon the standard SNMP-Util package.

Make sure that the SNMP-Util package is present in the system before you install the MegaRAID Storage Manager software.

The SNMP-Util package includes the RPM's `net-snmp-libs`, `net-snmp-utils`, and additional dependent RPM's. Make sure that these RPM's are installed from the operating system media before you install the MegaRAID Storage Manager software.

5.4.2 Installing MegaRAID Storage Manager Software on Microsoft Windows

Follow these steps if you need to install MegaRAID Storage Manager software on a system running Microsoft Windows 2000, Microsoft Windows Server 2003R2, Microsoft Windows XP, or Microsoft Windows Vista:

1. Insert the MegaRAID Storage Manager software installation CD in the CD-ROM drive.

If necessary, find and double-click the `setup.exe` file to start the installation program.


2. When the Welcome screen appears, click **Next**.

If the MegaRAID Storage Manager software is already installed on this system, then an upgraded installation occurs.

3. Read and accept the user license in the following figure and click **Next**.

The Customer Information screen appears, as shown in the following figure.

Figure 5.1 License Agreement



The Customer Information window appears, as shown in the following figure.


Figure 5.2 Customer Information Screen

The screenshot shows a Windows-style dialog box titled "MegaRAID Storage Manager v11.08.02.0200 - InstallShield Wizard". The main heading is "Customer Information". Below the heading, it says "Please enter your information:". There are two text input fields: "User Name:" with "Administrator" entered, and "Organization:" which is empty. Below these fields, there is a section titled "Allow availability of this application for:" with two radio button options: "All users" (which is selected) and "Only for current user (Administrator)". At the bottom left, the "InstallShield" logo is visible. At the bottom right, there are three buttons: "< Back", "Next >", and "Cancel".

4. Enter your user name and organization name. In the bottom part of the screen, select an installation option:
 - If you select **All users**, any user with administrative privileges can use this version of MegaRAID Storage Manager software to view or change storage configurations.
 - If you select **Only for current user**, the MegaRAID Storage Manager shortcuts and associated icons are available only to the user with this user name.
5. Click **Next** to continue.
6. On the next screen, accept the default Destination Folder, or click **Change** to select a different destination folder.
7. Click **Next** to continue.

The Setup Type screen appears, as shown in the following figure.

Figure 5.3 Setup Type Screen



8. Select one of the Setup options.
The options are explained in the screen text.
 - Select **Complete** if you are installing MegaRAID Storage Manager software on a server.
 - Select **Custom Installation** if you want to select individual program components.
9. Click **Next** to continue.
 - If you selected **Custom Installation** as your setup option, the second Setup Type screen appears, as shown in the following figure.
 - If you selected **Complete** as your setup option, the **LDAP Logon Information** dialog appears, as shown in the following figure.

Figure 5.4 LDAP Logon Information

The image shows a 'LDAP Login' dialog box with the LSI logo in the top right corner. The dialog contains several input fields: 'LDAP Server IP Address:', 'User Name:', 'Password:', and 'Distinguished Name :'. There is a text instruction 'Use your LDAP server credentials to login' with a help icon. Below these are two checkboxes: 'Use Default Port' (checked) and 'Remember my Login Details' (unchecked). The 'Port:' field is set to '389'. At the bottom are 'Login' and 'Cancel' buttons.

10. To specify LDAP configuration details, select **Yes**, and perform the following substeps, or if you do not want to specify LDAP configuration details, click **No** and click **Next**.
 - a. Enter the LDAP server's IP address in the **Server IP** field.
 - b. Enter the LDAP server's user name in the **User name** field.

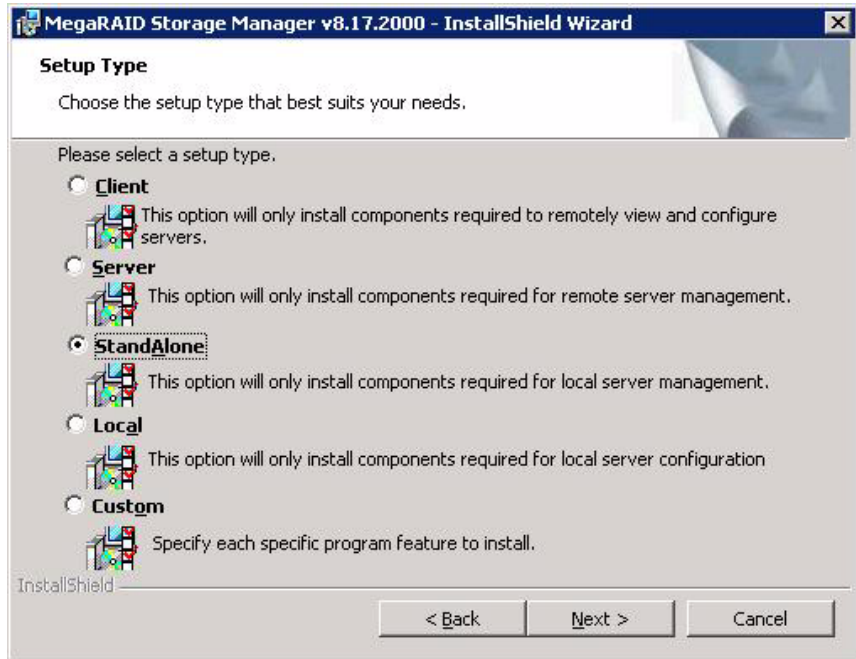
An example of a user name can be `username@testldap.com`.
 - c. Enter the name of the Domain Controller in the **Distinguished User name** field.

As an example, the Domain Controller name can be `dc=TESTLDAP, dc=com`.
 - d. Enter the LDAP server's port number in the **Port** field.
 - e. Select the **Use LDAP as default Login** check box to always connect to the LDAP server.

All of the values entered in this dialog are saved in the `ladp.properties` file.
11. Click **Next**.

- The Setup Type dialog box appears.
- Click **Install** to begin the installation.

Figure 5.5 Setup Type Screen



- Select one of the custom setup options.

The screen text explains the options.

- Select **Client** if you are installing MegaRAID Storage Manager software on a PC that will be used to view and configure servers over a network.

In the Client mode of installation, the MegaRAID Storage Manager software installs only client-related components, such as the MegaRAID Storage Manager GUI, and monitor configurator.

Use this mode when you want to manage and monitor servers remotely. When you install the MegaRAID Storage Manager software in Client mode on a laptop or a desktop, you can log in to a specific server by providing the IP address.


- Select **Server** to install only those components required for remote server management. Server mode is used to configure alerts and e-mail settings. To begin installation, click on **Install** on the next screen that appears.
- Select **StandAlone** if you will use MegaRAID Storage Manager software to create and manage storage configurations on a standalone workstation.

Note: If you select Client or Standalone as your setup option, the LDAP Logon Information dialog appears, as shown in [Figure 5.4](#).

- Select **Local** if you want to view only the workstation that has the MegaRAID Storage Manager software installed. You cannot discover other remote servers, and other remote servers cannot connect to your workstation. In a local mode installation, you use the loopback address instead of the IP address.
- Select **Custom** if you want to specify individual program features to install.

If you select **Custom**, a window listing the installation features appears, as shown in the following figure. Select the features you want on this screen.

Figure 5.6 Custom Setup Screen



14. Click **Next**.

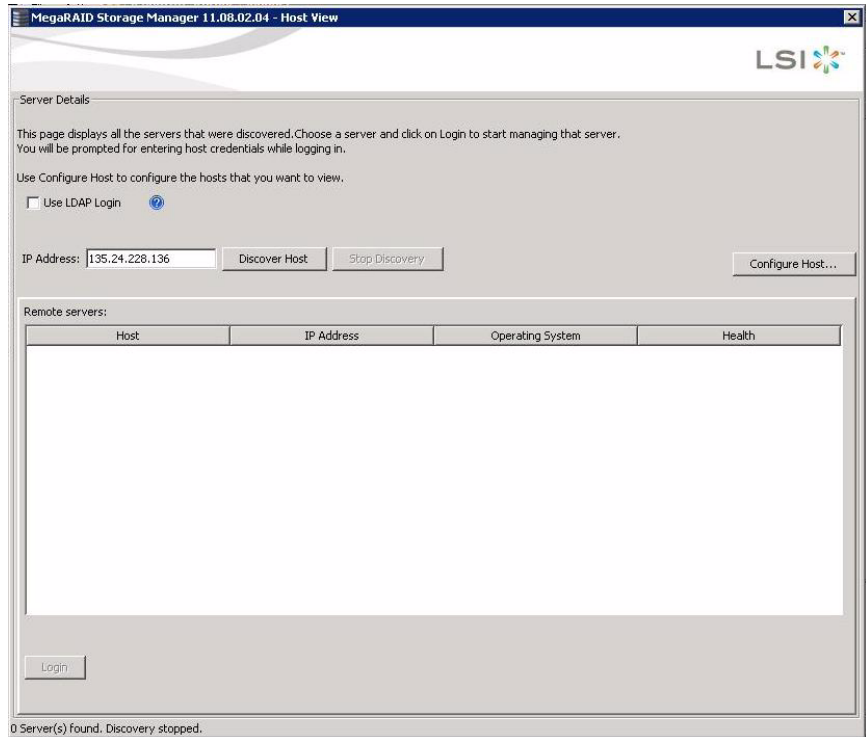
15. Click **Install** to install the program.

16. When the final Configuration Wizard window appears, click **Finish**.

If you select **Client** installation for a computer that is used to monitor servers, and if no available servers exist with a registered framework on the local subnet (that is, servers with a complete installation of the MegaRAID Storage Manager software), the server window appears.

The MegaRAID Storage Manager - Host View window does not list any servers. You can use the MegaRAID Storage Manager - Host View window to manage systems remotely.

Figure 5.7 MegaRAID Storage Manager - Host View Window



5.4.3 Source Port and Destination Port for MSM Components

[Table 5.1](#) lists the sending port information and destination port for the MSM components.

Table 5.1 Source Port and Destination Port for MSM Components

Component	Sending Port/Source Port	Listening Port/Dest Port
MSM GUI Client	49258 - 50058 (TCP)	49258 - 50058 (TCP)
Framework	49258 - 50058 (TCP)	3071(TCP/UDP)
Popup	49152 - 65535 (TCP)	49152 - 65535 (TCP)
MrMonitor	49152 - 65535 (TCP)	49152 - 65535 (TCP)

5.4.4 Prerequisites for Installing MegaRAID Storage Manager on the RHEL6.X x64 Operating System

Before installing the MegaRAID Storage Manager software on RHEL 6.X x64 system, install the following RPMs. Without these RPMs the MegaRAID Storage Manager software might not install properly or might not work as expected.

- libstdc++-4.4.4-13.el6.i686.rpm
- compat-libstdc++-33-3.2.3-69.i686.rpm
- libXau-1.0.5-1.el6.i686.rpm
- libxcb-1.5-1.el6.i686.rpm
- libX11-1.3-2.el6.i686.rpm
- libXext-1.1-3.el6.i686.rpm
- libXi-1.3-3.el6.i686.rpm
- libXtst-1.0.99.2-3.el6.i686.rpm

The RHEL6.X x64 complete operating system installation is required for the MegaRAID Storage Manager software to work. The above mentioned rpm's come as part of RHEL6.X x64 Operating System DVD. These RPMs might need additional dependent RPMs as well, and you must install all the dependent RPMs on the target system.

Note: The RPM's versions mentioned above might get changed in the future RHEL6.x releases. Install the corresponding RPM's from the operating system installation media.

Note: The MegaRAID Storage Manager software now provides an additional binary to run it in a native 64-bit Linux environment.

5.4.5 Installing MegaRAID Storage Manager for Linux Operating System

Follow these steps if you need to install the MegaRAID Storage Manager software on a system running Red Hat Linux 3.0, 4.0,5.0, 6.0 or SUSE Linux/SLES 9, 10, and 11:

1. Copy the `MSM_linux_installer-11.02.00-00.tar.gz` file to a temporary folder.
2. Untar the `MSM_linux_installer-11.02.00-00.tar.gz` file using the following command:

```
tar -zxvf MSM_linux_installer-11.02.00-00-...tar.gz
```

A new disk directory is created.

3. Go to the new disk directory.
4. In the disk directory, find and read the `readme.txt` file.
5. To start the installation, enter the following command:

```
cd install.csh -a
```

The preceding command works only if `csh` shell is installed; otherwise, use the following command:

```
install.csh
```

If you select **Client** installation for a computer that is used to monitor servers, and if no available servers exist with a registered framework on the local subnet (that is, servers with a complete installation of the MegaRAID Storage Manager software), the MegaRAID Storage Manager - Host Name window appears. The MegaRAID Storage Manager - Host Name window does not list any servers. You can use this window to manage systems remotely.

To install the software using an interactive mode, execute the command `./install.csh` from the installation disk.

To install the product in a non-interactive mode or silent mode, use the command `./install.csh [-options] [-ru popup]` from the installation disk. The installation options are as follows:

- Complete
- Client Component Only

- StandAlone
- Local
- Server

The `-ru popup` command removes the pop-up from the installation list.

You also can run a non-interactive installation using the `RunRPM.sh` command.

The installer offers the following setup options:

- **Complete** – This installs all the features of the product.
- **Client Components Only** – The storelib feature of the product is not installed in this type of installation. As a result, the resident system can only administer and configure all of the servers in the subnet, but it cannot serve as a server.
- **StandAlone** – Only the networking feature will not be installed in this case. But the system can discover other servers in the subnet and can be discovered by the other servers in the subnet.
- **Local** - This option enables you to view only the workstation that has the MegaRAID Storage Manager software installed. You will not be able to discover other remote servers and other remote servers will also not be able to connect to your workstation. In a local mode installation, you will be using the loopback address instead of the IP address.
- **Server** - This option installs components required for remote server management

This installation helps you select any of the setup types, but if you run `RunRPM.sh`, it installs the complete feature.

Note: To install and run the MegaRAID Storage Manager software on RHEL 5, you need to disable SELinux.

5.4.6 Linux Error Messages

The following messages can appear while you are installing the MegaRAID Storage Manager software on a Linux operating system:

- More than one copy of MegaRAID Storage Manager software has been installed.

This message indicates that the user has installed more than one copy of the MegaRAID Storage Manager software. (This step can be done by using the `rpm-force` command to install the `rpm` file directly, which is not recommended, instead of using the `install.sh` file.) In such cases, the user must uninstall all of the `rpm` files manually before installing the MegaRAID Storage Manager software with the procedure listed previously.

- The version is already installed.

This message indicates that the version of the MegaRAID Storage Manager software you are trying to install is already installed on the system.

- The installed version is newer.

This message indicates that a version of the MegaRAID Storage Manager software is already installed on the system, and it is a newer version than the version you are trying to install.

- Exiting installation.

This message appears when the installation is complete.

- RPM installation failed.

This message indicates that the installation failed for some reason. Additional message text explains the cause of the failure.

5.4.7 Kernel Upgrade

If you want to upgrade the kernel in the Linux operating system, you must restart the MegaRAID Storage Manager Framework and Services in the same order by entering the following command:

```
/etc/init.d/vivaldiframeworkd restart
```

5.4.8 Uninstalling MegaRAID Storage Manager Software on Linux

To uninstall the MegaRAID Storage Manager software on a system running Linux, follow these steps:

1. Go to `/usr/local/MegaRAID Storage Manager`.
2. Run `./uninstaller.sh`.

This procedure uninstalls the MegaRAID Storage Manager software.

5.4.8.1 Executing a CIM Plug-in on Red Hat Enterprise Linux 5

To execute a Common Information Model (CIM) plug-in on Red Hat Enterprise Linux 5, you must create the following symbolic links:

1. `cd /usr/lib` on RHEL 5
2. Search for `libcrypto`, `libssl`, and `libsysfs` libraries as follows:

```
ls -lrt libcrypto*, ls -lrt libssl*, ls -lrt libsysfs*
```
3. If the files `libcrypto.so.4`, `libssl.so.4`, and `libsysfs.so.1` are missing, manually create sym links as follows:
 - `ln -s libcrypto.so libcrypto.so.4`
 - `ln -s libssl.so libssl.so.4`
 - `ln -s libsysfs.so libsysfs.so.1`

For more information about CIM, see Section [5.5.3, Limitations](#).

If the `.so` files are not present in the `/usr/lib` directory, create a link with the existing version of the library. For example, if `libcrypto.so.6` is present and `libcrypto.so` is not, create the link as follows:

```
ln -s libcrypto.so.6 libcrypto.so.4
```

On a 64-bit operating system, the system libraries are present in the `/usr/lib64` directory by default. However, for supporting CIM Plug-in, make sure that the libraries are also present in `/usr/lib` by installing the appropriate RPMs.

5.5 MegaRAID Storage Manager Support and Installation on the VMware Operating System

This section documents the installation of the MegaRAID Storage Manager software on VMware ESX (also known as Classic) and on the VMware® ESXi operating system.

5.5.1 Prerequisites for Installing MegaRAID Storage Manager for VMware

For the VMware 4.1 operating system, it is necessary to create a soft link as follows before installing the MegaRAID Storage Manager software. Run the following command to create the necessary soft link required for the MegaRAID Storage Manager software to work.

```
sudo ln -sf  
/lib/libgcc_s.so.1/usr/lib/vmware/lib/libgcc_s.so.1
```

For VMware ESXi 5.0 to work with the MegaRAID Storage Manager software, the SMI-S Provider must be installed.

5.5.2 Uninstalling MegaRAID Storage Manager on the VMware Operating System

To uninstall the Server Component of the MegaRAID Storage Manager software on the VMware operating system, either use the `Uninstall` command in the Program menu, or run the script `/usr/local/MegaRAID Storage Manager/uninstaller.sh`.

You need to keep in mind the following points:

- A MegaRAID Storage Manager upgrade is supported in this release. Future releases can update this release.
- To shut down the MegaRAID Storage Manager Framework service, run the following command:

```
/etc/init.d/vivaldiframeworkd stop
```

The Linux RPM of the MegaRAID Storage Manager software works under the console with minimal changes. Hardware RAID is currently supported in ESX 4.x.

Note: There is a known limitation that virtual drives that are created or deleted will not be reflected to the kernel. The workaround is to reboot the server or to run `esxcfg-rescan <vmhba#>` from COS shell.

5.5.3 Limitations

The following are the limitations of this installation and configuration.

- No status information exists for the controller
- Events are collected as long as the MegaRAID Storage Manager software runs on the client.
- The MegaRAID Storage Manager software on VMware responds slower as compared to the response of the MegaRAID Storage Manager software on Windows/Linux/Solaris. Events are collected from the time a client logs in to an ESXi machine for the first time, and it continues to be collected as long as the Framework is running.

5.5.3.1 Differences in MegaRAID Storage Manager for the VMware ESXi System

The following are some of the differences in the MegaRAID Storage Manager utility when you manage a VMware server.

- The following limitations apply to the system information exposed through the application:
 - Only the IP address and the host name appear.
 - No support exists for the controller health information.
- Authentication support:
 - The MegaRAID Storage Manager software allows CIMOM server authentication with the user ID and the password for VMware.
 - Access to VMware ESXi hosts is controlled based on the user privileges. Only root users can have full access, while the non-root users can have only view only access.
 - Multiple root users can simultaneously login using 'Full Access' mode to access the VMware ESXi server.
- Event logging:

- Event logging support is available for the VMware ESXi operating system, but it works differently than the normal MegaRAID Storage Manager framework mode. The event logging feature for the MegaRAID Storage Manager Client connected to a VMware ESXi system behaves as follows
- The support for retrieving initial logs is limited to 30 events. Only those events that occur after a client logs in for the first time to an VMware ESXi server appear in the Event Logger dialog.
- The System logs are not displayed.
- The *Save log* feature is not supported; however, the *Save Log as Text* feature is supported.
- The *View Log* option allows you to view the logs saved in a text file on the Event Logger dialog.
- Refreshing of the MegaRAID Storage Manager GUI after any updates on the firmware is slower for a client connected to VMware ESXi hosts, compared to one that is connected to a Windows/Linux/Solaris host.
- VMware ESXi is supported only on a full installation of the MegaRAID Storage Manager software; standalone, client-only, server-only, and local modes do not support VMware ESXi management.
- VMware ESXi is supported on following operating systems:
 - Microsoft Windows Server
 - RHEL
 - SuSE Linux

5.6 Installing and Configuring an SNMP Agent

A Simple Network Management Protocol (SNMP)-based management application can monitor and manage devices through SNMP extension agents. The MegaRAID SNMP subagent reports the information about the RAID controller, virtual drives, physical devices, enclosures, and other items per SNMP request. The SNMP application monitors these devices for issues that might require administrative attention.

This section describes the installation and configuration of the MegaRAID SNMP agent on Linux and Windows operating systems.

Note: The complete installation of the MegaRAID Storage Manager software installs the SNMP agent. However, you can install the SNMP agent (installer) on a system separately, without the MegaRAID Storage Manager software being installed.

5.6.1 Prerequisite for LSI SNMP Agent RPM Installation

The LSI SNMP agent application depends upon the standard SNMP Utils package. Make sure that the SNMP-Util package is present in the system before you install LSI SNMP agent RPM.

The SNMP-Util package includes the RPM's net-snmp-libs, net-snmp-utils, and additional dependent RPMs.

Make sure that these RPM's are installed from the operating system media before you install the LSI SNMP agent RPM.

5.6.2 Prerequisite for Installing SNMP Agent on Linux Server

The SNMP application requires the standard library libsysfs. Make sure that this library is present in the system before installing the SNMP RPM.

The minimum library versions required for installing SNMP server are as follows.

- libsysfs version 2.0. This library is available in the rpm `<Lib_Utils-1.xx-xx.noarch.rpm>`.
`<Lib_Utils-1.xx-xx.noarch.rpm>` is packaged in the SNMP zip file.
- libstdc++.so.6. This library is present in `/usr/lib` directory.

You can install the SNMP application from the Linux software component RPM that provides these libraries. These RPM's are available in the Linux OS DVD.

5.6.3 Installing and Configuring an SNMP Agent on Linux

This section explains how to install and configure SAS SNMP Agent for the SUSE Linux and Red Hat Linux operating systems.

Note: This procedure requires that you have Net-SNMP agent installed on the Linux machine.

Note: The RPM has not been created to support -U version. The RPM -U will probably fail with this RPM.

To install SAS SNMP Agent, perform the following steps.

1. Install the LSI SAS SNMP Agent by using the following command:
`rpm -ivh <sas rpm>`

Note: Before installation, check whether any pass command exists that starts with 1.3.6.1.4.1.3582 OID in `snmpd.conf`. If so, delete all of the old pass commands that start with 1.3.6.1.4.1.3582 OID. (This situation could occur if an earlier version of LSI SNMP Agent was installed in the system.) After installation, find the SAS MIB file `LSI-AdapterSAS.mib` under the `/etc/lsi_mrdsnmp/sas` directory.

RPM makes the necessary modification needed in the `snmpd.conf` file to run the agent.

The `snmpd.conf` file structure should be the same as `lsi_mrdsnmpd.conf`. For reference, a sample configuration file (`lsi_mrdsnmpd.conf`) is in the `/etc/lsi_mrdsnmp` directory.

2. To run an SNMP query from a remote machine, add the IP address of that machine in the `snmpd.conf` file, as in this example:

```
com2sec      snmpclient      172.28.136.112      public
```

Here, the IP address of the remote machine is 172.28.136.112.

3. To receive an SNMP trap to a particular machine, add the IP address of that machine in the `com2sec` section of the `snmpd.conf` file.

For example, to get a trap in 10.0.0.144, add the following to `snmpd.conf`.

```
#      sec.name      source      community
com2sec snmpclient  10.0.0.144  public
```


4. To send SNMPv1 traps to a custom port, add the following configuration information to the `snmpd.conf` file:

```
Trapsink HOST [community [port] ]
```

Specify the custom port number; otherwise, the default SNMP trap port, 162, is used to send traps.

5. To start or stop the `snmpd` daemon, enter one of the following command:

```
/etc/init.d/snmpd start
```

```
/etc/init.d/snmpd stop
```

6. To start or stop the SAS SNMP Agent daemon before issuing a SNMP query, enter the following command:

```
/etc/init.d/lsi_mrdsnmpd start
```

```
/etc/init.d/lsi_mrdsnmpd stop
```

You can check the status of the SAS SNMP Agent daemon by issuing the following command:

```
/etc/init.d/lsi_mrdsnmpd status
```

7. Issue an SNMP query in the following format:

```
snmpwalk -v1 -c public localhost .1.3.6.1.4.1.3582
```

8. You can receive the SNMP trap from the local machine by issuing the following command:

```
snmptrapd -P -F "%02.2h:%02.2j TRAP#w.%q from %A %v\n"
```

Note: To receive a trap in a local machine with Net-SNMP version 5.3, you must modify the `snmptrapd.conf` file (generally located at `/var/net-snmp/snmptrapd.conf`). Add "disableAuthorization yes" in `snmptrapd.conf` and then execute "sudo snmptrapd -P -F "%02.2h:%02.2j TRAP#w.%q from %A %v\n".

Note: It is assumed that `snmpd.conf` is located at `/etc/snmp` for Red Hat and `/etc` for SuSE. You can change the file location from `/etc/init.d/lsi_mrdsnmpd` file.

You can install SNMP without the trap functionality. To do so, set the "TRAPIND" environment variable to "N" before running RPM.

Before you install a new version, you must uninstall all previous versions.

For SLES 10, perform the following steps to run SNMP:

1. Copy `/etc/snmp/snmpd.conf` to `/etc/snmpd.conf`.
2. Modify the `/etc/init.d/snmpd` file and change the `SNMPDCONF=/etc/snmp/snmpd.conf` entry to `SNMPDCONF=/etc/snmpd.conf`.
3. Run `LSI SNMP rpm`.

5.6.4 Installing an SNMP Agent on Windows

This section explains how to install and configure SAS SNMP Agent for the Microsoft Windows operating system.

5.6.4.1 Installing SNMP Agent

Perform the following steps to install SNMP Agent:

1. Run `setup.exe` from DISK1.
2. Use SNMP Manager to retrieve the SAS data (it is assumed that you have compiled `LSI-AdapterSAS.mib` file already).

The `LSI-AdapterSAS.mib` file is available under `%ProgramFiles%\LSI Corporation\SNMPAgent\SAS` directory.

3. Use a trap utility to get the traps.

Note: Before you install the Agent, make sure that SNMP Service is already installed in the system.

5.6.4.2 Installing SNMP Service for Windows

If you do not have SNMP Service installed on your system, perform the following steps to install SNMP Service for a Windows system:

1. Select **Add/Remove Programs** from Control Panel.
2. Select **Add/Remove Windows Components** in the left side of the **Add/Remove Programs** window.
3. Select **Management and Monitoring Tools**.
4. Click **Next** and follow any prompts to complete the installation procedure.

5.6.4.3 Configuring SNMP Service on the Server Side

Perform the following steps to configure SNMP Service on the server side.

1. Select **Administrative Tools** from Control Panel.
2. Select **Services** from the Administrative Tools window.
3. Select **SNMP Service** in the Services window.
4. Open **SNMP Service**.
5. Click the **Security** tab and make sure that **Accept SNMP Packets from any host** is selected.
6. Click the **Traps** tab and select the list of host IPs to which you want the traps to be sent with the community name.

5.6.4.4 Installing SNMP Service for the Windows 2008 Operating System

Before you install the LSI Agent, make sure that SNMP Service is already installed in the system.

If you do not have SNMP Service installed on your system, perform the following steps to install SNMP Service for Windows 2008 operating system.

1. Select **Program and Features** from the Control Panel.
2. Click **Turn windows feature on/off** to select the windows components to install.
3. Select **Features** from the menu.
4. Click **Add Features**.
5. Select **SNMP Services**.
6. Click **Next**.
7. Click **Install**.

The SNMP installation starts. You are prompted for the Windows 2008 CD during the installation.

8. Insert the CD, and click **Ok**.

The installation resumes.

After the installation is finished, the system displays a message saying that the installation is successful.

5.6.4.5 Configuring SNMP Service on the Server Side for the Windows 2008 Operating System

To configure SNMP service on the server side for Windows 2008 operating system, perform the following steps:

1. Select **Administrative Tools** from the **Control Panel**.
2. Select **Services** from Administrative Tools window.
3. Select **SNMP Service** from the Services window.
4. Open **SNMP Service**, and go to its properties.
5. Go to the **Security** tab, and make sure that **Accept SNMP Packets from any host** is selected.
6. Click the **Traps** tab, and select the list of host IP addresses to which you want the traps to be sent with the community name.

Chapter 6

MegaRAID Storage Manager Screens and Menus

This chapter explains how to start the MegaRAID Storage Manager configuration utility, and it describes the MegaRAID Storage Manager screens and menus. This chapter has the following sections:

- [Section 6.1, “Starting MegaRAID Storage Manager Software”](#)
 - [Section 6.2, “Discovery and Login”](#)
 - [Section 6.3, “LDAP Support”](#)
 - [Section 6.4, “Configuring LDAP Support Settings”](#)
 - [Section 6.5, “MegaRAID Storage Manager Main Menu Screen”](#)
-

6.1 Starting MegaRAID Storage Manager Software

You must have administrative privileges in order to use MegaRAID Storage Manager software in either full-access mode or view-only mode. Start MegaRAID Storage Manager using the method required for the Windows operating system or Linux operating system:

- To start MegaRAID Storage Manager software on a Microsoft Windows system, select **Start >> Programs >> MegaRAID Storage Manager >> StartupUI**, or double-click the MegaRAID Storage Manager shortcut on the desktop.

Note: If a warning appears stating that Windows Firewall has blocked some features of the program, click **Unblock** to allow MegaRAID Storage Manager software to start. (The Windows Firewall sometimes blocks the operation of programs that use Java.)

- To start MegaRAID Storage Manager software on a Red Hat Linux system, select **Applications >> System Tools >> MegaRAID Storage Manager >> StartupUI**.

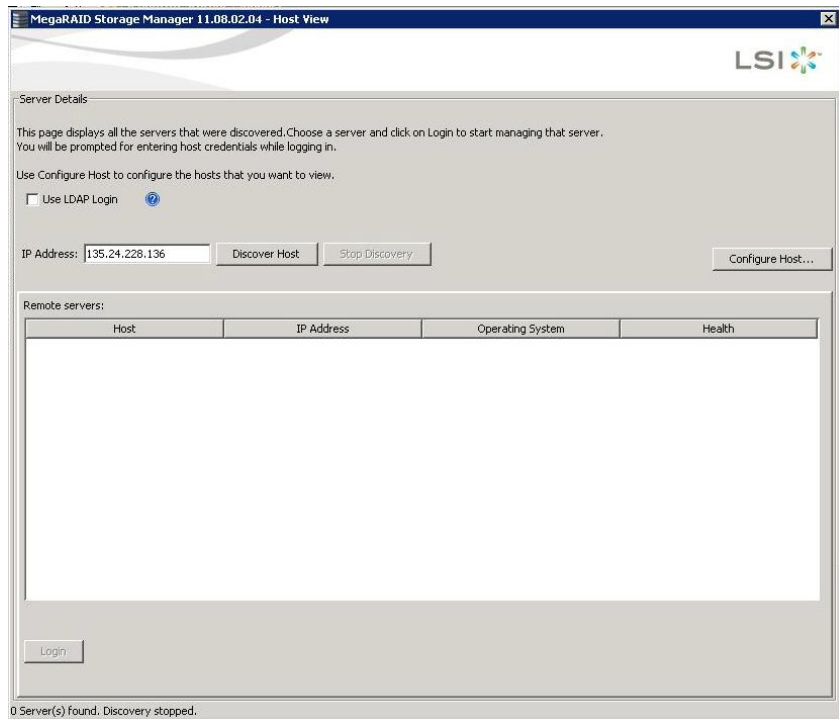
- To start MegaRAID Storage Manager software on a SUSE SLES system, select **Start >> System >> More Programs >> MegaRAID Storage Manager**.

6.2 Discovery and Login

You can start the MegaRAID Storage Manager software from a remote Windows/Linux machine that has the MegaRAID Storage Manager software installed in complete mode. When the program starts, the **Select Server** dialog appears, as shown in the following figure. The remote servers are displayed, along with their IP addresses, operating system, and health status.

If you do a local mode installation, as shown in Section Installing MegaRAID Storage Manager Software on Microsoft Windows, the following figure does not appear. You are prompted to the login dialog as shown in the Server Login.

Figure 6.1 Select Server



The **Select Server** dialog shows an icon for each server on which the MegaRAID Storage Manager software is installed. The servers are color-coded with the following definitions:

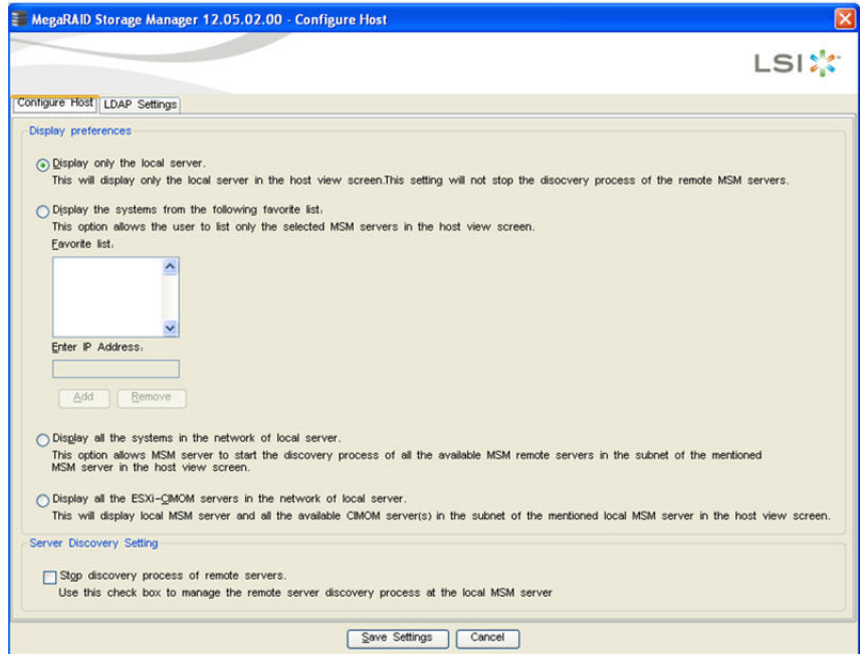
- Green: The server is operating properly.
- Yellow: The server is running in a partially degraded state (possibly because a drive in a virtual drive has failed).
- Orange: The server is running in a degraded state.
- Red: The server storage configuration has failed.

Note: Do not enter the VMware ESXi server's **IP address** in the IP Address field in the previous figure. Instead enter a valid MegaRAID Storage Manager server's IP address and select the **Display all the systems in the Network of the local server** option in the following figure.

1. Click **Configure Host** to select which systems to display.

The **Configure Host** dialog appears, as shown in the following figure.

Figure 6.2 Configure Host Window



2. Select any one of the following options from **Display preferences**.

—**Display only the local server** – Discovers only the local MegaRAID Storage Manager server.

—**Display the systems from the following favorite list** – Allows you to enter IP addresses of the MegaRAID Storage Manager servers and discovers only those servers. You can enter an IP address in the Enter IP Address field and click Add. The server corresponding to the IP address appears in the Favorite list.

—**Display all the systems in the network of the local server** – Discovers all the MegaRAID Storage Manager servers available in the network.

—**Display all the ESXi-CIMOM servers in the network of local server** - Discovers the local MegaRAID Storage Manager server and all the available ESXi servers in the network.

Note: On some Windows machines, the discovery of VMware ESXi servers fail as a result of a bug in the third-party application that is used for discovery. This failure is caused by one of the Windows servers in the network that contains a service called IBM SLP SA, which gets installed along with the IBM Director. If you stop this service on all the Windows servers in the network, the MegaRAID Storage Manager can discover all the ESXi servers.

3. Click **Save Settings** to save your setting.

A dialog box appears and asks you to confirm your settings.

4. Click **OK** in the confirmation dialog to start the discovery process.

After the discovery process is completed, the servers appear in the **Select Server** dialog.

To abort the discovery process which has already begun, select the **Stop discovery process of remote servers** check box and click **Save Settings**. This check box is enabled only when there is a active discovery process.

Note: For the VMware ESXi, the server icon does not denote the health of the server. The icon is always green regardless of the health of the system. The VMware server does not show the system health and the operating system labels. It shows only the host name and the IP address of the server. When connecting to a VMware server on a different subnet, one or more frameworks have to be running in the subnet to connect to the CIMOM.

5. Double-click the icon of the server that you want to access.

The **Server Login** window appears, as shown in the following figure.

Figure 6.3 Server Login Window

Enter User Name & Password

LSI

Server : 135.24.228.178

Use your Operating System's login username and password to login the MSM server

User Name: ?

Password:

Login Mode: Full Access

Login Cancel

6. Enter the root account name and the password of the host in the **User Name** and **Password** fields, respectively.

Note: In the **User Name** field, you can also enter the domain name along with the user name; for example, LSI\abc, where LSI is the domain name and abc is your user name.

The question mark icon opens a dialog box that explains what you need for full access to the server and for view-only access to the server. You are allowed three attempts to log in.

Note: When connected to VMware operating system, the **Server Login** window shows only one label for access, Full Access. Multiple users can have full access to the VMware server.

7. Select an access mode from the drop-down menu for **Login Mode**, and click **Login**.
 - Select **Full Access** if you need to both view and change the current configuration.
 - Select **View Only** if you need to only view and monitor the current configuration.

Note: If the computer is networked, this login is for the computer itself; it is not the network login.

8. Enter the root or administrator user name and password to use Full Access mode.

Note: In Linux, users belonging to the root group can log in. You do not have to be the user root.

If your user name and your password are correct for the Login mode you have chosen, the MegaRAID Storage Manager main menu screen appears.

6.3 LDAP Support

The MegaRAID Storage Manager application supports the use of the Lightweight Directory Access Protocol (LDAP) to discover remote MegaRAID Storage Managers servers. To enable LDAP support, the MegaRAID Storage Manager servers must be registered with the LDAP server.

Note: LDAP supports only Windows Active Directory LDAP Server Implementation.

Note: VMware ESXi servers are not discovered during LDAP discovery.

To register the MegaRAID Storage Manager servers with the LDAP server, define a new attribute, `ou`, on the machine on which the LDAP server is configured, and give this attribute the value `MSM`. This registration enables the discovery of only the MegaRAID Storage Manager servers that have been registered with the LDAP server.

Perform the following steps to use LDAP support:

1. Double-click the MegaRAID Storage Manager software shortcut icon on your desktop.

The **Select Server** dialog appears.

2. Select the **Use LDAP Login** check box, and click **Discover Host**.

All the MegaRAID Storage Manager servers that are registered with the LDAP server appear in the **Remote servers** box.

Note: If the **Use LDAP Login** check box is selected, the **IP Address** field is disabled.

3. Click on a server link to connect to the LDAP server.

Note: Based on the privileges allotted to you, the MegaRAID Storage Manager servers are launched with full access rights or read-only rights.

If you select the **Do not prompt for credentials when connecting to LDAP** check box in the LDAP Settings tab on the Configure Host dialog (see [Figure 6.5](#)), you are connected directly to the LDAP server. Otherwise, the LDAP Login dialog appears, as shown in the following figure.

Figure 6.4 LDAP Login



The screenshot shows the 'LDAP Login' dialog box. At the top, there's a title bar with 'LDAP Login' and a close button. Below it is the LSI logo. The main area contains several input fields: 'LDAP Server IP Address', 'User Name', 'Password', and 'Distinguished Name'. To the right of the 'User Name' field is a help icon. Below the 'Password' field is a 'Port' field with '389' entered. There are two checkboxes: 'Use Default Port' (checked) and 'Remember my Login Details' (unchecked). At the bottom are 'Login' and 'Cancel' buttons.

Follow these steps to enter the LDAP login details:

1. Enter the IP address of the LDAP server in the **LDAP Server IP Address** field.
2. Enter the LDAP server's user name and password in the **User Name** and **Password** fields, respectively.

An example of a user name can be `username@testldap.com`.

3. Enter the name of the Domain Controller in the **Distinguished Name** field.

As an example, the Domain Controller name can be `dc=TESTLDAP, dc=com`.

Note: The **LDAP Server IP Address**, **User Name**, **Password**, and **Distinguished Name** fields are already populated if their corresponding values have been stored in the LDAP Settings tab on the **Configure Host** dialog (see [Figure 6.5](#)).

4. Perform one of these actions:
 - If you want to use the default port number, select the **Use Default Port** check box. The default port number, 389, appears in the **Port** field.
 - If you do not want to use the default port number, uncheck the **Use Default Port** check box, and enter a port number in the **Port** field.
5. Select the **Remember my Login Details** check box if you want to save all the values entered in this dialog in the LDAP Settings tab in the **Configure Host** dialog.
6. Click **Login** to log in to the LDAP server.

6.4 Configuring LDAP Support Settings

Perform the following steps to configure settings for LDAP support.

1. Navigate to the **Configure Host** dialog, and click the **LDAP Settings** tab.

The following screen appears.

Figure 6.5 Configuring Host LDAP

MegaRAID Storage Manager 11.08.02.04 - Configure Host

LSI

Configure Host: **LDAP Settings**

☐ Use LDAP login as default login mode

☐ Do not prompt for credentials when connecting to LDAP

Server

IP Address: Port:

Distinguished Name :

Connection

User Name :

Password :

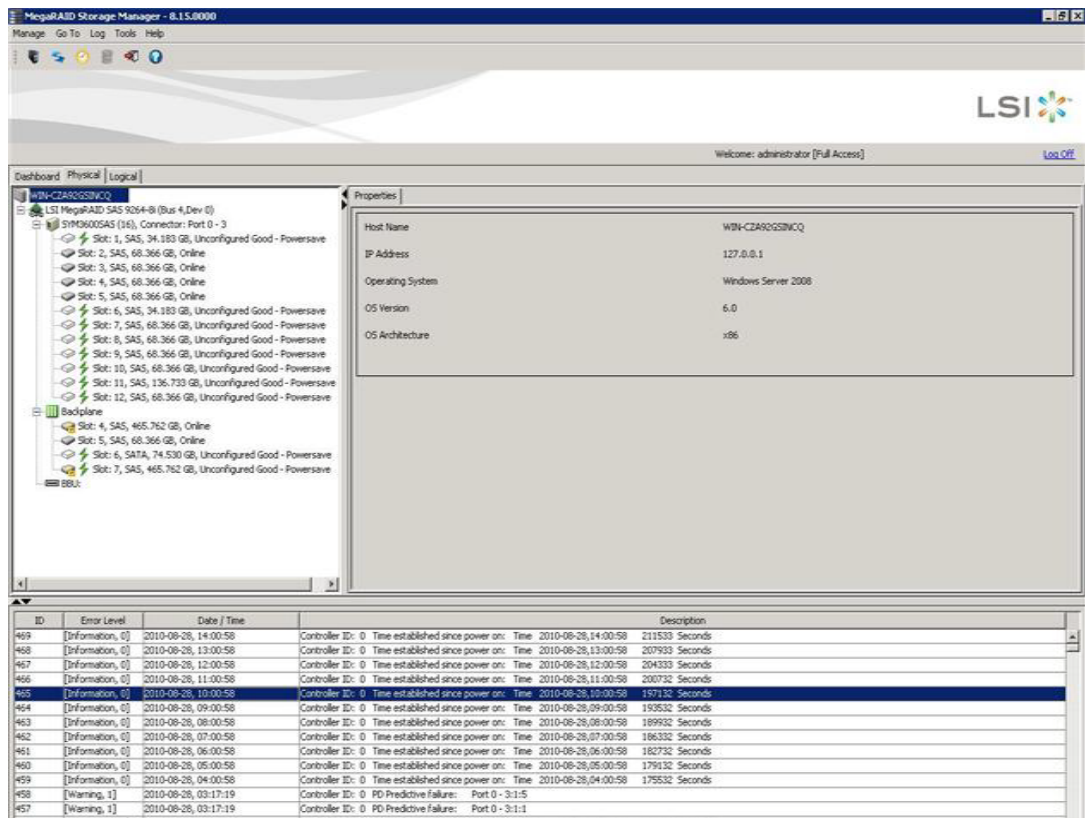
Save Settings Cancel

2. Select the **Use LDAP login as default login mode** check box to always connect to the LDAP server.
3. Select the **Do not prompt for credentials when connecting to LDAP** check box if you do not want the LDAP Login dialog (see [Figure 6.4](#)) to appear when you connect to the LDAP server.
4. Enter the IP address of the LDAP server in the **IP Address** field.
5. Enter the port number in the **Port** field.
6. Enter the name of the Domain Controller in the **Distinguished Name** field.
7. Enter the user name and password for logging into the LDAP server in the **User Name** and **Password** fields, respectively.
8. Click **Save Settings** to save all the values entered in the fields in the `msm.properties` file.

6.5 MegaRAID Storage Manager Main Menu Screen

This section describes the MegaRAID Storage Manager main menu screen, which is shown in the following figure.

Figure 6.6 MegaRAID Storage Manager Main Screen



The following topics describe the panels and menu options that appear in this screen.

6.5.1 Dashboard/Physical View/Logical View

The left panel of the MegaRAID Storage Manager window displays the *Dashboard*, the *Physical* view or the *Logical* view of the system and the attached devices, depending on which tab is selected.

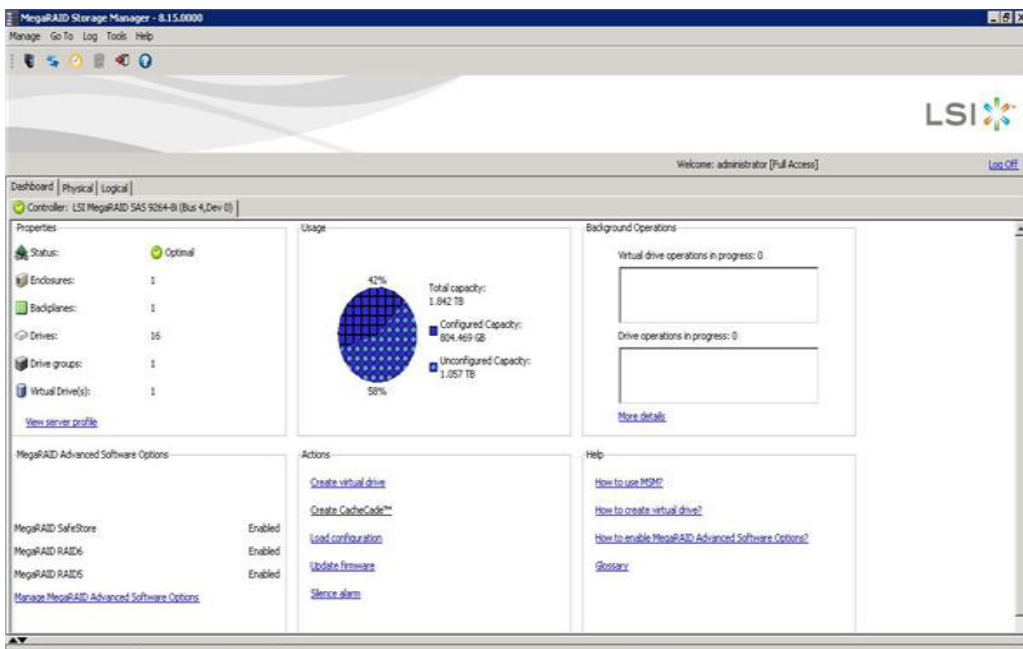
6.5.2 Dashboard

The *Dashboard* view shows an overview of the system and covers the following features:

- Properties of the virtual drives and the physical drives
- Total capacity, configured capacity, and unconfigured capacity
- Background operations in progress
- MSM features and their status (enabled or disabled)
- Actions you can perform, such as creating a virtual drive and updating the firmware
- Links to Online Help

The following figure shows the Dashboard view.

Figure 6.7 MSM Dashboard View

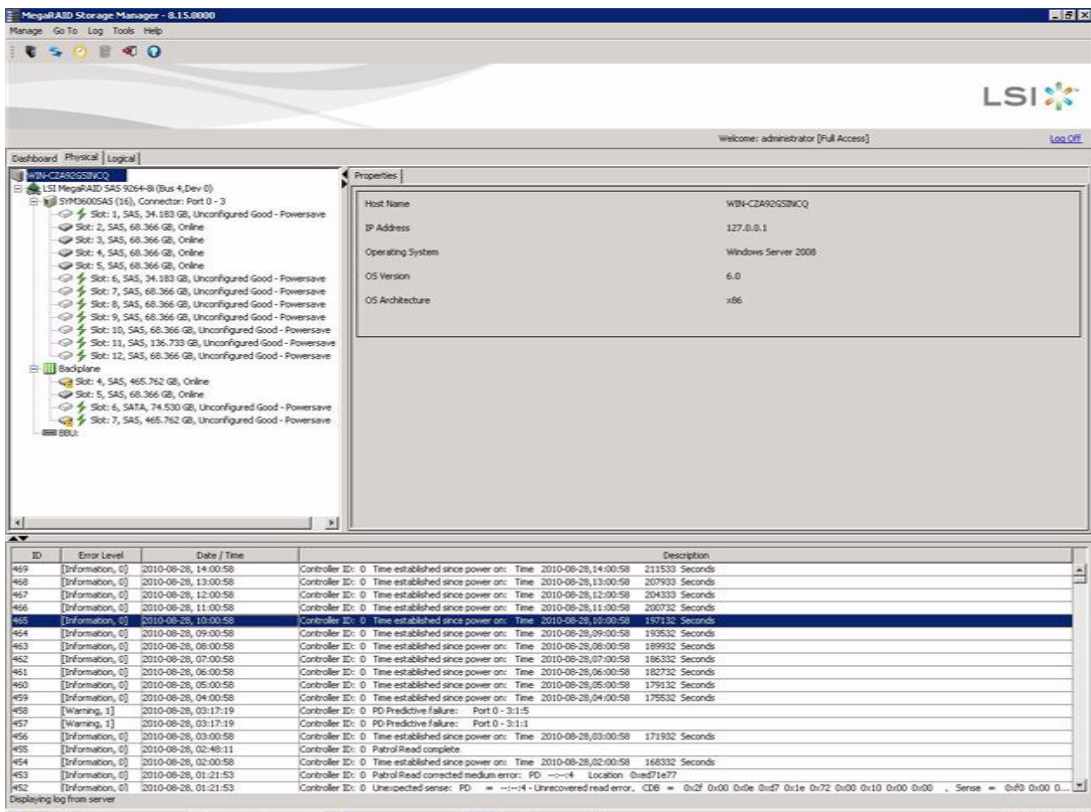


6.5.3 Physical View

The *Physical* view shows the hierarchy of physical devices in the system. At the top of the hierarchy is the system itself. One or more controllers are installed in the system. The controller label identifies the ServerRAID controller so that you can easily differentiate between multiple controllers. Each controller has one or more ports. Drives and other devices are attached to the ports. The properties for each item appear in the right panel of the screen.

The following figure shows the Physical view.

Figure 6.8 MSM Physical View

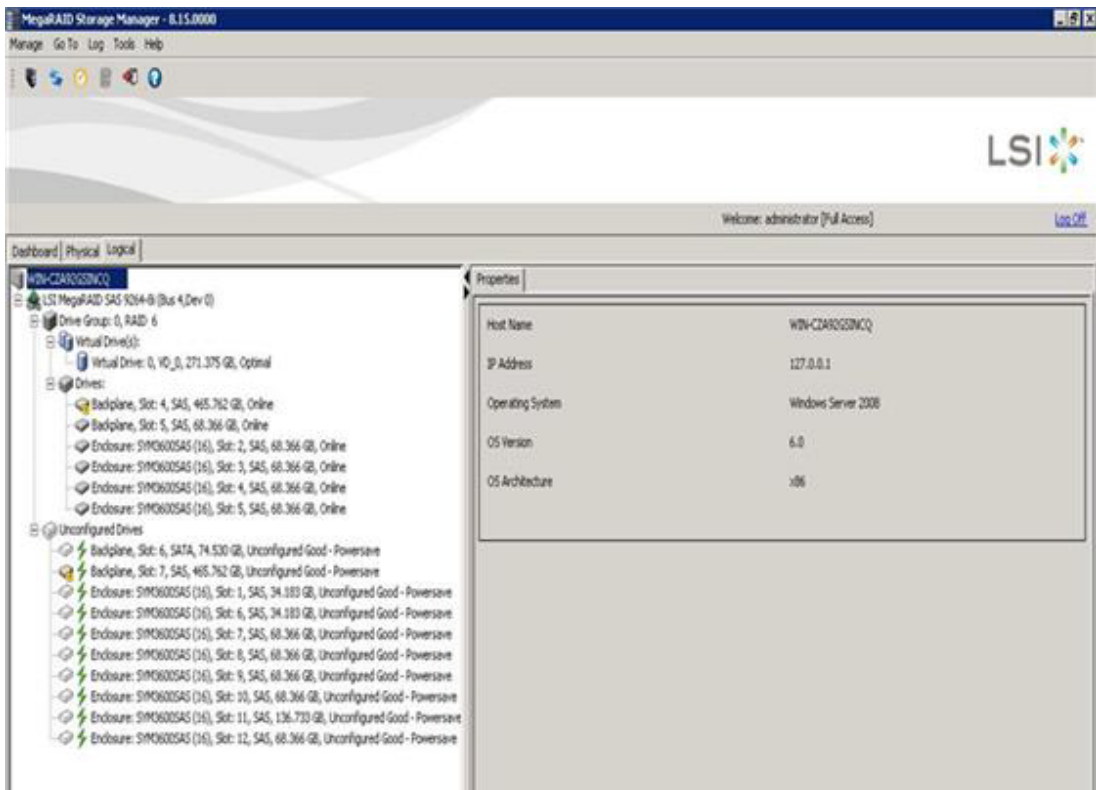


6.5.4 Logical View


The *Logical* view shows the hierarchy of controllers, virtual drives, drives, and drive groups that make up the virtual drives. The properties for these components appear in the right panel. (Drives also appear in the *Logical* view, so you can see which drives are used by each virtual drive.)













The following figure shows the Logical view.

Figure 6.9 Logical View





The following icons in the left panel represent the controllers, drives, and other devices:

-  - State
-  - System

-  - Controller
-  - Backplane
-  - Enclosure
-  - Port
-  - Drive group
-  - Virtual drive
-  - Drive
-  - Dedicated hot spare
-  - Global hot spare
-  - Battery backup unit (BBU)
-  - Tape drive
-  - CD-ROM

Note: MegaRAID Storage Manager shows the icons for tape drive devices; however, no tape-related operations are supported by the utility. If these operations are required, use a separate backup application.

A red circle to the right of an icon indicates that the device has failed. For example, this icon indicates that a drive has failed: .

A yellow circle to the right of an icon indicates that a device is running in a degraded state. For example, this icon indicates that a virtual drive is running in a degraded state because a drive has failed: .

An orange circle to the right of an icon indicates that a device is running in a degraded state.

6.5.4.1 Controller Properties

In the Physical view and the Logical view, you can view the chip temperature and controller temperature under the controller properties for the controller, as shown in the following figure.

Figure 6.10 Chip and Controller Temperature

The screenshot displays the MegaRAID Storage Manager interface. The left sidebar shows a tree view of the hardware configuration, including the RAID controller and its associated drives. The main pane shows the 'Properties' tab for the selected controller. The properties are organized into several sections: Coercion Mode, BBU Present, NVRAM Present, NVRAM Size, BIOS Version, Native Command Queuing, Flash Size, Chip Temperature, and Controller Temperature. The 'Power State Properties' section is also visible. The 'Chip Temperature' and 'Controller Temperature' are both shown as 36 C (96.8 F) and 31 C (87.8 F) respectively.

Property	Value
Coercion Mode	128 MB
BBU Present	No
NVRAM Present	Yes
NVRAM Size	32,000 KB
BIOS Version	4.24.00_4.11.05.00_0x05020000
Native Command Queuing	Enabled
Flash Size	16,000 MB
Chip Temperature	36 C (96.8 F)
Controller Temperature	31 C (87.8 F)

Power State Properties:

Property	Value
Power savings on unconfigured drives	Disabled
Power savings on hot spares	Enabled

Drive Security Properties:

Property	Value
Cluster Enable	No
Cluster Active	No
SSD Guard	Disable
Drive security capable	No


Background Operation Properties:

Property	Value
Rebuild Rate	30
Patrol Read Rate	30
Reconstruction Rate	30
BGI Rate	30
Consistency Check Rate	30

6.5.4.2 Physical Drive Temperature

The following figure shows the temperature for the physical drive.

Figure 6.11 Physical Drive Temperature



6.5.5 Shield State

This section describes the Shield state in the MegaRAID Storage Manager software.

Physical devices in MegaRAID firmware transit between different states. If firmware detects a problem or a communication loss for a physical drive, it transitions the physical drive to a bad (FAILED/UNCONF BAD) state. To avoid transient failures, an interim state called the Shield state appears before a physical drive is changed to a bad state.

The Shield state is an interim state of a physical drive for diagnostic operations. The results of the diagnostic tests determine if the physical drive is good or bad. If any of the diagnostics tests fail, the physical drive will transition to BAD state (FAILED or UNCONF BAD).

The three possible Shield states are **Unconfigured - Shielded**, **Configured - Shielded**, and **Hotspare - Shielded**.

6.5.5.1 Shield State Physical View

To view the Shield state under the Physical view tab, click the **Physical** tab in the device tree.


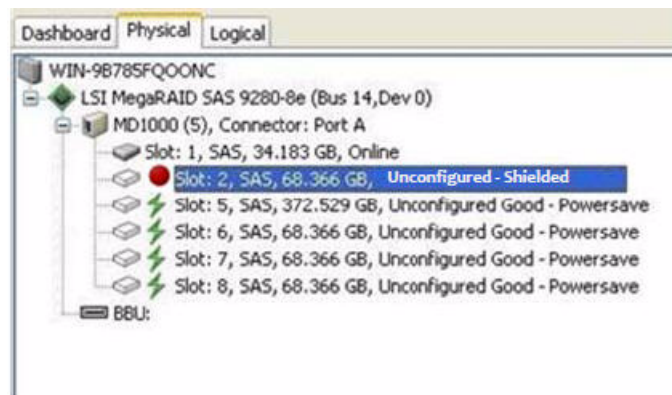
The  icon indicates a Shield state, as shown in the following figure.

Figure 6.12 Physical View Shield State



6.5.5.2 Logical View Shield State

Follow these steps to view the Shield state under the **Logical** tab.



1. Click the **Logical** tab in the device tree.
2. The  icon indicates a Shield state, as shown in the following figure.

Figure 6.13 Logical View Shield State



6.5.5.3 Viewing the Physical Drive Properties

Follow these steps to view the physical properties of the drive in the Shield state.


1. Click the **Physical** tab or **Logical** tab in the device tree.
The  icon indicates a Shield state.
2. Click the physical drive in Shield state on the Physical view or Logical view of the device tree to view the properties, as shown in the following figure.

Figure 6.14 Physical Drive Properties of a Drive in Shield State

Properties			
General:		Temperature	31 C (87.8 F)
Usable Capacity	67.044 GB	Power Status	Powersave
Raw Capacity	68.366 GB	Revision Level	A48B
Certified	No	Media Error Count	0
Product ID	HU5151473VL5300	Pred Fail Count	0
Vendor ID	HITACHI	Enclosure Properties	
Device ID	14	Enclosure ID	5
Status	Unconfigured - Shielded	Enclosure Model	MD1000
Drive Speed	3.0 Gbps	Enclosure Location	External
Negotiated Link Speed	3.0 Gbps	Connector	Port A
SCSI Device Type	Disk	Slot Number	2
SAS Address 0	0x5000CCA0074D2718	Drive Security Properties:	
SAS Address 1	0x0	Full Disk Encryption capable	No

Note: The Status of the drive must be of the Shield type.

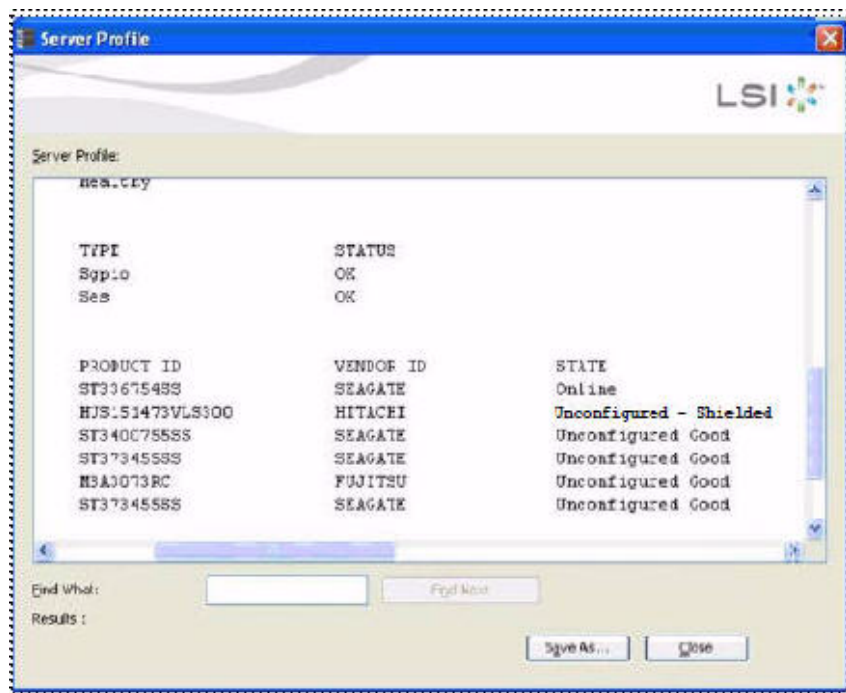
6.5.5.4 Viewing Server Profile of a Drive in Shield State

Follow these steps to view the server properties of the drive in Shield state.

1. Click the **Dashboard** tab in the device tree.
2. Click the **View Server Profile** link in the dashboard view.

The server profile information appears, as shown in the following figure.

Figure 6.15 Server Profile View of a Drive in Shield State



6.5.6 Displaying the Virtual Drive Properties

The MegaRAID Storage Manager application displays the following additional virtual drive statistics under controller properties.

- Parity size
- Mirror date size

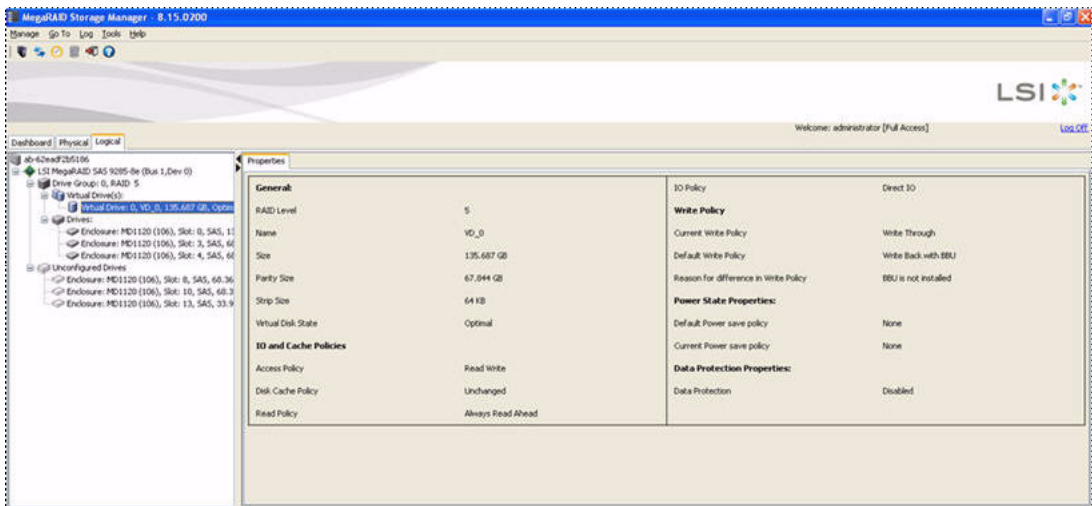
- Metadata size

6.5.6.1 Parity Size

Parity size is used for storing parity information on RAID 5, RAID 6, RAID 50, and RAID 60 virtual drives. Follow these steps to view the Parity Size.

1. In the Logical view, click the **Virtual Drive** node.
2. For RAID 5, RAID 6, RAID 50, and RAID 60, the **Parity Size** appears, as shown in the following figure.

Figure 6.16 Parity Size




6.5.6.2 Mirror Data Size

Mirror Data Size is used to determine the size used for storing redundant information on RAID 1 and RAID 10 virtual drives. Follow these steps to view the **Mirror Data Size**.

1. In the **Logical** view, click on the Virtual Drive node.

The Mirror data size displays for RAID 1 and RAID 10 volumes, as shown in the following figure.

Figure 6.17 Mirror Data Size



Note: The parity size and mirror data size are not displayed for RAID 0 volumes.

6.5.6.3 Metadata Size

The metadata size field displays the total space used for metadata. Follow these steps to view the Metadata Size.

1. In the Logical view or the Physical view, click the controller node.
The total space used for metadata appears in this field, as shown in the following figure.

Figure 6.18 Metadata Size

Properties			
Alarm Present	Yes	Backend SAS Address 7	0x0
Alarm Enabled	No	Correctable Error Count	0
Cache Flush Interval	4 sec	Memory uncorrectable count	0
Coordin Mode	None	Cluster Enable	No
BBU Present	Yes	Cluster Active	No
NVRAM Present	Yes	SSD Guard	Enabled
NVRAM Size	32,000 KB	Drive Security Properties:	
BIOS Version	3.18.00_4.09.05.00_0x04260000	Drive security capable	No
Native Command Queuing	Enabled	Background Operation Properties:	
Flash Size	8,000 MB	Rebuild Rate	55
Memory Size	512,000 MB	Patrol Read Rate	10
Metadata Size	500 MB	Reconstruction Rate	10
Power State Properties:		ERL Rate	14
Power savings on unconfigured drives	Enabled	Consistency Check Rate	16
Power savings on hot spares	Enabled	MegaRAID Recovery Properties:	
Power Save Policy for Configured Drives	Auto	MegaRAID Recovery	Enabled
Drive Standby Time	43mins		
Firmware Properties:			

Note: The size units appears as follows:

- If the size is less than 1 MB (1024 KB), the size is displayed in KB.
- If the size is greater than or equal to 1 MB but less than 1 GB (1024 MB), the size is displayed in MB.
- If the size is greater than or equal to 1 GB, but less than 1 TB (1024 GB), the size displays in GB.

6.5.7 SSD Disk Cache Policy

The MegaRAID firmware provides support to change the write-cache policy for SSD media of individual physical drives.

The MegaRAID firmware does not allow any user application to modify the write-cache policies of any SSD media. The host applications can modify this property through a new logical device (LD) addition or a LD property change. When SSDs are configured in a mixed disk group with

HDDs, the Physical Device Write-Cache Policy setting of all the participating drives are changed to match the SSD cache policy setting.

Follow these steps to view the SSD cache property.

1. Click the controller node in the device tree.

The Controller Properties screen appears, as shown in the following figure.

Figure 6.19 Controller Properties – SSD Disk Cache Policy

Properties			
Host Port Count	0	Backend SAS Address 6	0x0
FRU		Backend SAS Address 7	0x0
Alarm Present	Yes	Correctable ErrorCount	0
Alarm Enabled	Yes	Memory uncorrectable count	0
Cache Flush Interval	4 sec	Cluster Enable	No
Coercion Mode	None	Cluster Active	No
BBU Present	No	SSD Guard	Enabled
NVRAM Present	Yes	SSD Disk Cache Setting	Disabled
NVRAM Size	32,000 KB	Drive Security Properties:	
BIOS Version	3.18.00_4.09.05.00_0x0416A000	Drive security enabled	No
Native Command Queuing	Enabled	Drive security method	FDE Only
Flash Size	8,000 MB	Drive security capable	Yes
Memory Size	256,000 MB	EKM Supported	Yes
Power State Properties:		Key Management Mode	N/A
Power savings on unconfigured drives	Enabled	Background Operation Properties:	
Power savings on hot spares	Enabled	Rebuild Rate	30
Drive Standby Time	30mins	Patrol Read Rate	30
Firmware Properties:		Reconstruction Rate	30
Firmware Package Version	12.10.0-0015	BGI Rate	30
		Consistency Check Rate	30

6.5.7.1 Virtual Drive Settings

If the SSD cache property is enabled in the controller properties screen as shown, in [Figure 6.19](#), then you cannot select the disk cache policy for the virtual drives that have only SSD drives or a mix of SSD drives


and HDD drives during virtual drive creation. The value of the disk cache policy is unchanged and the drop-down menu is disabled.

Follow these steps to view the **Virtual Drive Settings**.

1. Right-click the controller node in the device tree.
2. Select the **Create Virtual Drive** menu option.
3. Select **Advanced Configuration**, and click **Next**.
4. Create **Drive Group**, and click **Next**.

The Create Virtual Drive – Virtual drive settings dialog appears, as shown in the following figure.

Figure 6.20 Virtual Drive Settings



The value of the disk cache policy is Unchanged, and the drop-down list is disabled.

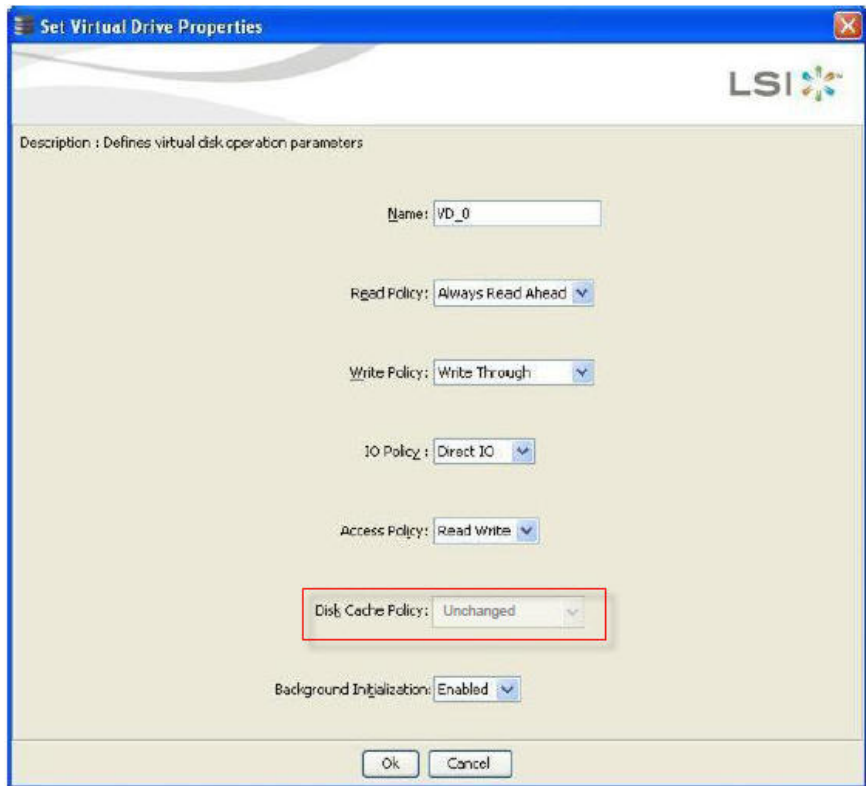
6.5.7.2 Set Virtual Drive Properties

Follow these steps to set virtual drive properties.

1. Right click on virtual drive node in the logical view of the device tree.
2. Select **Set Virtual Drive Properties**.

The Set Virtual Drive Properties dialog appears, as shown in the following figure.

Figure 6.21 Virtual Drive Properties



Note: You cannot select the Disk cache policy for the virtual drives having only SSD drives or a mix of SSD and HDD during VD creation. The value of the Disk Cache Policy is Unchanged and can be set for only HDD drives.

6.5.8 Non-SED Secure Erase Support

This section describes the firmware changes required to securely erase data on non-SEDs (normal HDDs).

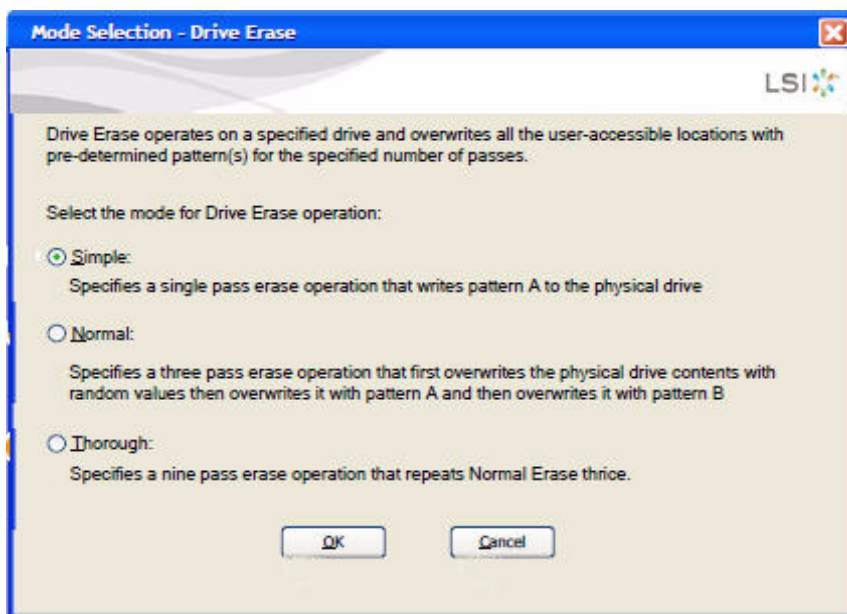
SEDs securely erase their internal encryption keys, effectively destroying all of the data present on the drive. For Non-SED drives, the erase operation consists of a series of write operations to a drive that overwrites every user-accessible sector of the drive with specified patterns. It can be repeated in multiple passes using different data patterns for enhanced security. The sanitization technique is more secure than a simple format operation and is commonly called a “clearing” operation, similar to the existing physical drive clear command.

Follow these steps to set physical drive properties.

1. In the Physical view, right click the **Physical Drive** node.
2. Select the **Drive Erase** option (**Alt+E**).

The **Mode Selection - Drive Erase** dialog box appears, as shown in the following figure.

Figure 6.22 Mode Selection - Drive Erase Window

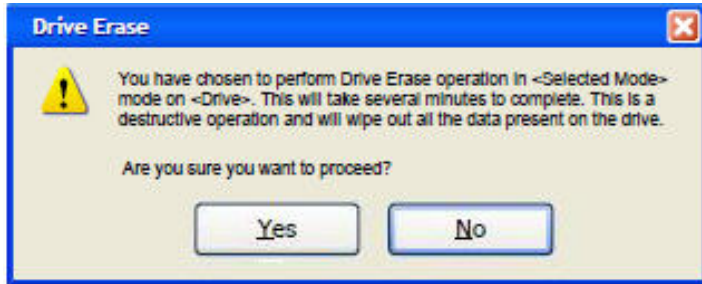


3. Select one of the available modes.

- **Simple** (Alt + S)

When you select this option and click **OK**, the Drive Erase message box appears, as shown in the following figure.

Figure 6.23 Drive Erase Message



- **Normal** (Alt + N)

Select this option and click **OK**. [Figure 6.23](#) appears.

- **Thorough** (Alt + T)

Select this option and click **OK**. [Figure 6.23](#) appears.

6.5.8.1 Group Show Progress

Physical drive erase operation is a time-consuming operation and is performed as a background task. It posts events to notify users of the progress.

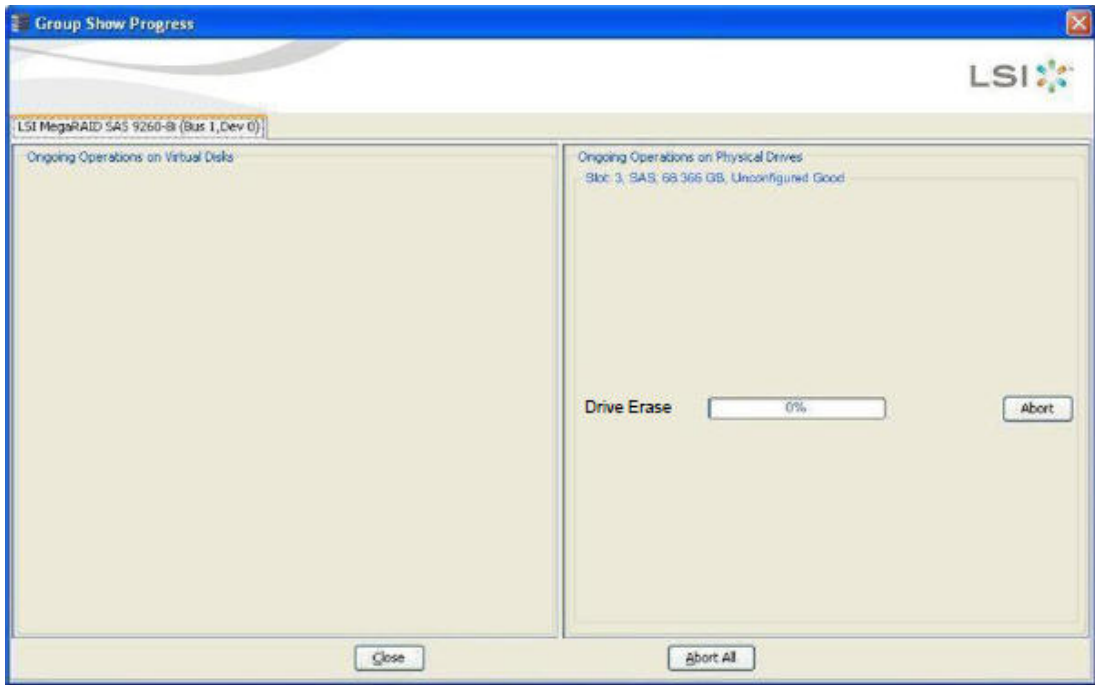
Follow these steps to check the progress of physical drive erase operation.

1. Click the **Show Progress** toolbar icon in the MegaRAID Storage Manager. You can also select **Show Progress** from the dashboard or select **Show Progress** from the Manage menu.

2. Click the **More info** link under the Background Operations portlet.

The progress bar appears, as shown in the following figure.

Figure 6.24 Group Show Progress



When you click the **Abort All** button, all Drive Erase operations stop, and the progress bar is not displayed.

6.5.8.2 Virtual Drive Erase

Virtual drive erase operates on a specified virtual drive and overwrites all user-accessible locations. It supports non-zero patterns and multiple passes. Virtual drive erase optionally deletes the virtual drive and erases the data within the virtual drive's LBA range. Virtual drive erase is a background operation, and it posts events to notify users of their progress.

Follow these steps to open the Virtual Drive Erase menu.

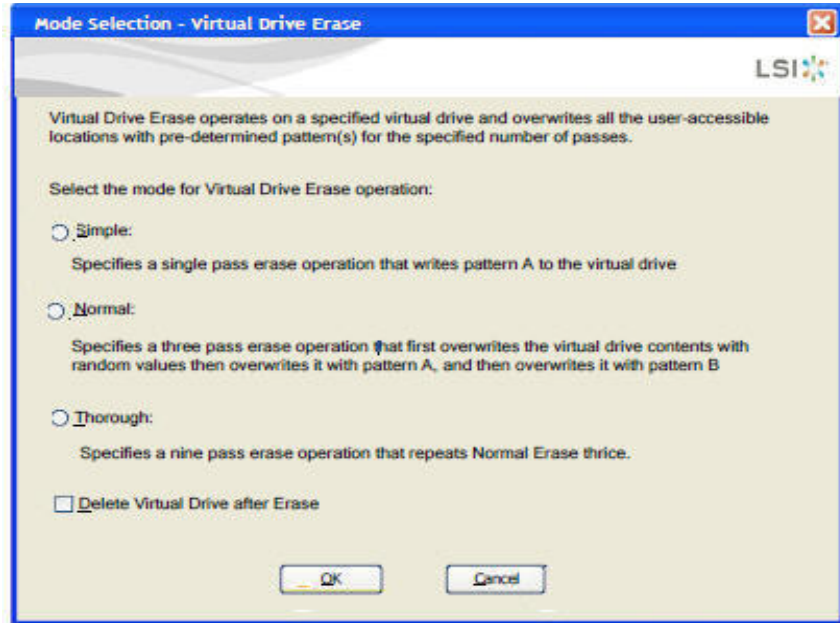
1. In the Logical view, right-click the Virtual Drive node.
2. Click on the Virtual Drive node, select top level navigation and click **Go to**.
3. Select **Virtual Drive** and select **Events & Response**.

The Logical View - Virtual Drive Erase menu appears.

4. Select **Virtual Drive Erase**.

The Virtual Drive Erase Menu opens, as shown in the following figure.

Figure 6.25 Mode Selection – Virtual Drive Erase Dialog




The menu has the following options.

- **Simple (Alt + S)** – After you select this option and click **OK**, and if **Delete Virtual Drive after Erase** is selected, [Figure 6.26](#) appears; otherwise, [Figure 6.27](#) appears.
- **Normal (Alt + N)** – After you select this option and click **OK**, and if **Delete Virtual Drive after Erase** is selected, [Figure 6.26](#) appears; otherwise, [Figure 6.27](#) appears.
- **Thorough (Alt + T)** – After you select this option and click **OK** and if **Delete Virtual Drive after Erase** is selected, [Figure 6.26](#) appears; otherwise, [Figure 6.27](#) appears.
- **Delete Virtual Drive after Erase (Alt + D)** – When you select this option, the virtual drive is erased and [Figure 6.26](#) appears; otherwise, [Figure 6.27](#) appears.


- **OK (Alt + O)** – Click **OK** and if **Delete Virtual Drive after Erase** is checked, the following figure appears; otherwise, [Figure 6.27](#) appears.
- **Cancel (Alt + C)** – When you select this option, the dialog closes, and the MegaRAID Storage Manager navigates back to Physical view.

Figure 6.26 Warning Message for Virtual Drive Erase



- Click **Yes** to erase the virtual drive.
- Click **No** to cancel the erase and close the dialogue.

Figure 6.27 Warning Message for Virtual Drive Erase Without Virtual Drive Delete



- Click **Yes** to erase the virtual drive.
- Click **No** to cancel the erase and close the dialogue.

6.5.8.3 Group Show Progress for Virtual Drive Erase


The virtual drive erase operation is a time-consuming operation and is performed as a background task. It posts events to notify users of the progress.

To view the progress of Group Show Progress-Virtual Drive, click the **Show Progress** toolbar icon.

You can also either select **Show Progress** from the Manage menu, or select the **More info** Link under Background Operations portlet on the dashboard.

The Virtual Drive Erase progress bar appears, as shown in the following figure.

Figure 6.28 Group Show Progress – Virtual Drive



6.5.9 Rebuild Write Cache

MegaRAID firmware supports drive cache properties during a rebuild operation. The MegaRAID solution temporarily enables drive cache for the physical drive that is being rebuilt for the duration of the rebuild operation. Users can enable or disable this feature using the Mega CLI feature.

The MegaRAID software automatically changes the setting for a drive that is being rebuilt. If the PD_CACHE for the rebuilt drive is already set, the firmware does not need to do anything extra.

The firmware identifies and sets the cache policy of the drives whenever a rebuild operation starts and the cache policy is reflected in the event logs. The firmware also makes sure to flush the cache just before committing the drive to the disk group.

6.5.10 Background Suspend or Resume Support

MegaRAID provides a background Suspend or Resume Support feature that enhances the functionality where in the background operations running on a physical drive or a virtual drive can be suspended for some time, and resumed later using the Resume option.

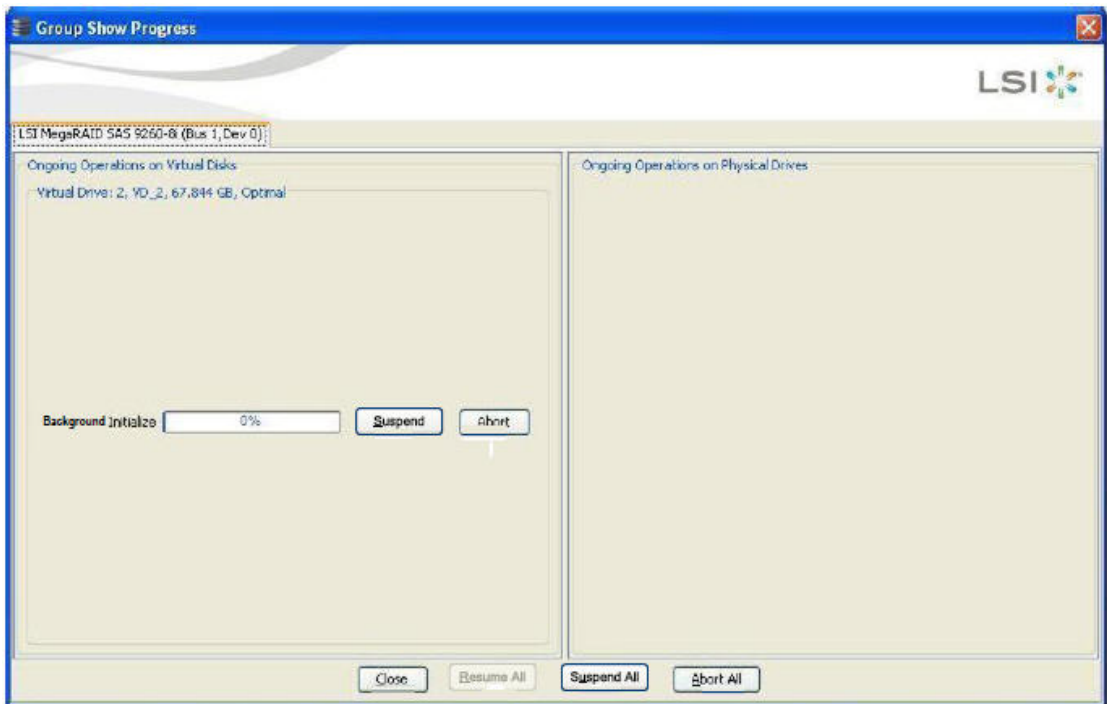
The background operations, including consistency-check, rebuild, background initialization, and patrol read, are supported by an abort operation. If any operation is stopped before completion, it is considered to be aborted. An aborted operation cannot be resumed from the place where it was stopped.

A suspended operation can be resumed later by using the Resume option, and the suspended operation resumes from the point where the operation was suspended last.

To perform a suspend and resume operation, go to the Group Show Progress dialog, and perform the tasks mentioned below. You also can select **Show Progress** from the Manage menu, or select the **More info** link under the Background Operations portlet on the dashboard

The Group Show Progress dialog appears, as shown in the following figure.

Figure 6.29 Group Show Progress Dialog




- **Suspend (Alt + S)** – Click the **Suspend** button to suspend the background operation taking place at that particular point of time. When the operations gets suspended, the **Resume** button appears instead of the **Suspend** button.
- **Resume (Alt + E)** – Click the **Resume** button to resume the operation from the point where it was suspended last.
- **Abort (Alt + B)** – Click the **Abort** button to abort the ongoing active operation.
- **Resume All (Alt + R)** – Click the **Resume All** button to resume all the suspended operations from the point they were suspended. This button is disabled if no operations are suspended.
- **Suspend All (Alt +S)** – Click the **Suspend All** button to suspend all the active operations. The **Suspend All** button is enabled only if one or more operations are in active state.

- **Abort All (Alt + A)** – Click the **Abort All** button to abort all the active operations.
- **Close (Alt + C)** – Click the **Close** button to close the dialog.

Note: The Suspend, Resume, Suspend All, and Resume All will be applicable only for background initialization and rebuild patrol read operations.

6.5.11 Enclosure Properties


To view the enclosure properties, in the Physical View click the **Enclosure**  node.

The Enclosure Properties are displayed, as shown in the following figure.

Figure 6.30 Enclosure Properties

Vendor ID	DELL	FRU Number	42R3133
Enclosure ID	5	Part Number	CP-111-006-020
Enclosure Type	SES	Component Properties	
Enclosure Model	MD1000	Number of Temperature Sensors	4
Enclosure Location	External	Number of Fans	4
Firmware Version	A.04	Number of Power Supplies	2
Serial Number	0802V16VTE	Number of Voltage Sensors	0
Connector	Port A		
Number of Slots	15		

6.5.12 Monitoring Battery Backup Units

When the MegaRAID Storage Manager software is running, you can monitor the status of all of the BBUs connected to controllers in the server. If a BBU is operating normally, the icon looks like this . If it fails, a red dot appears next to the icon.

To show the properties for a BBU, perform the following steps.

1. On the main menu screen, click the **Physical** tab to open the physical view.
2. Select the BBU icon in the left panel.

The BBU properties appear in the right panel. The BBU properties include the following:

- The number of times the BBU has been recharged (cycle count).
- The full capacity of the BBU, plus the percentage of its current state of charge, and the estimated time until it will be depleted.
- The current BBU temperature, voltage, current, and remaining capacity.
- If the battery is charging, the estimated time until it is fully charged.
- The battery state, which says if it is in operational state.
- If battery replacement is required.
- The BBU retention time, which gives the total number of hours the battery can support the current capacity reserve.

The BBU Properties display, as shown in [Figure 6.31](#) and [Figure 6.32](#).

Figure 6.31 Battery Backup Unit Properties for iBBU Batteries






















Properties			
BBU Battery Type	iBBU	Cycle Count	32
Battery State	Operational	Automatic Learn Cycle	Enabled
Battery Replacement	Required / Not required	Auto Learn Period	30 days
Temperature	29.0 C (84.2 F) - Normal	Next Learn Cycle	Aug 10 2010 20:52:13
Voltage	4055 mV	Relative State of Charge	99%
Current	0 mA	Absolute State of Charge	35%
Full Capacity	<value> mAh	Run Time to Empty	Battery is not being discharged
Remaining Capacity	<value> mAh	Average Time to Empty	Battery is not being discharged
BBU Retention Time	 48+ Hours	Average Time to Full	Battery is not being discharged
Estimated Time to Recharge	<value> Mins	Maximum Error Margin	25%
FRU	None		

Figure 6.32 Battery Backup Unit Properties for TMM-C Batteries

Properties			
BBU Battery Type	TMM-C (Not activated) 2	Estimated Time to Recharge	<value> Mins
Battery State	Operational	FRU	None
Battery Replacement	Required / Not required	Memory Module FRU	<value>
Temperature	29.0 C (84.2 F) - Normal	Automatic Learn Cycle	Enabled
Voltage	4055 mV	Auto Learn Period	30 days
Current	0 mA	Next Learn Cycle	Aug 10 2010 20:52:13
Full Capacity	<value> Joules		
Remaining Capacity	<value> Joules		
BBU Retention Time	 48+ Hours		


The icons in the left panel of the following table represent the controllers, drives, and other devices:


Table 6.1 Device Icons

Icon	Definition
	Status
	System
	Controller
	Backplane
	Enclosure
	Port
	Drive group
	Virtual drive
	Online drive
	Power save mode
	Dedicated hotspare
	Global hotspare
	Battery backup unit (BBU)
	Tape drive
	CD-ROM
	Foreign drive
	Unconfigured drive
	Locked SED
	Unlocked SED

Note: The MegaRAID Storage Manager software shows the icons for tape drive devices; however, no tape-related operations

are supported by the utility. If these operations are required, use a separate backup application.

A red circle to the right of an icon indicates that the device has failed. For example, this icon indicates that a drive has failed: .

A yellow circle to the right of an icon indicates that a device is running in a partially degraded state. For example, this icon indicates that a virtual drive is running in a degraded state because a controller has failed: .

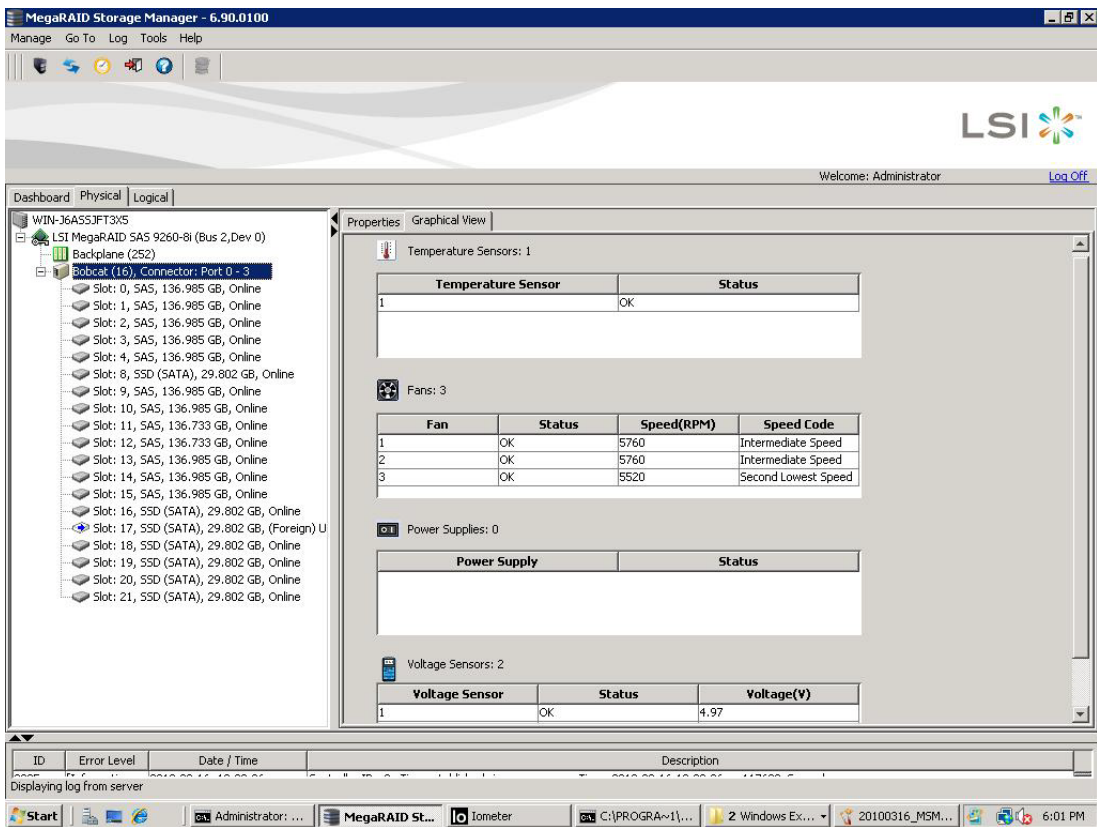
An orange circle to the right of an icon indicates that a device is running in a degraded state.

6.5.13 Properties/Graphical View Tabs

The right panel of the MegaRAID Storage Manager window has one tab or two tabs, depending on which kind of device you select in the left panel. The following figure shows the MSM main menu screen.

- The *Properties* tab displays information about the selected device. For example, if you select a controller icon in the left panel, the Properties tab lists information about the controller, such as the controller name, NVRAM size, and device port count.
- The *Graphical View* tab displays information about the temperature, fans, power supplies, and voltage sensors. To display a graphical view of a drive, click an enclosure icon in the left panel of the MegaRAID Storage Manager main menu screen, and click the Graphical View tab.

Figure 6.33 Properties Tab and Graphical View Tab



6.5.14 Event Log Panel

The lower part of the MegaRAID Storage Manager window displays the system event log entries. New event log entries appear during the session. Each entry has an ID, an error level indicating the severity of the event, the timestamp and date, and a brief description of the event.

6.5.14.1 Menu Bar

The following are brief descriptions of the main selections on the MegaRAID Storage Manager menu bar.

6.5.14.2 Manage Menu

The Manage menu has a Refresh option for updating the display in the MegaRAID Storage Manager window (refresh is seldom required; the display normally updates automatically) and an Exit option to end your session on MegaRAID Storage Manager. The Server menu item shows all the servers that were discovered by a scan. In addition, you can perform a check consistency, initialize multiple virtual groups, and show the progress of group operations on virtual drives.

6.5.14.3 Go To Menu

The Go To menu is available when you select a controller, drive group, physical drive, virtual drive, or battery backup unit in the main menu screen. The menu options vary depending on the type of device selected in the left panel of the MegaRAID Storage Manager main menu. The options also vary depending on the current state of the selected device. For example, if you select an offline drive, the Make Drive Online option appears in the Physical Drive menu.

Configuration options are also available. This is where you access the Configuration Wizard that you use to perform configuration drive groups and virtual drives. To access the Wizard, select the controller in the left panel, and then select **Go To >> Controller >> Create Virtual Drive**.

6.5.14.4 Log Menu

The Log menu includes options for saving and clearing the message log. For more information about the Log menu, see Appendix A, [Events and Messages](#).

6.5.14.5 Tools Menu

On the Tools menu you can select **Tools >> Configure Alerts** to access the Configure Alerts screen, which you can use to set the alert delivery rules, event severity levels, exceptions, and email settings.

6.5.14.6 Help Menu

On the Help menu you can select **Help >> Contents** to view the MegaRAID Storage Manager online help file. You can select

Help >> About MegaRAID Storage Manager to view version information for the MegaRAID Storage Manager software.

Note: When you use the MegaRAID Storage Manager online help, you might see a warning message that Internet Explorer has restricted the file from showing active content. If this warning appears, click on the active content warning bar and enable the active content.

Note: If you are using the Linux operating system, you must install the Firefox® or Mozilla® browser for the MegaRAID Storage Manager online help to display.

Chapter 7

Configuration

You can use the MegaRAID Storage Manager (MSM) software to create and modify storage configurations on ServeRAID-M controllers. These controllers support RAID 0, RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, and RAID 60 storage configurations. To learn more about RAID and RAID levels, see [Chapter 2, “Introduction to RAID.”](#)

The Modify Drive Group Wizard allows you to easily change RAID levels or to expand or reduce the capacity of existing virtual drives.

Note: You cannot create or modify a storage configuration unless you are logged on to a server with administrator privileges.

This chapter explains how to use MegaRAID Storage Manager software to perform the following configuration tasks:

- [Section 7.1, “Creating a New Storage Configuration”](#)
- [Section 7.2, “Selecting Self-Encrypting Disk Security Options”](#)
- [Section 7.3, “Adding Hot Spare Drives”](#)
- [Section 7.4, “Changing Adjustable Task Rates”](#)
- [Section 7.5, “Recovering and Clearing Punctured Block Entries”](#)
- [Section 7.6, “Changing Virtual Drive Properties”](#)
- [Section 7.7, “Changing a Virtual Drive Configuration”](#)
- [Section 7.8, “Deleting a Virtual Drive”](#)

7.1 Creating a New Storage Configuration

You can use the MegaRAID Storage Manager to create new storage configurations on systems with IBM ServeRAID SAS controllers. You can create the following types of configurations:

- **Simple configuration** specifies a limited number of settings and has the system select drives for you. This option is the easiest way to create a virtual drive.
- **Advanced configuration** lets you choose additional settings and customize virtual drive creation. This option provides greater flexibility when creating virtual drives for your specific requirements.

The following subsections describe the virtual drive parameters and explain how to create simple and advanced storage configurations:

- [Section 7.1.1, “Selecting Virtual Drive Settings”](#)
- [Section 7.1.2, “Creating a Virtual Drive Using Simple Configuration”](#)
- [Section 7.1.3, “Creating a Virtual Drive Using Advanced Configuration”](#)

7.1.1 Selecting Virtual Drive Settings

This section describes the virtual drive settings that you can select when you use the advanced configuration procedure to create virtual drives. You should change these parameters only if you have a specific reason for doing so. It is usually best to leave them at their default settings.

- **Initialization state:** Initialization prepares the storage medium for use. Specify the initialization status:
 - ◇ *No Initialization:* (the default) The new configuration is not initialized and the existing data on the drives is not overwritten.
 - ◇ *Fast Initialization:* The firmware quickly writes zeroes to the first and last 8-Mbyte regions of the new virtual drive and then completes the initialization in the background. This allows you to start writing data to the virtual drive immediately.

- ◇ *Full Initialization*: A complete initialization is done on the new configuration. You cannot write data to the new virtual drive until the initialization is complete. This can take a long time if the drives are large.
- **Stripe size**: Stripe sizes of 8, 16, 32, 64, 128, 256, 512, and 1024 Kbytes are supported. The default is 64 Kbytes. For more information, see the *striping* Glossary entry.

Note: The ServeRAID M1015 and M1115 SAS/SATA controllers do not support the cache policies. These policies include access policy, read policy, write policy, IO policy, and drive cache.

- **Read policy**: Specify the read policy for this virtual drive:
 - ◇ *Always read ahead*: Read ahead capability allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data, but there is little improvement when accessing random data.
 - ◇ *No read ahead*: (the default) Disables the read ahead capability.
- **Write policy**: Specify the write policy for this virtual drive:
 - ◇ *Write through*: (the default) In this mode the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction. This option eliminates the risk of losing cached data in case of power failure.
 - ◇ *Always Write Back*: In this mode, the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction.
 - ◇ *Write Back with BBU*: (the default) In this mode, the controller enables Write Back caching when the battery backup unit (BBU) is installed and charged. It provides a good balance between data protection and performance.

Note: The Write Policy depends on the status of the battery backup unit (BBU). If the BBU is not present, is low, is

failed, or is being charged, the default Write Policy will be Write through. This provides better data protection.

- **I/O policy:** The IO policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.

- ◊ *Cached IO:* In this mode, all reads are buffered in cache memory.
- ◊ *Direct IO:* (the default) In this mode, reads are not buffered in cache memory. Data is transferred to the cache and the host concurrently. If the same data block is read again, it comes from cache memory.

Cached IO provides faster processing than Direct IO. Direct IO ensures that the cache and the host contain the same data.

- **Access policy:** Select the type of data access that is allowed for this virtual drive.

- ◊ *Read/Write:* (the default) Allow read/write access. This is the default.
- ◊ *Read Only:* Allow read-only access.
- ◊ *Blocked:* Do not allow access.

- **Disk cache policy:** Select a cache setting for this drive:

- ◊ *Enable:* Enable the disk cache.
- ◊ *Disable:* Disable the disk cache.
- ◊ *Unchanged:* (the default) Leave the current disk cache policy unchanged.

7.1.2 Creating a Virtual Drive Using Simple Configuration

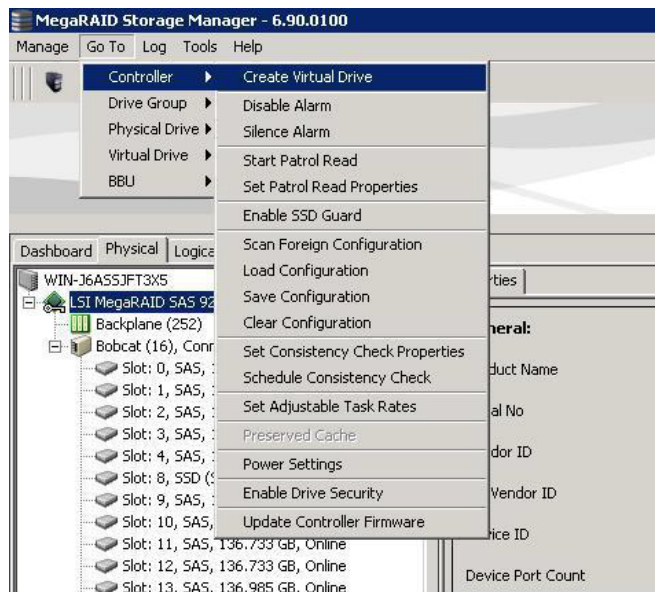
Simple configuration is the quickest and easiest way to create a new storage configuration. When you select simple configuration mode, the system creates the best configuration possible using the available drives.

Note: You cannot create spanned drives using the simple configuration procedure. To create spanned drives, use the advanced configuration procedure described in [Section 7.1.3, “Creating a Virtual Drive Using Advanced Configuration”](#).

Follow these steps to create a new storage configuration in simple configuration mode.

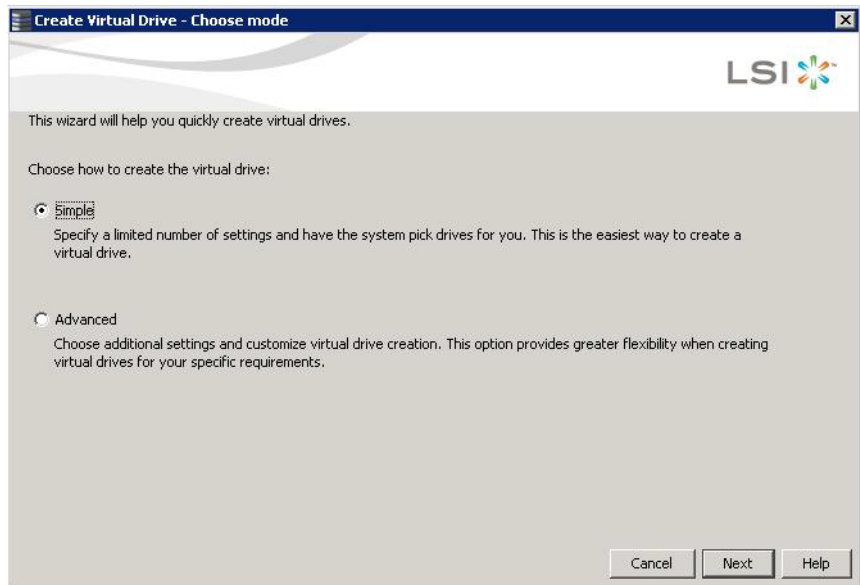
1. Perform either of the following steps:
 - Right-click the controller node in the device tree in the left frame of the MegaRAID Storage Manager window and select **Create Virtual Drive**
 - Select the controller node and select **Go To >> Controller >> Create Virtual Drive** in the menu bar, as shown in the following figure.

Figure 7.1 Virtual Drive Creation Menu



The dialog box for the configuration mode (simple or advanced) appears, as shown in the following figure.

Figure 7.2 Virtual Drive Creation Mode

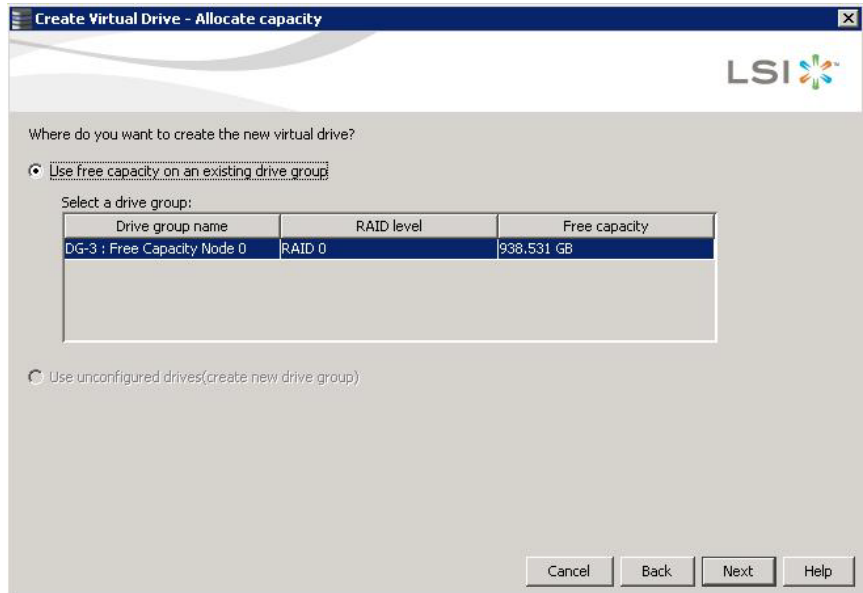


2. Select **Simple** and click **Next**.

If there are no unconfigured drives available, you have the option to use free capacity of an existing drive group, as shown in the following figure. The Create Virtual drive-Summary screen appears as shown in [Figure 7.5](#).

If unconfigured drives are available, [Figure 7.4](#) appears, and you can go to [step 4](#).

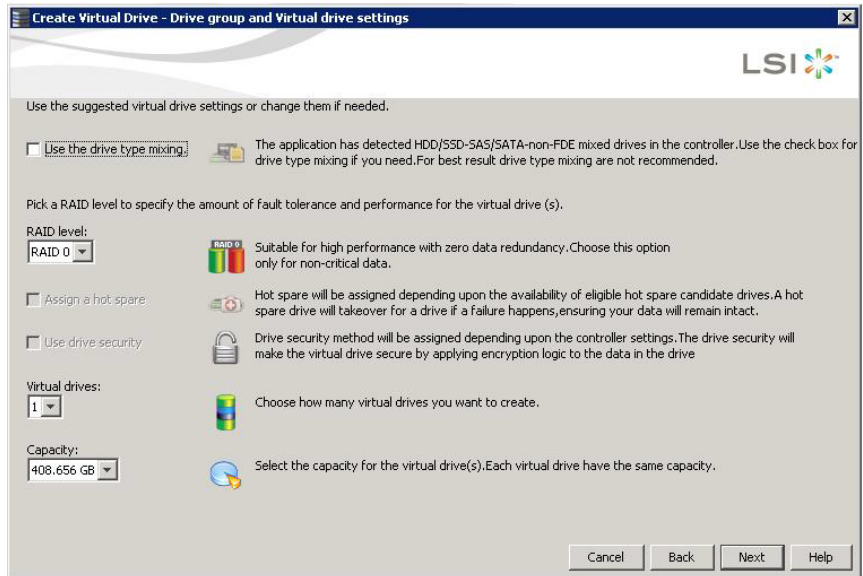
Figure 7.3 Using the Free Capacity of an Existing Drive Group



3. Check the option **Use Unconfigured drives (create new drive group)** and click **Next**.

The Create Virtual Drive screen appears, as shown in the following figure. If there are different types of drives attached to the controller, such as HDD, SSD, SAS, and SATA, there is an option to allow drive type mixing.

Figure 7.4 Create Virtual Drive Screen



4. If you want to allow different types of drives in a configuration, click the checkbox **Use the drive type mixing**.

Note: For best results, do not use drive type mixing.

5. Select a RAID level for the virtual drive.

When you use simple configuration, the RAID controller supports RAID levels 1, 5, and 6 virtual drives. In addition, it supports independent drives (configured as RAID 0). The screen text gives a brief description of the RAID level that you select. The RAID levels that you can choose depend on the number of drives available. To learn more about RAID levels, see [Chapter 2, “Introduction to RAID.”](#)

6. Click the **Assign a hot spare** check box if you want to assign a dedicated hot spare to the new virtual drive.

If an unconfigured good drive is available, that drive is assigned as a hot spare. Hot spares are drives that are available to replace failed drives automatically in a redundant virtual drive (RAID 1, RAID 5, or RAID 6).

Note: In the simple configuration procedure, you can assign dedicated hot spares to a maximum of 16 drive groups at one time. In ServeRAID, dedicated hot spares can support

only up to 16 drive groups. If you try to create more than 16 drive groups at one time, dedicated hot spares are not assigned to drive groups beyond the first 16.

To create more than 16 drive groups with hot spares, you need at least 35 drives of the same capacity.

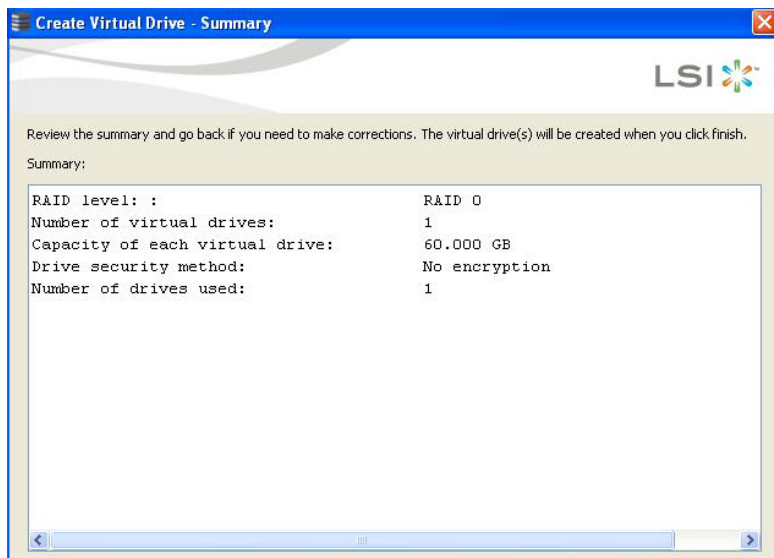
7. Click the box next to the text **Use drive security** if you want to set a drive security method.

The Self-Encrypting Disk (SED) feature provides the ability to encrypt data and use disk-based key management for your data security solution. This solution provides protection to the data in the event of theft or loss of drives.

8. Use the drop-down menu in the **Virtual drives** field to choose how many virtual drives you want to create.
9. Select the capacity of the virtual drive(s).
Each virtual drive has the same capacity.
10. Click **Next**.

The **Create Virtual Drive - Summary** window appears, as shown in the following figure. This window shows the selections you made for simple configuration.

Figure 7.5 Create Virtual Drive - Summary Window



11. Click **Back** to return to the previous screen to change any selections or click **Finish** to accept and complete the configuration.

The new virtual drive is created after you click **Finish**. After the configuration is completed, a dialog box notifies you that the virtual drives were created successfully.

Note: If you create a large configuration using drives that are in powersave mode, it could take several minutes to spin up the drives. A progress bar appears as the drives spin up. If any of the selected unconfigured drives fail to spin up, a box appears that identifies the drive or drives.

After the configuration is completed, a dialog box notifies you that the virtual drives were created successfully.

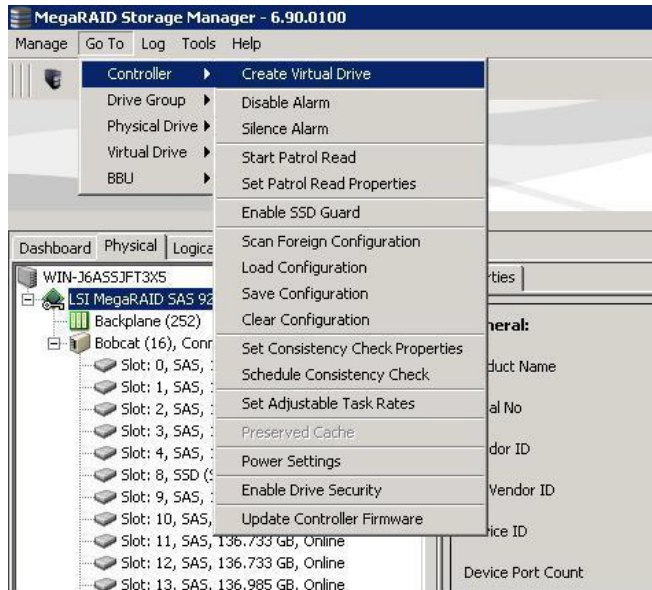
7.1.3 Creating a Virtual Drive Using Advanced Configuration

The advanced configuration procedure provides an easy way to create a new storage configuration. Advanced configuration gives you greater flexibility than simple configuration because you can select the drives and the virtual drive parameters when you create a virtual drive. In addition, you can use the advanced configuration procedure to create spanned drive groups.

Follow these steps to create a new storage configuration in the advanced configuration mode. This example shows the configuration of a spanned drive group.

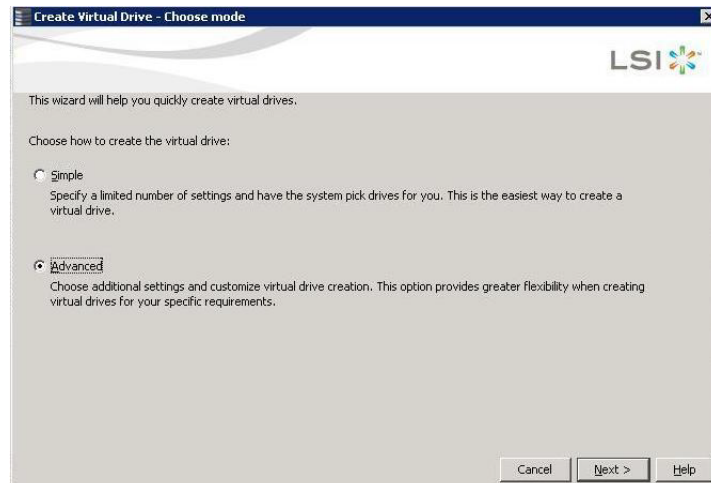
1. Perform either of the following steps:
 - Right click on the controller node in the device tree in the left frame of the MegaRAID Storage Manager window and select **Create Virtual Drive**
 - Select the controller node and select **GoTo >> Controller >> Create Virtual Drive** in the menu bar, as shown in the following figure

Figure 7.6 Virtual Drive Creation Menu



The dialog box shown in the following figure appears.

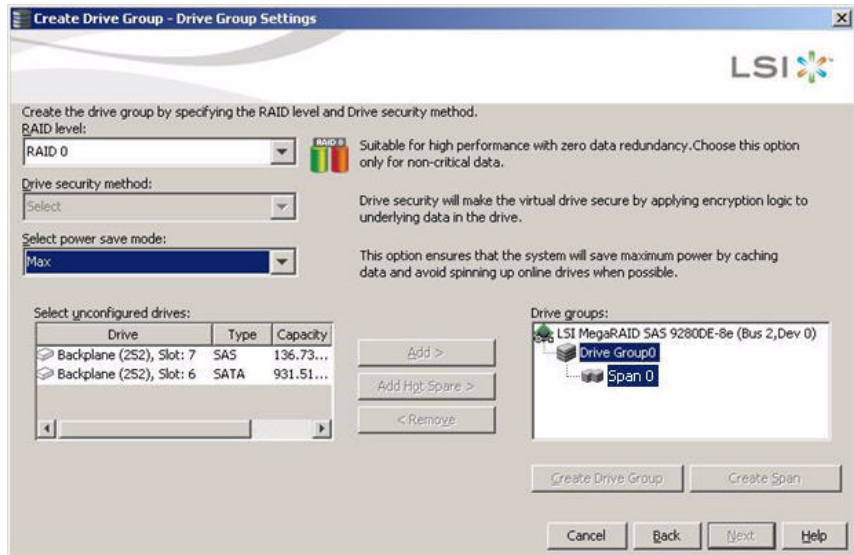
Figure 7.7 Virtual Drive Advanced Configuration Mode



2. Click **Advanced** and then click **Next**.

The Create Drive Group Settings screen appears, as shown in the following figure.

Figure 7.8 Create Drive Group Settings Screen



3. Select the following items on the Create Drive Group Settings screen:
 - a. Select the RAID level desired for the drive group from the drop-down menu. To make a spanned drive, select **RAID 10**, **RAID 50**, or **RAID 60** in the **RAID level** field.

Drive Group 0 and **Span 0** appear in the **Drive groups** field when you select RAID 10, 50, or 60.

The RAID controller supports RAID levels 1, 5, 6, 10, 50, and 60. In addition, it supports independent drives (configured as RAID 0). The screen text gives a brief description of the RAID level you select. RAID levels you can choose depend on the number of drives available. To learn more about RAID levels, see [Chapter 2, "Introduction to RAID."](#)

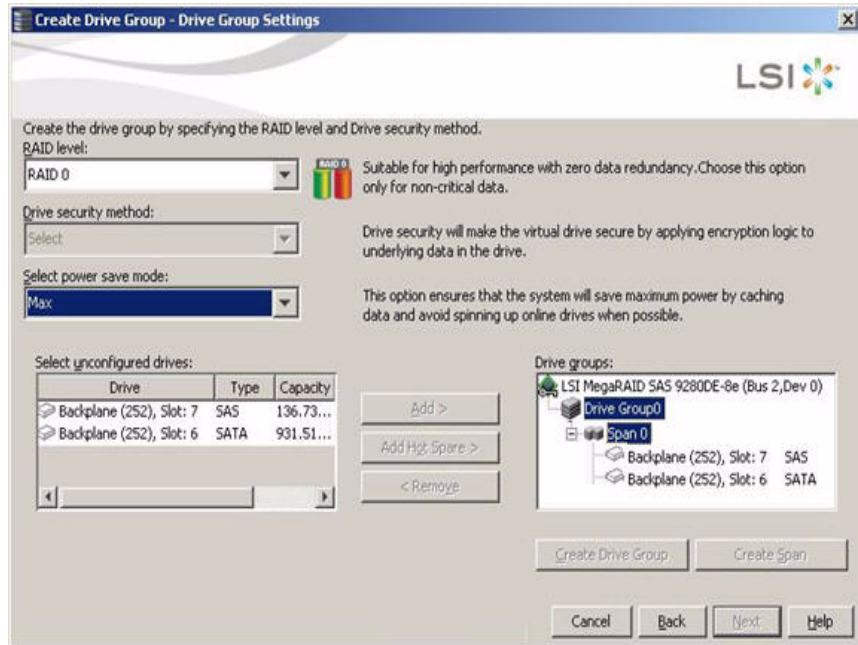
- b. Scroll down the menu for the **Drive security method** field if you want to set a drive security method.

The SED feature provides the ability to encrypt data and use disk-based key management for your data security solution. This solution provides protection to the data in the event of theft or loss of drives.

- c. Select *unconfigured* drives from the list of drives and click **Add >** to add them to the drive group.

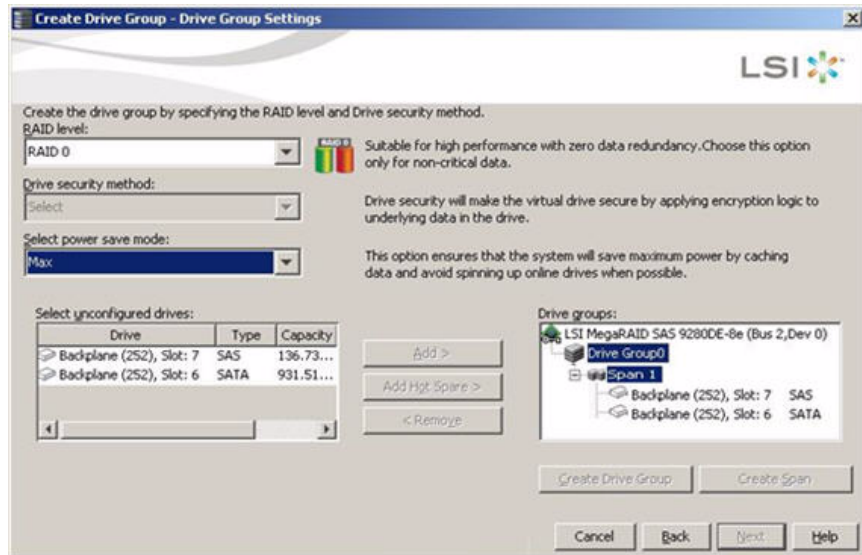
The selected drives appear under **Span 0** below **Drive Group 0**, as shown in the following figure.

Figure 7.9 Span 0 of Drive Group 0



- d. Click **Create Span** to create a second span in the drive group.
- e. Select *unconfigured* drives from the list of drives and click **Add >** to add them to the drive group.
- f. The selected drives appear under **Span 1** below **Drive Group 0**, as shown in the following figure.

Figure 7.10 Span 0 and Span 1 of Drive Group 0




- g. Click **Create Drive Group** to make a drive group with the spans.
- h. Click **Next** to complete this step.

The Virtual drive settings window appears, as shown in the following figure. The drive group and the default virtual drive settings appear. The options to update the virtual drive or remove the virtual drive are grayed out until you create the virtual drive.

Note: The parameters in the Virtual drive settings window display in disabled mode (grayed out) for SAS-Integrated RAID (IR) controllers because these parameters do not apply to SAS-IR controllers.

Figure 7.11 Virtual Drive Settings Window



Note: If you select Write Back with BBU as the Write policy, and there is no battery, or the battery is low or failed, or the battery is running through a re-learn cycle, the Write policy switches to Write Through. This eliminates the risk of data loss in case of power failure. A message screen notifies you of this change.

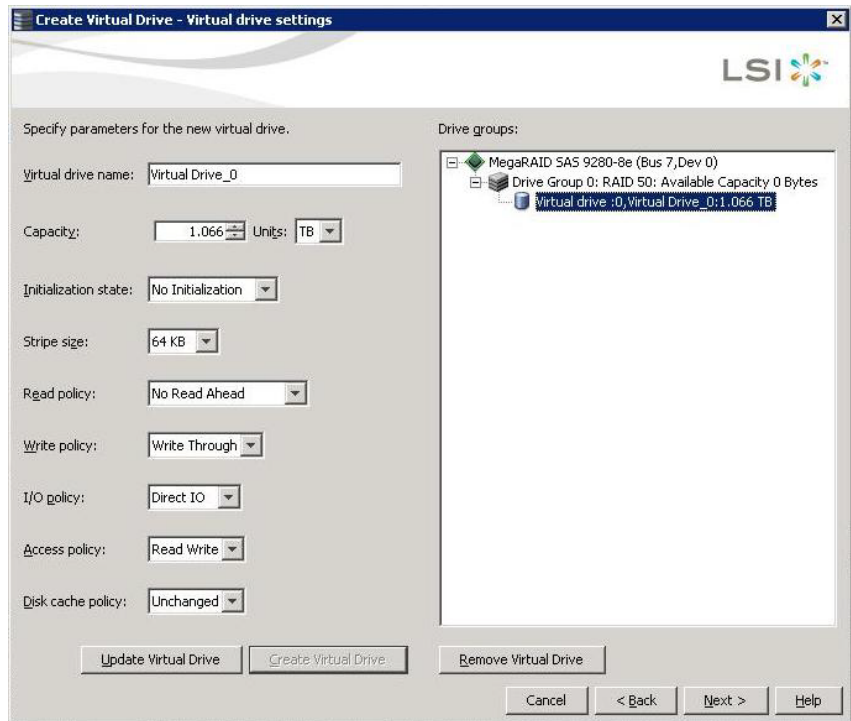
4. Change any virtual drive settings, if desired.

See [Section 7.1.1, “Selecting Virtual Drive Settings”](#) for more information about the virtual drive settings.

5. Click **Create Virtual Drive**.

The new virtual drive appears under the drive group, as shown in the following figure. The options **Update Virtual Drive** and **Remove Virtual Drive** are now available. **Update Virtual Drive** allows you to change the virtual drive settings and **Remove Virtual Drive** allows you to delete the virtual drive.

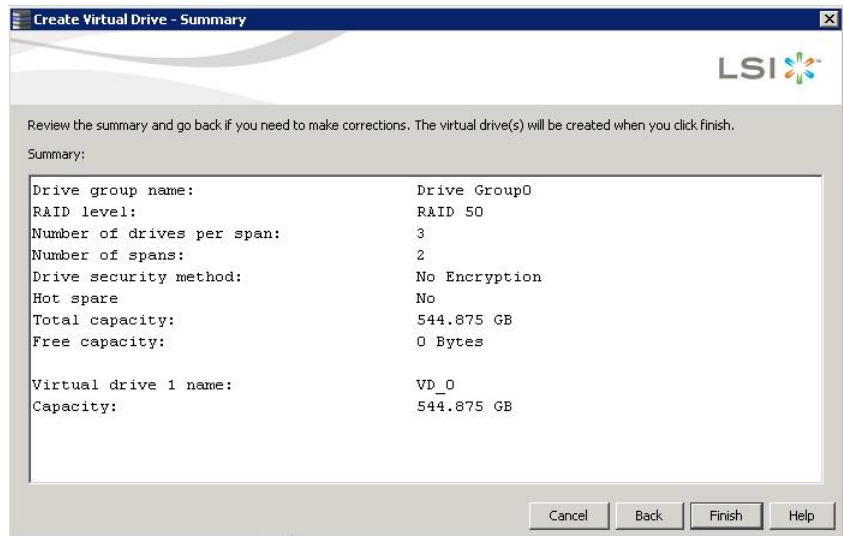
Figure 7.12 New Virtual Drive 0



6. Click **Next**.

The **Create Virtual Drive - Summary** window appears, as shown in the following figure. This window shows the selections you made for advanced configuration.

Figure 7.13 Create Virtual Drive Summary Window



7. Click **Back** to return to the previous screen to change any selections or click **Finish** to accept and complete the configuration.

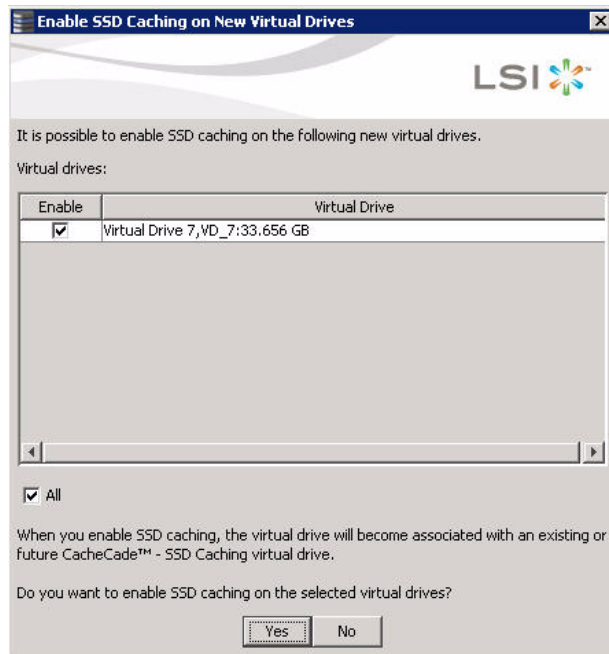
After you click **Finish**, the new storage configuration is created and initialized.

Note: If you create a large configuration using drives that are in powersave mode, it could take several minutes to spin up the drives. A progress bar appears as the drives spin up. If any of the selected unconfigured drives fail to spin up, a box appears to identify the drive or drives.

After the configuration is completed, a dialog box notifies you that the virtual drives were created successfully.

8. Click **OK**. The Enable SSD Caching on New Virtual Drives dialog appears, as shown in the following figure.

Figure 7.14 Enable SSD Caching on New Virtual Drives



The newly created virtual drive is enabled for SSD caching by default.

9. Click **OK** to confirm SSD caching on the virtual drive. Click **No** if you want to disable SSD caching on the virtual drive.

The **All** check box is selected by default. To disable SSD caching on the virtual drives, deselect the **All** check box.

If more drive capacity exists, the dialog asks whether you want to create more virtual drives. If no more drive capacity exists, you are prompted to close the configuration session.

10. Select either **Yes** or **No** to indicate whether you want to create additional virtual drives.

If you select **Yes**, the system takes you to the Create Virtual Drive window, as shown in [Figure 7.4](#). If you select **No**, the utility asks whether you want to close the wizard.

11. If you selected **No** in the previous step, select either **Yes** or **No** to indicate whether you want to close the wizard.

If you select **Yes**, the configuration wizard closes. If you select **No**, the dialog closes, and you remain on the same page.

7.2 Selecting Self-Encrypting Disk Security Options

The Self-Encrypting Disk feature provides the ability to encrypt data and use disk-based key management for your data security solution. This solution provides protection to the data in the event of theft or loss of physical drives. This section describes how to enable, change, or disable drive security, and how to import a foreign configuration.

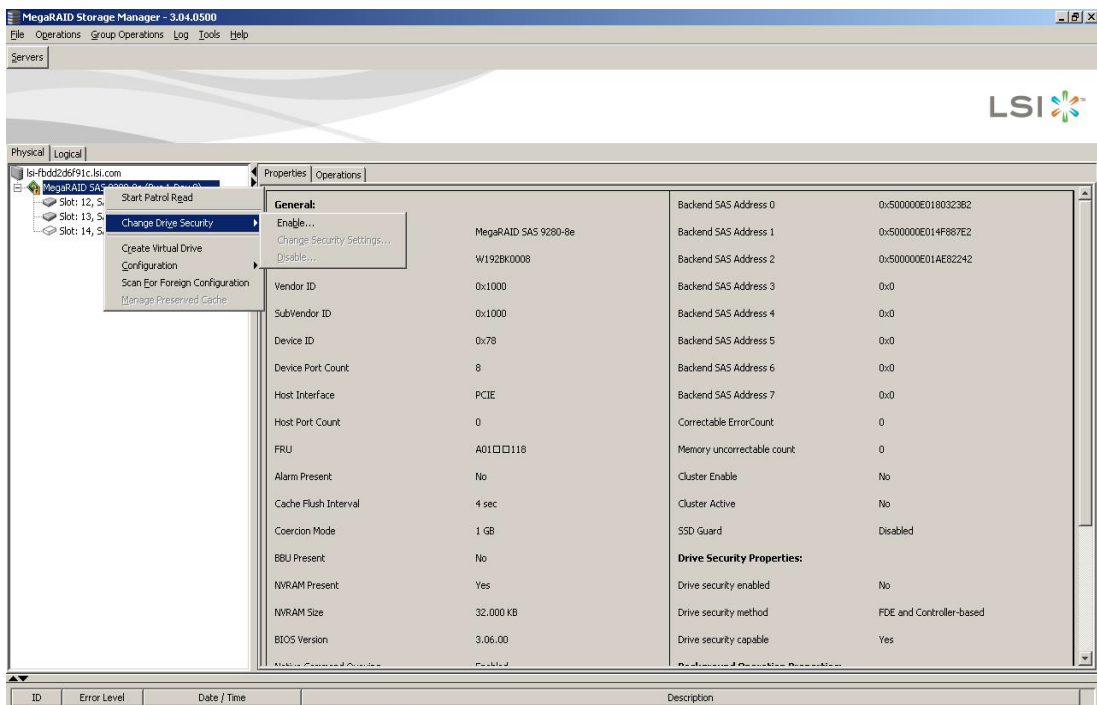
7.2.1 Enabling Drive Security

Perform the following steps to enable drive security. To do this, you create a security key identifier, security key, and (optional) passphrase.

1. Select the **Physical View** tab in the left panel of the MegaRAID Storage Manager window, and click a controller icon.
2. Right-click on the controller icon to display the menu of operations available.
3. Select **Change Drive Security >> Enable**, as shown in the following figure.

Note: You can access the drive security settings menu by clicking the Operations menu on the menu bar and selecting **Change Drive Security >> Enable**, also.

Figure 7.15 Drive Security Settings Menu



The Enable Drive Security – Introduction screen appears as shown in the following figure. This screen describes how the wizard will help you create a security key on the controller. After you create a security key, you have the option to create secure virtual drives using the security key.

First, create the security key identifier. The identifier appears whenever you have to enter the security key. If you have more than one security key, the identifier helps you determine which security key to enter.

Next, create a security key. You need the security key to perform certain operations. Finally, you have the option to create a passphrase for additional security. If you create a passphrase, you must enter it whenever you boot your server.

Figure 7.16 Enable Drive Security - Introduction Screen



4. On the introduction screen, click **Next**.

The Enter Security Key ID screen appears, as shown in the following figure.

Figure 7.17 Enter Security Key ID Screen



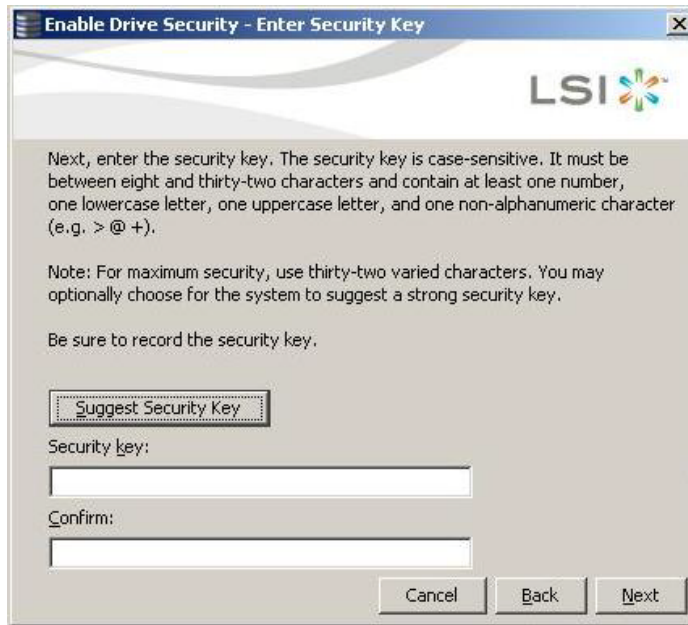
5. Use the default security key identifier or enter a new security key identifier.

Note: If you create more than one security key, it is highly recommended that you change the security key identifier. Otherwise, you cannot differentiate between the security keys.

6. Click **Next**.

The Enable Security Key ID screen appears as shown in the following figure.

Figure 7.18 Enter Security Key Screen



Enable Drive Security - Enter Security Key

LSI

Next, enter the security key. The security key is case-sensitive. It must be between eight and thirty-two characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (e.g., > @ +).

Note: For maximum security, use thirty-two varied characters. You may optionally choose for the system to suggest a strong security key.

Be sure to record the security key.

Suggest Security Key

Security key:

Confirm:

Cancel Back Next

7. Click **Suggest Security Key** to have the systems create the security key or enter a new security key. Enter the new security key again to confirm.

Attention: **If you forget the security key, you will lose access to your data.** Be sure to record your security key information. You might need to enter the security key to perform certain operations.

The security key is case-sensitive. It must be between eight and thirty-two characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (for example, < > @ +). The space character is not permitted.

Note: Non-US keyboard users must be careful not to enter DBCS characters in the security key field. Firmware works with the ASCII character set only.

8. Click **Next**.

The Enter Pass Phrase screen appears, as shown in the following figure.

Figure 7.19 Enable Drive Security - Enter Pass Phrase Screen



Enable Drive Security - Enter Pass Phrase

LSI

Optionally, you may enter a pass phrase to provide additional security. If you choose to require a pass phrase, you will need to enter it every time you reboot the server, and whenever you provide the security key. The pass phrase should be different from the security key.

The pass phrase is case-sensitive. It must be between eight and thirty-two characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (e.g. ! * @ + #).

☒ Use a pass phrase in addition to the security key

Pass phrase:

Confirm:

Cancel Back Finish

9. Click **Use a pass phrase in addition to the security key** if you want to use the pass phrase for additional security.
10. Enter a passphrase in the **Pass phrase** field and then enter the passphrase in the **Confirm** field.

The passphrase is case-sensitive. It must be between eight and thirty-two characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (for example, < > @ +). The space character is not permitted.

Warning messages appear if there is a mismatch between the characters entered in the Passphrase field and the Confirm field, or if there is an invalid character entered.

Attention: Be sure to record the passphrase. If you lose the passphrase, you could lose access to your data.

11. Click **Next**.

The Confirm Enable Drive Security screen appears, as shown in the following figure, to show the changes requested to the drive security settings.

Attention: **If you forget the security key, you will lose access to your data.** Be sure to record your security key. You might need to enter the security key to perform certain operations.

Figure 7.20 Confirm Create Security Key Screen



12. Confirm that you want to enable drive security on this controller and have recorded the security settings for future reference.

MSM enables drive security and returns you to the main menu.

7.2.2 Changing the Security Key Identifier, Security Key, and Pass Phrase

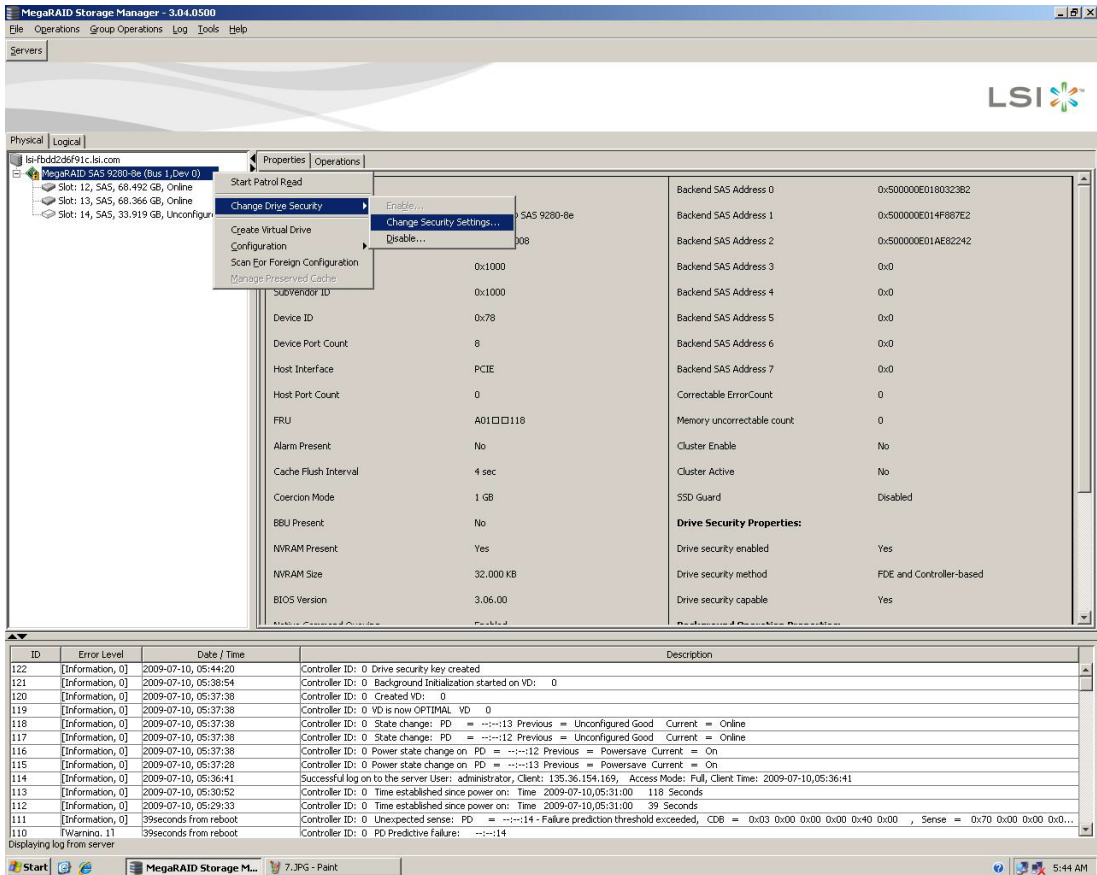
Perform the following steps to change the encryption settings for the security key identifier, security key, and pass phrase.

1. Select the **Physical View** tab in the left panel of the MegaRAID Storage Manager window, and click a controller icon.
2. Right-click on the controller icon to display the menu of operations available.
3. Select **Change Drive Security >> Change Security Settings**, as shown in the following figure.

Note: You can also access the drive security settings menu by clicking the Operations menu on the menu bar and


selecting **Change Drive Security >> Change Security Settings**.

Figure 7.21 Change Drive Security Menu



The Change Security Settings – Introduction screen appears as shown in the following figure. This screen lists the actions you can perform, which include editing the security key identifier, security key, and the passphrase.

Figure 7.22 Change Security Settings - Introduction Screen



4. On the introduction screen, click **Next**.

The Change Security Settings - Security Key ID screen appears, as shown in the following figure.

Figure 7.23 Change Security Settings - Security Key ID Screen

Change Security Settings - Security Key ID

LSI

Select whether you want to keep the existing drive security key identifier or enter a new one.

Note: If you plan to change the security key, it is highly recommended that you change the security key identifier. Otherwise, you will not be able to differentiate between the security keys.

☒ Use the existing security key identifier.

Current security key identifier:
MR9240-8i_L000651809_11e91b1d

☐ Enter a new security key identifier

New security key identifier:

Cancel Back Next

5. Keep the existing security key identifier or enter a new security key identifier.

Note: If you change the security key, it is highly recommended that you change the security key identifier. Otherwise, you cannot distinguish between the security keys.

6. Click **Next**.

The Change Security Settings - Security Key screen appears as shown in the following figure.

Figure 7.24 Change Security Settings - Security Key Screen



7. Click **Use the existing drive security key** to use the existing drive security key or enter a new security key and then enter the new security key again to confirm.

Attention: **If you forget the security key, you will lose access to your data.** Be sure to record your security key information. You might need to enter the security key to perform certain operations.

The security key is case-sensitive. It must be between eight and thirty-two characters and contain at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (for example, < > @ +). The space character is not permitted.

Note: Non-US keyboard users must be careful not to enter DBCS characters in the security key field. Firmware works with the ASCII character set only.

8. Click **Next**.

The Authenticate Drive Security Settings Screen appears, as shown in the following figure. Authentication is required for the changes that you requested to the drive security settings.

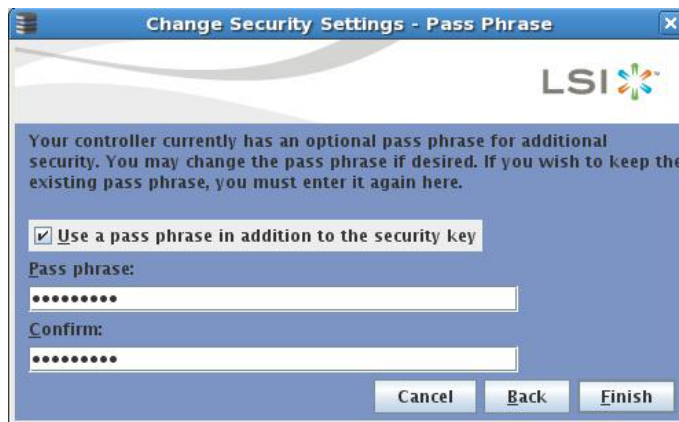
Figure 7.25 Authenticate Drive Security Settings Screen



9. Enter the current security key to authenticate the changes.

The Change Security Settings - Pass Phrase screen appears, as shown in the following figure.

Figure 7.26 Change Security Settings - Pass Phrase Screen



10. If you choose to, click the option to use a passphrase in addition to the security key.
11. If you chose to use a passphrase, either enter the existing passphrase or enter a new passphrase, and enter the passphrase again to confirm. If not, proceed to [step 13](#).

The text box for the passphrase can hold up to 32 characters. The key must be at least eight characters.

12. Click **Finish**.

The next screen that appears describes the changes you made and asks you whether you want to confirm these changes.

13. Click the checkbox to confirm that you have recorded the security settings for future reference.
14. Click **Yes** to confirm that you want to change the drive security settings.

MSM updates the existing configuration on the controller to use the new security settings and returns you to the main menu.

7.2.3 Disabling Drive Security

Note: If you disable drive security, your existing data will not be secure and you cannot create any new secure virtual drives. Disabling drive security does not affect the security of data on foreign drives. If you removed any drives that were previously secured, you still need to enter the passphrase when you import them. Otherwise, you cannot access the data on those drives.

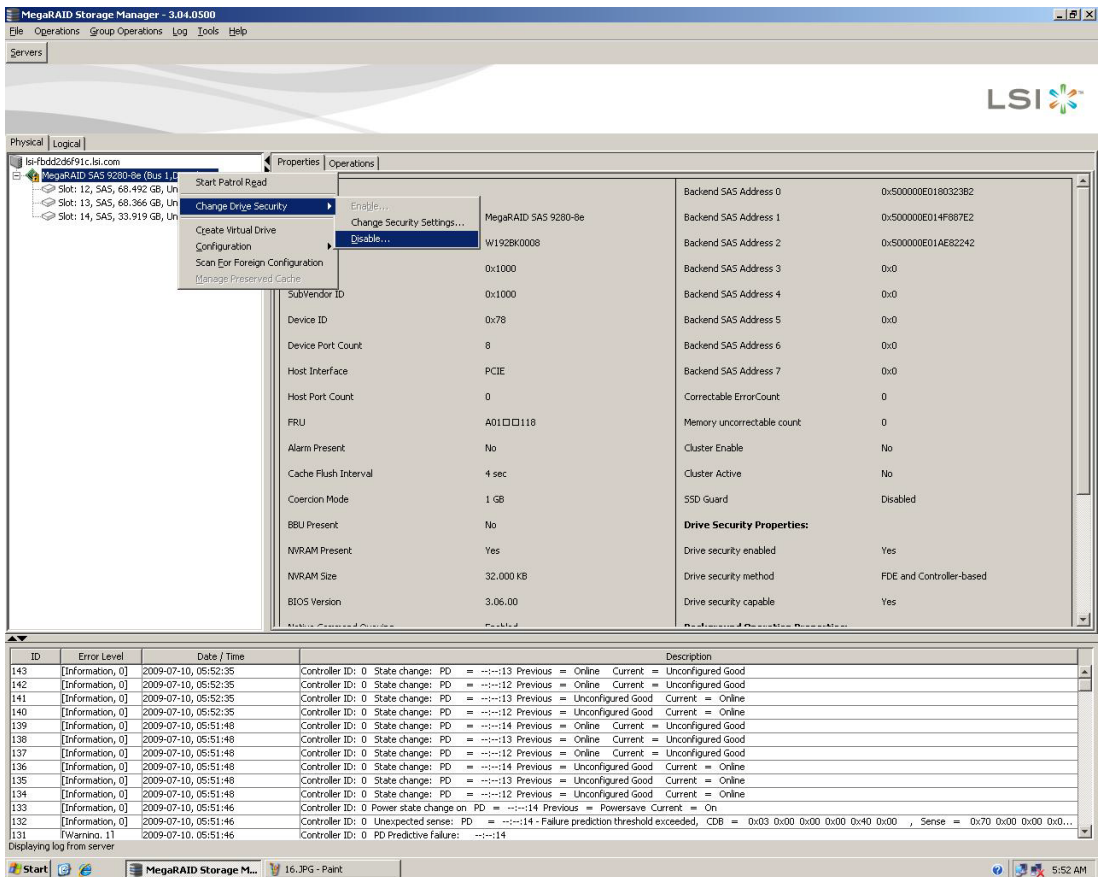
Note: If there are any secure drive groups on the controller, you cannot delete the security key and a warning screen appears. In order to delete the security key, you must first delete the virtual drives on all of the secure drive groups.

Perform the following steps to disable drive security.

1. Select the **Physical View** tab in the left panel of the MegaRAID Storage Manager window, and click a controller icon.
2. Right-click on the controller icon to display the menu of operations available.
3. Select **Change Drive Security >> Disable**, as shown in the following figure.

Note: You can also access the drive security settings menu by clicking the Operations menu on the menu bar and selecting **Change Drive Security >> Disable**.

Figure 7.27 Change Drive Security Menu



The Confirm Disable Drive Security screen appears as shown in the following figure.

Figure 7.28 Confirm Disable Drive Security Screen



4. To disable drive security, click **Yes**.

MSM disables drive security and returns you to the main menu.

7.2.4 Importing or Clearing a Foreign Configuration

A foreign configuration is a RAID configuration that already exists on a replacement set of drives that you install in a computer system. You can use MSM to import the foreign configuration to the RAID controller or to clear the foreign configuration so you can create a new configuration using these drives.

To import a foreign configuration, you must do the following:

- Enable security to allow importation of locked foreign configurations. (You can import unsecured or unlocked configurations when security is disabled.)
- Run a scan for foreign configurations.
- If a locked foreign configuration is present and security is enabled, enter the security key and unlock the configuration.
- Import the foreign configuration.

In addition, if one or more drives are removed from a configuration, by a cable pull or drive removal for example, the configuration on those drives is considered a foreign configuration by the RAID controller.

Verify whether any drives are left to import as the locked drives can use different security keys. If there are any drives left, repeat the import process for the remaining drives. After all the drives are imported, there is no configuration to import.

Note: When you create a new configuration, MSM shows only the unconfigured drives. Drives that have existing configurations, including foreign configurations, do **not** appear. To use drives with existing configurations, you must first clear the configuration on those drives.

Perform the following steps to import or clear a configuration.

1. Enable drive security to allow importation of locked foreign drives.
See [Section 7.2.1, “Enabling Drive Security”](#) for the procedure used to enable drive security.
2. After you create a security key, right-click on the controller and click **Scan for Foreign Configuration**.
If there are locked drives (security is enabled), the Unlock foreign drives dialog box appears.
3. Enter the security key and unlock the configuration.
The The Foreign Configuration Detected screen appears, as shown in the following figure.

Figure 7.29 Foreign Configuration Detected Screen



4. Click one of the following:
 - a. **Import** to import the foreign configurations from all of the foreign drives
 - b. **Clear** to remove the configurations from all of the foreign drives
 - c. **Advanced** to preview and import specific foreign configurations.
5. Click **OK**.

Note: The operation cannot be reversed after it is started. Imported drives display as *Online* in the MegaRAID Storage Manager window.

6. Repeat the import process for any remaining drives.

Because locked drives can use different security keys, you must verify whether there are any remaining drives to imported.

Note: When you create a new configuration, MSM shows only the unconfigured drives. Drives that have existing configurations, including foreign configurations, do not appear. To use drives with existing configurations, you must first clear the configuration on those drives.

7.2.4.1 Foreign Configurations in Cable Pull and Drive Removal Scenarios

If one or more drives are removed from a configuration, by a cable pull or drive removal, for example, the configuration on those drives is considered a foreign configuration by the RAID controller.

The following scenarios can occur with cable pulls or drive removals. Use the **Foreign Configuration Preview** screen to import or clear the foreign configuration in each case.

Note: If you want to import the foreign configuration in any of the following scenarios, you should have all of the drives in the enclosure before you perform the import operation.

1. Scenario #1: If all of the drives in a configuration are removed and re-inserted, the controller considers the drives to have foreign configurations.

Import or clear the foreign configuration. If you select **Import**, automatic rebuilds occur in redundant virtual drives.

Note: Start a consistency check immediately after the rebuild is complete to ensure data integrity for the virtual drives. See [Section 9.3, “Running a Consistency Check,”](#) for more information about checking data consistency.

2. Scenario #2: If some of the drives in a configuration are removed and re-inserted, the controller considers the drives to have foreign configurations.

Import or clear the foreign configuration. If you select **Import**, automatic rebuilds occur in redundant virtual drives.

Note: Start a consistency check immediately after the rebuild is complete to ensure data integrity for the virtual drives. See [Section 9.3, “Running a Consistency Check,”](#) for more information about checking data consistency.

3. Scenario #3: If all of the drives in a virtual drive are removed, but at different times, and re-inserted, the controller considers the drives to have foreign configurations.

Import or clear the foreign configuration. If you select **Import**, all drives that were pulled *before* the virtual drive became offline will be imported and then automatically rebuilt. Automatic rebuilds occur in any redundant virtual drives.

4. Scenario #4: If the drives in a non-redundant virtual drive are removed, the controller considers the drives to have foreign configurations.

Import or clear the foreign configuration. No rebuilds occur after the import operation because there is no redundant data to rebuild the drives with.

7.3 Adding Hot Spare Drives

Hot spares are drives that are available to automatically replace failed drives in a RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, or RAID 60 virtual drive. *Dedicated hot spares* are used to replace failed drives in a selected drive group only. *Global hot spares* are available to any virtual drive on a specific controller.

To add a dedicated hot spare or a global hot spare drive, follow these steps:

1. Select the **Physical** tab in the left panel of the MegaRAID Storage Manager main menu, and click the icon of an unused drive.

For each drive, the screen displays the port number, enclosure number, slot number, drive state, drive capacity, and drive manufacturer.

2. Select **Go To >> Physical Drive >> Assign (G)lobal Hot Spare** or **Go To >> Physical Drive >> Assign (D)edicated Hot Spare**.

3. If you selected **Assign Dedicated Hotspare**, select a drive group from the list that appears. The hot spare is dedicated to the drive group that you select.

If you selected **Assign Global Hotspare**, skip this step and go to the next step. The hot spare is available to any virtual drive on a specific controller.

4. Click **Go** to create the hot spare.

The drive state for the drive changes to dedicated hot spare or global hot spare, depending on your selection.

7.4 Changing Adjustable Task Rates

If you want to change the Rebuild rate and other task rates for a controller, you must first log onto the server in Full Access mode.

Note: Leave the adjustable task rates at their default settings to achieve the best system performance. If you raise the task rates above the defaults, foreground tasks run more slowly and it might seem that the system is not responding. If you lower the task rates below the defaults, rebuilds and other background tasks might run very slowly and might not complete within a reasonable time. If you decide to change the values, record the original default value here so you can restore them later, if necessary:

Rebuild Rate: _____

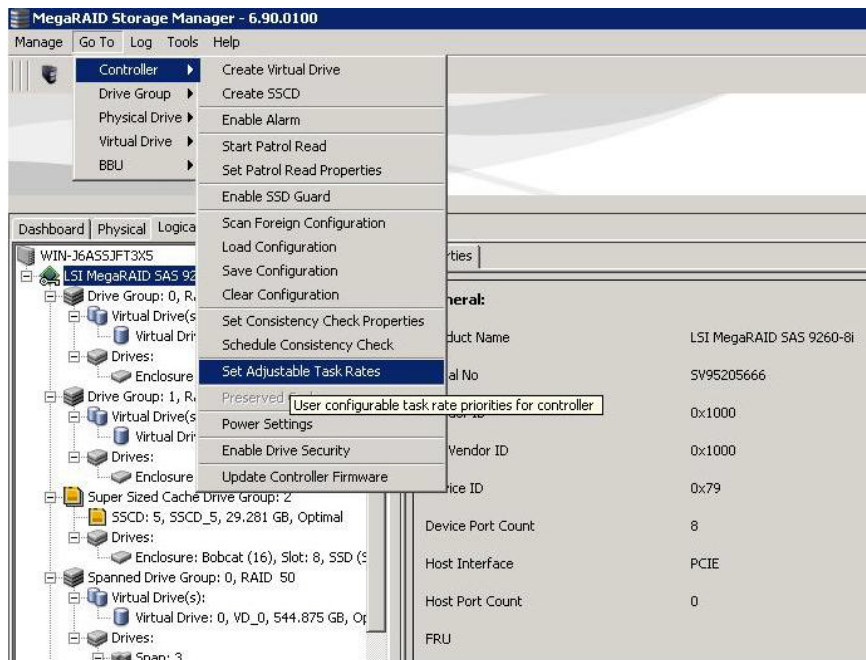
Background Initialization (BGI) Rate: _____

Check Consistency Rate: _____

Follow these steps if you need to change the adjustable rates for rebuilds, and other system tasks that run in the background:

1. Select a controller icon in the **Physical** tab or the **Logical** tab in the left panel of the MegaRAID Storage Manager window.
2. Select **Go To >> Controller >> Set Adjustable Task Rates** from the menu bar, as shown in the following figure.

Figure 7.30 Set Adjustable Task Rates Menu



The default task rates appear in the right panel, as shown in [Figure 7.31](#).

Figure 7.31 Set Adjustable Task Rates

Set Adjustable Task Rates

LSI

Description : User configurable task rate priorities for controller

Rebuild Rate (%) 30

Patrol Rate (%) 30

BGI Rate (%) 30

Check Consistency Rate (%) 30

Reconstruction Rate (%) 30

Ok Cancel

3. Enter changes, as needed, to the following task rates
 - Rebuild Rate
 - Patrol Read
 - Background Initialization (BGI) (for fast initialization)
 - Check Consistency (for consistency checks)
 - Reconstruction

Each task rate can be set from 0 to 100. The higher the number, the faster the activity runs in the background, possibly impacting other system tasks.

4. Click **OK** to accept the new task rates.
5. When the warning message appears, click **OK** to confirm that you want to change the task rates.

7.5 Recovering and Clearing Punctured Block Entries

You can recover/clear the punctured block area of a virtual drive.

Caution: This operation removes any data stored on the physical drives. Back up the good data on the drives before making any changes to the configuration.

When a Patrol Read or a Rebuild operation encounters a media error on the source drive, it punctures a block on the target drive to prevent the use of the data with the invalid parity. Any subsequent read operation to the punctured block completes but with an error. Consequently, the puncturing of a block prevents any invalid parity generation later while using this block.

To recover or clear the punctured block area of a virtual drive, run a Slow (or Full) Initialization to zero out and regenerate new parity, causing all bad block entries to be removed from the bad block table. To run a Slow (or Full) Initialization, see [Section 7.1.1, “Selecting Virtual Drive Settings”](#) or [Section 7.6, “Changing Virtual Drive Properties”](#).

7.6 Changing Virtual Drive Properties

You can change the Read Policy, Write Policy, and other virtual drive properties at any time after a virtual drive is created.

Attention: Do not enable drive caching on a mirrored drive group (RAID 1 or RAID 1E). If you do, data can be corrupted or lost in the event of a sudden power loss. A warning appears if you try to enable drive caching for a mirrored drive group

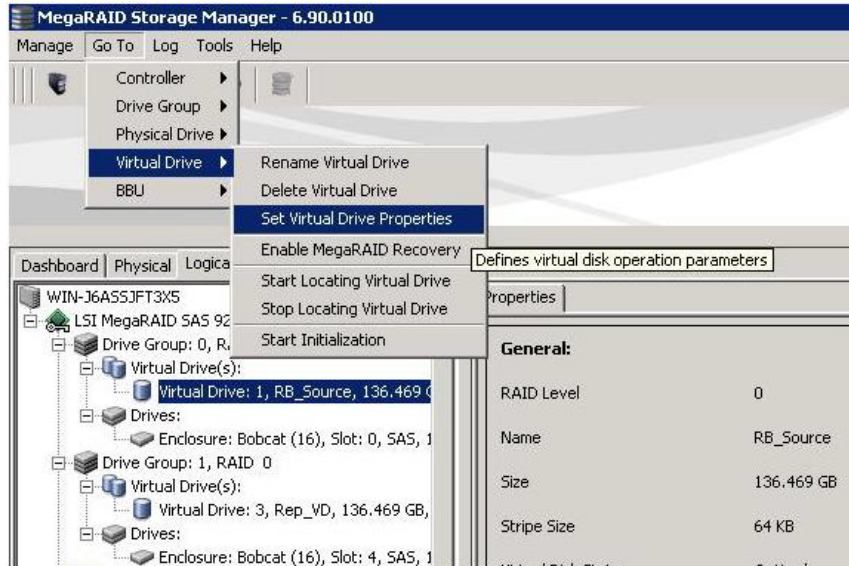
Note: For virtual drives with SAS drives only, set the drive write cache policy set to *Disabled*, by default. For virtual drives with SATA drives only, set the drive write cache policy to *Enabled*, by default.

Follow these steps to change the virtual drive properties.

1. Select a virtual drive icon in the **Physical** tab or the **Logical** tab in the left panel of the MegaRAID Storage Manager window.

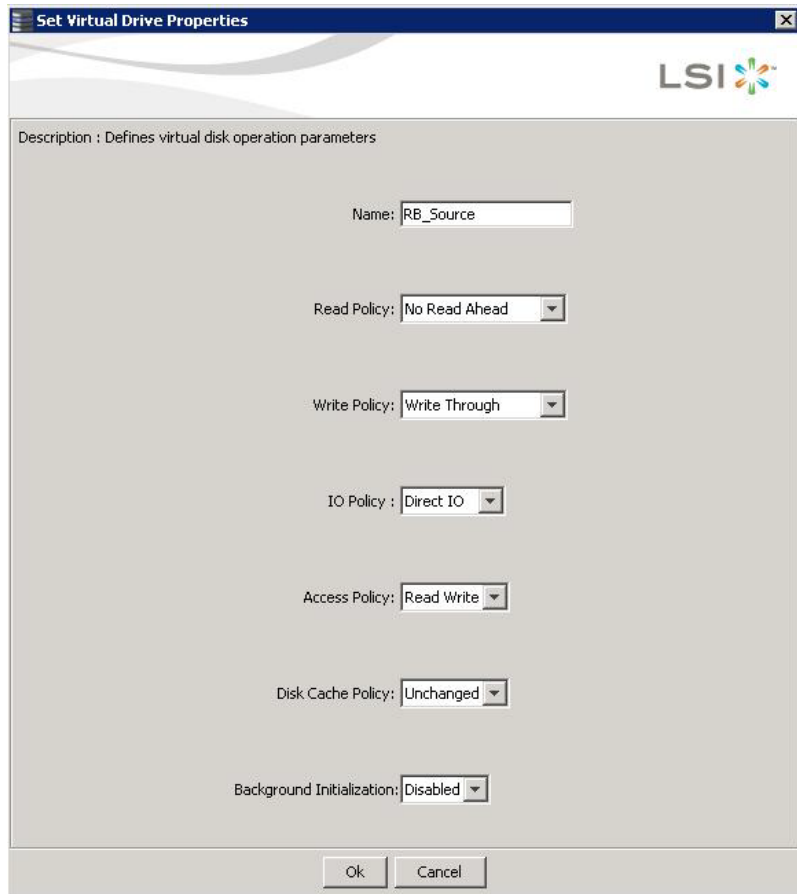
2. Select a virtual drive icon in the left panel of the MegaRAID Storage Manager window.
3. Select **Go To >> Virtual Drive >> Set Virtual Drive Properties** from the menu bar, as shown in the following figure.

Figure 7.32 Set Virtual Drive Properties Menu



The Set Virtual Drive Properties dialog box appears, as shown in [Figure 7.33](#).

Figure 7.33 Set Virtual Drive Properties Screen



The image shows a Windows-style dialog box titled "Set Virtual Drive Properties". In the top right corner is the LSI logo. Below the title bar, a description reads: "Description : Defines virtual disk operation parameters". The main area of the dialog contains several configuration options, each with a label and a text or dropdown field:

- Name: RB_Source
- Read Policy: No Read Ahead
- Write Policy: Write Through
- IO Policy : Direct IO
- Access Policy: Read Write
- Disk Cache Policy: Unchanged
- Background Initialization: Disabled

At the bottom of the dialog are two buttons: "Ok" and "Cancel".

4. Change the virtual drive properties as needed.

For information about these properties, see [Section 7.1.1, "Selecting Virtual Drive Settings"](#).

Note: The ServeRAID M1015 and M1115 SAS/SATA controllers do not support the cache policies. These policies include access policy, read policy, write policy, IO policy, and drive cache.

5. Click **Ok** to accept the changes.

The virtual drive settings are updated.

7.7 Changing a Virtual Drive Configuration

You can use the Modify Drive Group Wizard in MSM to change the configuration of a virtual drive by adding drives to the virtual drive, removing drives from it, or changing its RAID level.

Attention: Be sure to back up the data on the virtual drive before you change its configuration.

Note: You cannot change the configuration of a RAID 10, 50 or 60 virtual drive. You cannot change a RAID 0, 1, 5, or 6 configuration if two or more virtual drives are defined on a single drive group. (The **Logical** view tab shows which drive groups and drives are used by each virtual drive.)

7.7.1 Accessing the Modify Drive Group Wizard

Note: The Modify Drive Group Wizard was previously known as the Reconstruction Wizard.

Perform the following steps to access the Modify Drive Group Wizard options.

1. Click the **Logical** tab in the left panel of the MegaRAID Storage Manager window.
2. Select a drive group in the left panel of the window.
3. Select **Go To >> Drive Group >> Modify Drive Group** from the menu bar, or right-click the virtual drive icon to access the Modify Drive Group Wizard.

The message in the following figure warns about rebooting virtual drives containing boot partitions that are undergoing RAID level migration or capacity expansion operations. Back up your data before you proceed.

Figure 7.34 Reboot Warning Message




4. Click the **Confirm** check box and click **Yes**.

A warning to back up your data appears

5. Click the **Confirm** check box and click **Yes**.

The Modify Drive Group Wizard screen appears, as shown in the following figure.

Figure 7.35 Modify Drive Group Wizard



This section has the following subsections explaining the Modify Drive Group Wizard options:

- [Section 7.7.2, “Adding a Drive or Drives to a Configuration”](#)
- [Section 7.7.3, “Removing a Drive from a Configuration”](#)
- [Section 7.7.4, “Replacing a Drive”](#)
- [Section 7.7.5, “Migrating the RAID Level of a Virtual Drive”](#)

7.7.2 Adding a Drive or Drives to a Configuration

Attention: Be sure to back up the data on the virtual drive before you add a drive to it.

Follow these steps to add a drive or drives to a configuration with the Modify Drive Group Wizard.

1. Click the **Logical** tab in the left panel of the MegaRAID Storage Manager window.
2. Select a drive group in the left panel of the window.
3. Select **Go To >> Drive Group >> Modify Drive Group** on the menu bar, or right-click the virtual drive icon to access the Modify Drive Group Wizard.

The Modify Drive Group Wizard appears, as shown in the following figure.

Figure 7.36 Modify Drive Group Wizard

Modify Drive Group - Select the possible RAID level

LSI

This wizard allows you to modify the drive group configuration by adding drives to the virtual drive or removing drives from it, or changing its RAID level.

Current drive group configuration


Drive group name: Drive Group 4

RAID level: RAID 0

Virtual drive state: Optimal

Drive	Type	Capacity	Status
Enclosure: Bobcat (16), Slot:...	SAS	136.733 GB	Online
Enclosure: Bobcat (16), Slot:...	SAS	136.985 GB	Online
Enclosure: Bobcat (16), Slot:...	SAS	136.985 GB	Online

Select the RAID level you want to migrate:

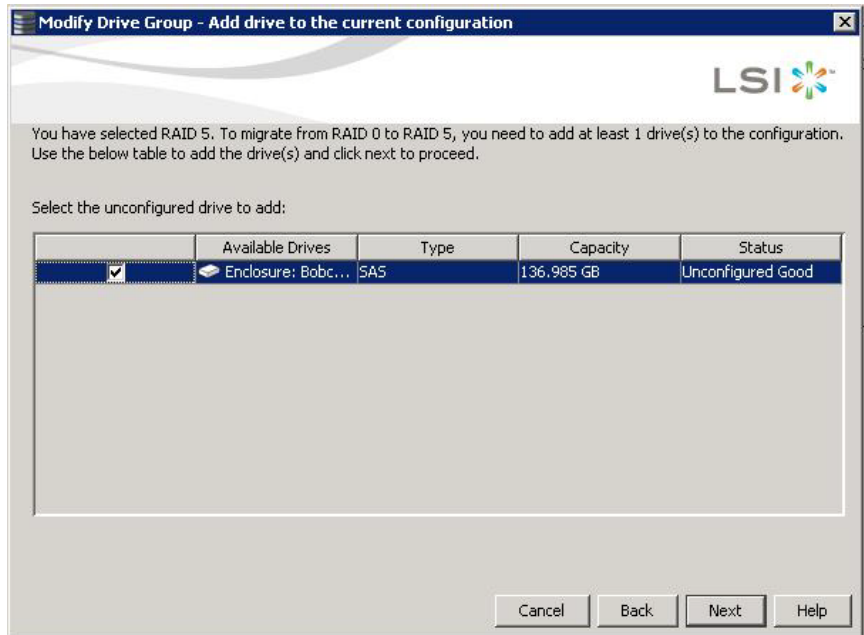
RAID 5  Suitable for multi-user environments(database or file system) with small I/O size and high proportion of read activity.

Cancel Next Help

4. Select the RAID level that you want to change ("migrate") the drive group to and click **Next**.

The following screen appears. It lists the drives you can add and it states whether you must add a minimum number of drives to change the RAID level from the current level to the new RAID level.

Figure 7.37 Add Drive(s) to the Current Configuration Screen



5. Click the check box next to any unconfigured drives that you want to add and then click **Next**.

Note: The drive(s) you add must have the same capacity as or greater capacity than the drives already in the drive group, or you cannot change the RAID level.

The Summary screen appears. This screen shows the current settings and also shows what the settings will be after the drives are added.

Figure 7.38 Modify Drive Group Summary Screen

Review the summary and go back if you need to make corrections. The Changes will be made when you click Finish.

Summary:

Current settings:	Post modification settings:
Drive group name: Drive Group: 4, RAID 0	Drive group name: Drive Group: 4, RAID 5
RAID level: RAID 0	RAID level: RAID 5
Virtual drive name:	Virtual drive name:
Total capacity: 408.656 GB	Total capacity: 408.656 GB
Number of drives: 3	Number of drives: 4

Cancel Back Finish Help

6. Review the configuration information.
Click **Back** if you need to change any selections.
7. Click **Finish** to accept the changes.
A confirmation message appears. The message states that this operation cannot be aborted and asks whether you want to continue.
8. Click **Yes** to accept and complete the addition of the drives to the drive group.

7.7.3 Removing a Drive from a Configuration

Attention: Be sure to back up the data on the virtual drive before you remove a drive from it.

Follow these steps to remove a drive from a RAID 1, RAID 5, or RAID 6 configuration with the Modify Drive Group Wizard.

Note: This option is not available for RAID 0 configurations.

1. Click the **Logical** tab in the left panel of the MegaRAID Storage Manager main menu screen.
2. Click a drive icon in the left panel of the screen.

3. Select **Go To >> Physical Drive >> Make Drive Offline** on the menu bar, or right-click the drive and select **Make Drive Offline** from the menu.

A confirmation message appears. The message states that this operation cannot be aborted and asks whether you want to continue.

4. Click **Yes** to accept and complete the removal of the drive from the drive group.

7.7.4 Replacing a Drive

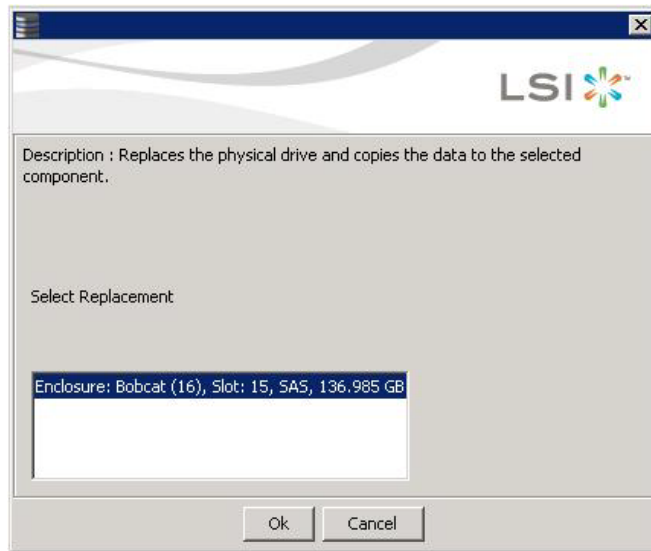
Note: Be sure to back up the data on the virtual drive before you replace a drive.

Follow these steps to add a replacement drive and copy the data from the drive that was removed to the replacement drive.

1. Click the **Logical** tab in the left panel of the MegaRAID Storage Manager window.
2. Select a drive in the left panel of the window.
3. Select **Go To >> Physical Drive >> Replace Physical Drive** on the menu bar, or right-click the virtual drive icon to access the Modify Drive Group Wizard.

The screen with the replacement drive appears, as shown in the following figure.

Figure 7.39 Drive Replacement Window



4. Select a replacement drive.
A confirmation message appears.
5. Click **Yes**.
This replaces a drive and copies the data to the selected component.

7.7.5 Migrating the RAID Level of a Virtual Drive

Attention: Be sure to back up the data on the virtual drive before you change the RAID level.

As the amount of data and the number of drives in your system increase, you can use RAID-level migration to change a virtual drive from one RAID level to another. You do not have to power down or reboot the system when you make this change.

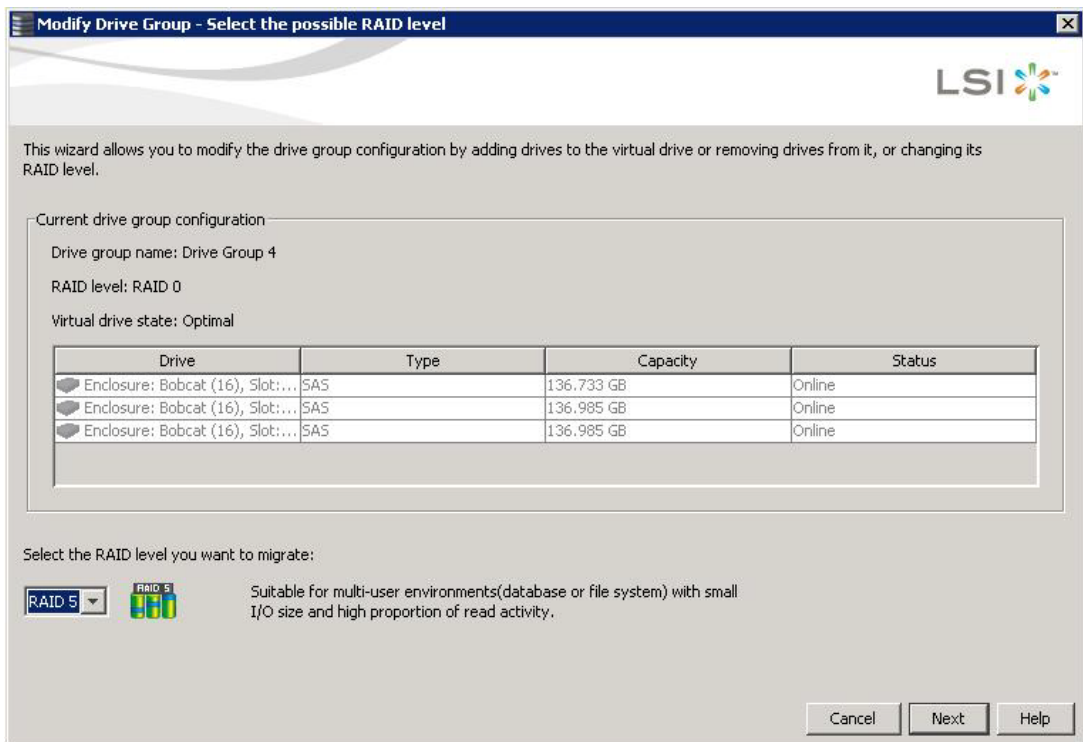
When you migrate a virtual drive to another RAID level, you can keep the same number of drives, or you can add drives. In some cases, you have to add a certain number of drives to migrate the virtual drive from one RAID level to another. The screen indicates the minimum number of drives you are required to add, if any.

Follow these steps to change the RAID level of the virtual drive with the Modify Drive Group Wizard:

1. Click the **Logical** tab in the left panel of the MegaRAID Storage Manager window.
2. Select a drive group in the left panel of the window.
3. Select **Go To >> Drive Group >> Modify Drive Group** on the menu bar, or right-click the virtual drive icon to access the Modify Drive Group Wizard.

The Modify Drive Group Wizard appears, as shown in the following figure.

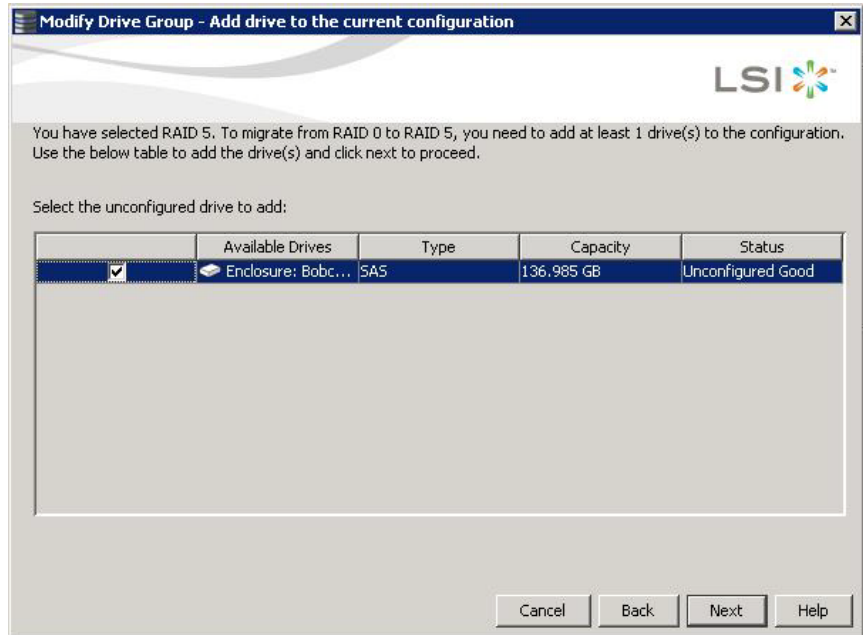
Figure 7.40 Modify Drive Group Wizard



4. On the Modify Drive Group Wizard screen, select the RAID level to which you want to change ("migrate") the drive group and click **Next**.

The following screen appears. The screen states the number of drives that you have to add to change the RAID level from the current level to a new RAID level that require more drives.

Figure 7.41 Add Drive(s) to the Current Configuration Screen



5. Select the unconfigured drive or drives to add and then click **Next**.

Note: The drive(s) you add must have the same capacity as or greater capacity than the drives already in the drive group, or you cannot change the RAID level.

The Summary screen appears. This screen shows the current settings and what the settings will be after the drives are added.

Figure 7.42 Modify Drive Group Summary Screen

Review the summary and go back if you need to make corrections. The Changes will be made when you click Finish.

Summary:

Current settings:	Post modification settings:
Drive group name: Drive Group: 4, RAID 0	Drive group name: Drive Group: 4, RAID 5
RAID level: RAID 0	RAID level: RAID 5
Virtual drive name:	Virtual drive name:
Total capacity: 408.656 GB	Total capacity: 408.656 GB
Number of drives: 3	Number of drives: 4

Cancel Back Finish Help

6. Review the configuration information.

You can click **Back** if you need to change any selections.

7. Click **Finish** to accept the changes.

A confirmation message appears. The message states that this operation cannot be aborted and asks whether you want to continue.

8. Click **Yes** to accept and complete the migration to the new RAID level.

The operation begins on the virtual disk. To monitor the progress of the RAID level change, select **Manage >> Show Progress** in the menu bar.

7.8 Deleting a Virtual Drive

Attention: Be sure to back up the data on the virtual drive before you delete it. Be sure that the operating system is not installed on this virtual drive.

You can delete virtual drives to rearrange the storage space. To delete a virtual drive, follow these steps.

1. Back up all user data on the virtual drive you want to delete.
2. On the MegaRAID Storage Manager main menu screen, select the **Logical** tab, and click the icon of the virtual drive you want to delete.
3. Select **Go To >> Virtual Drive >> Delete Virtual Drive**.
4. When the warning messages appear, click **Yes** to confirm that you want to delete the virtual drive.

Note: You are asked twice if you want to delete a virtual disk to avoid deleting the virtual disk by mistake.

Chapter 8

Monitoring System Events and Storage Devices

The MegaRAID Storage Manager (MSM) software enables you to monitor the status of drives, virtual drives, and other storage devices. This chapter explains how to use MegaRAID Storage Manager software to perform the following monitoring tasks.

- [Section 8.1, “Monitoring System Events”](#)
- [Section 8.2, “Configuring Alert Notifications”](#)
- [Section 8.3, “Monitoring Server Events”](#)
- [Section 8.4, “Monitoring Controllers”](#)
- [Section 8.5, “Monitoring Drives”](#)
- [Section 8.6, “Running a Patrol Read”](#)
- [Section 8.7, “Monitoring Virtual Drives”](#)
- [Section 8.8, “Monitoring Enclosures”](#)
- [Section 8.9, “Monitoring Battery Backup Units”](#)
- [Section 8.10, “Monitoring Rebuilds and Other Processes”](#)

8.1 Monitoring System Events

MegaRAID Storage Manager software monitors the activity and performance of all controllers in the system and the storage devices connected to them. When an event occurs (such as the creation of a new virtual drive or the removal of a drive) an event message appears in the log displayed at the bottom of the MegaRAID Storage Manager window.

You can use MegaRAID Storage Manager to alert you about events. There are settings for the delivery of alerts, the severity level of events, exceptions, and email settings. MSM must be set up in Server mode in order to configure alerts and email settings.

Each message that appears in the event log has a severity level that indicates the importance of the event, as shown in [Table 8.1](#), a date and timestamp, and a brief description. You can click on an event to display the same information in a window. (For a list of all events, see [Section Appendix A, “Events and Messages.”](#))

Table 8.1 Event Severity Levels

Severity Level	Meaning
Information	Informational message. No user action is necessary.
Warning	Some component might be close to a failure point.
Critical	A component has failed, but the system has not lost data.
Fatal	A component has failed, and data loss has occurred or will occur.

The Log menu has five options:

- **Save Log:** Saves the current log to a .log file.
- **Save Log Text:** Saves the current log as a .txt file.
- **Load:** Enables you to load a local .log file in the bottom of the MegaRAID Storage Manager main menu window. If you select the **Load** menu, you will not be able to view the current log.
- **Rollback to Current Log:** This menu appears if we have loaded the logs from a local .log file. Once you select this menu, you can view the current log.
- **Clear Log:** Clears the current log information. You have the option of saving the log first.

8.1.1 System Log

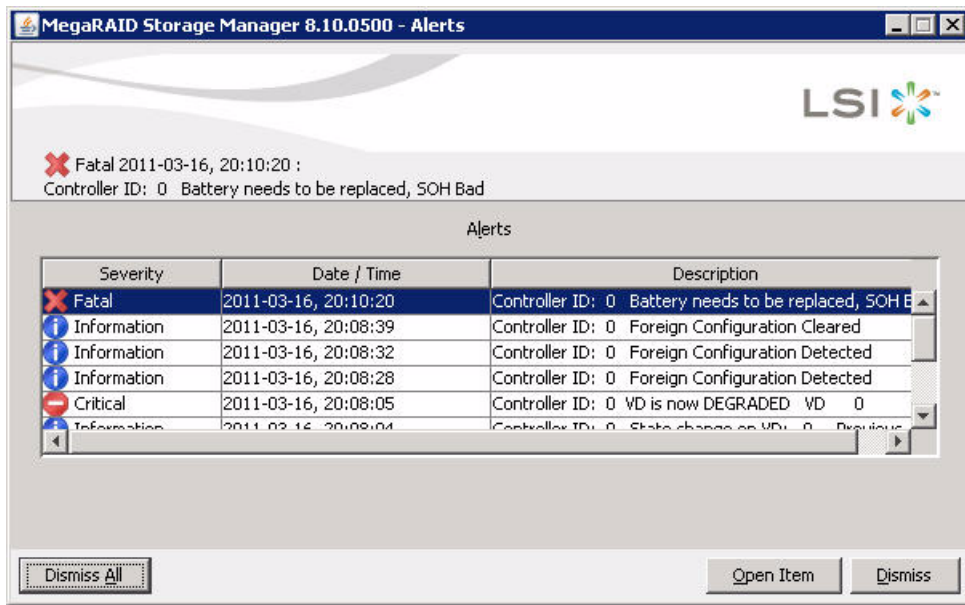
By default, all the severity events are logged in the local syslog. Based on the operating system you are using, the system log is logged in the following syslog locations:

- In Windows, the system log is logged in **Event Viewer >> Application**.
- In Linux, the system log is logged in `/var/log/messages`.
- In Solaris, the system log is logged in `/var/adm/messages`.

8.1.2 Pop-up Notification

By default, fatal and critical events are displaying in a pop-up notification. Pop-up notification is started automatically when you are login in to the operating system. Through this feature, you can view multiple events in a single pop-up window, as shown in the following figure.

Figure 8.1 Pop-Up Notification

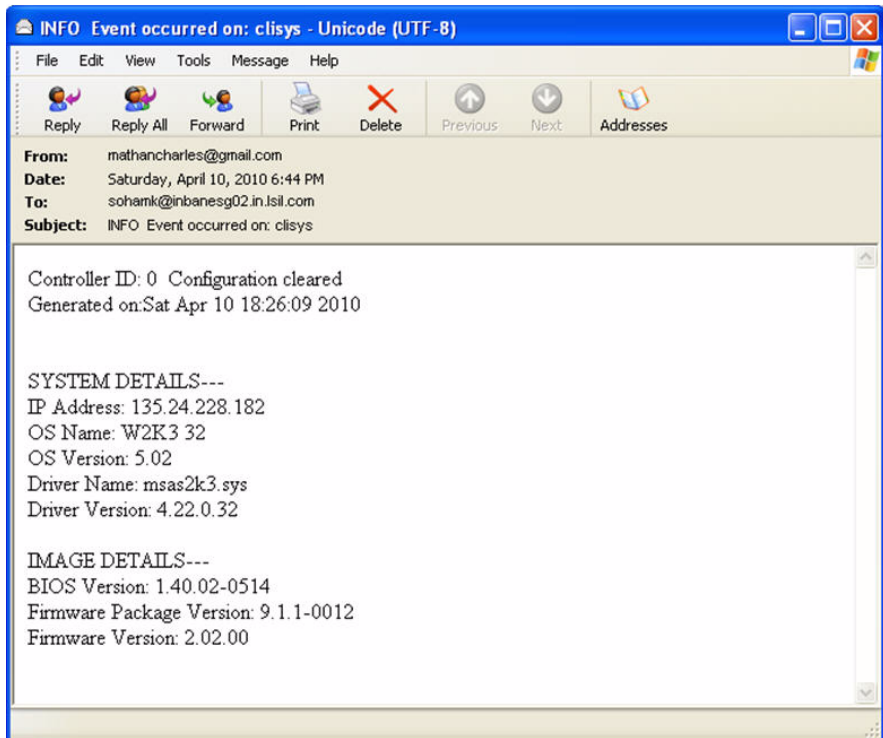


8.1.3 E-mail Notification

By default, fatal events are displayed as e-mail notifications. Based on your configuration, the e-mail notifications are delivered to you as shown in the following figure.

In the e-mail notification, besides the event's description, the email also contains system information and the controller's image details. Using this additional information, you can find out the system and the controller on which the fatal error occurred.

Figure 8.2 E-Mail Notification



8.2 Configuring Alert Notifications

The Alert Notification Configuration utility allows you to control and configure the alerts that MegaRAID Storage Manager software sends when various system events occur. You can use the Event Notification Configuration screen to perform the following functions:

- Set rules for the delivery of alerts
- Change the severity level of events
- Define selected events as exceptions
- Select the settings for the email addresses that are recipients of alert notifications

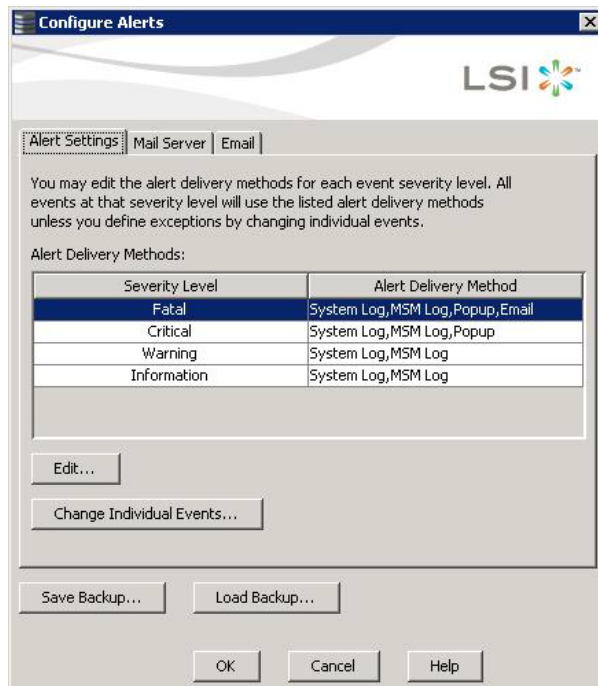
To access this screen, select **Tools >> Configure Alerts** on the main menu screen, as shown in the following figure.

Figure 8.3 Event Notification Configuration Menu



The Configure Alerts screen appears, as shown in the following figure. The screen contains three tabs: **Alert Settings**, **Mail Server**, and **Email**. You can use each tab to perform tasks for that topic.

Figure 8.4 Configure Alerts Screen



Select the **Alert Settings** tab to perform the following actions:

- Edit the alert delivery method for different severity levels.
- Change the method of delivery for each individual event.
- Change the severity level of each individual event.

- Save an `.xml` backup file of the entire alert configuration.
- Load all the values from a previously saved backup into the dialog to edit or save these values as the current alert notification configuration.

Note: When you load a saved backup file, all unsaved changes made in the current session are lost.

You can select the **Mail Server** tab to perform the following actions:

- Enter or edit the sender e-mail address.
- Enter the SMTP server name or the IP address.
- Require authentication of the email server.
- Save an `.xml` backup file of the entire alert configuration.
- Load all of the values from a previously saved backup into the dialog to edit or send to the monitor.

Note: When you load a saved backup file, all unsaved changes made in the current session are lost.

Select the **Email** tab to perform the following actions:

- Add new email addresses for recipients of alert notifications.
- Send test messages to the recipient email addresses.
- Remove email addresses of recipients of alert notifications.
- Save an `.xml` backup file of the entire alert configuration.
- Load all of the values from a previously saved backup into the dialog to edit or send to the monitor.

Note: When you load a saved backup file, all unsaved changes made in the current session are lost.

8.2.1 Setting Alert Delivery Methods

You can select the methods used to send alert deliveries, such as by popup, email, system log, or MSM log. You can select the alert delivery methods for each event severity level (Information, Warning, Critical and Fatal).

Perform the following steps to select the alert delivery methods:

1. On the Configure Alerts screen, click the **Alerts Setting** tab.
2. Under the **Alerts Delivery Methods** heading, select one of the severity levels.
3. Click **Edit**.

The Edit dialog box appears, as shown in the following figure.

Figure 8.5 Alert Notification Delivery Methods Dialog Box



4. Select the desired alert delivery methods for alert notifications at the event severity level.
5. Click **OK** to set the delivery methods used for the severity level that you selected.

8.2.2 Changing Alert Delivery Methods for Individual Events

You can change the alert delivery options for an event without changing the severity level.


1. On the Configure Alerts screen, click the **Alerts Setting** tab.

The **Alerts Setting** portion of the screen appears, as shown in [Figure 8.4](#).

2. Click **Change Individual Events**.

The **Change Individual Events** dialog box appears, as shown in the following figure. The dialog box shows the events by their ID number, description, and severity level.

Figure 8.6 Change Individual Events Dialog Box



3. Click an event in the list to select it.
The current alert delivery methods appear for the selected event under the **Alert Delivery Methods** heading.
4. Select the desired alert delivery methods for the event.
5. Click **OK** to return to the **Configure Alerts** window, or click **Cancel** to discard your current changes and to go back to the **Configure Alerts** window.
6. In the **Configure Alerts** window, click **OK**.

This action saves all of the changes made to the event.

Note: You can click **Restore Defaults** to revert back to the default alert delivery method and the default severity level of an individual event.

8.2.3 Changing the Severity Level for Individual Events

To change the event severity level for a specific event, perform the following steps:

Note: See [Table 8.1](#) for details about the severity levels.

1. On the Configure Alerts screen, click the **Alerts Setting** tab.

The **Alerts Setting** portion of the screen appears.

2. Click **Change Individual Events**.

The **Change Individual Events** dialog box appears, as shown in [Figure 8.6](#). The dialog box shows the events by their ID number, description, and severity level.


3. Click an event in the list to select it.

The current alert delivery methods appear for the selected event.

4. Click the **Severity** cell for the event.

The Event Severity drop-down menu appears for that event, as shown in the following figure.

Figure 8.7 Change Individual Events Severity Level Menu



5. Select a different severity level for the event from the menu.
6. Click **OK** to return to the **Configure Alerts** window, or click Cancel to discard your current changes and to go back to the **Configure Alerts** window.
7. In the **Configure Alerts** window, click **OK**.
This action saves all of the changes made to the events.

8.2.4 Reverting to a Default Individual Event Configuration

To revert to the default alert delivery method and the default severity level of an individual event, perform the following steps:

1. On the **Configure Alerts** window, click the **Alerts Setting** tab.
The Alerts Setting portion of the window appears.
2. Click **Change Individual Events**.
The Change Individual Events dialog appears, as shown in [Figure 8.6](#). The dialog shows the events by their ID number, description, and the severity level.
3. Click **Restore Defaults**.
The Change Individual Events dialog appears with the default alert delivery method and the default severity level of all individual events.
4. Click **OK** to return to the **Configure Alerts** window.
5. In the **Configure Alerts** window, click **OK** to save all the changes made to the events.

8.2.5 Entering or Editing the Sender Email Address and SMTP Server

You can use the **Configure Alerts** screen to enter or edit the sender e-mail address and the SMTP server.

1. On the Configure Alerts screen, click the **Mail Server** tab.
The **Mail Server** options appear, as shown in the following figure.

Figure 8.8 Mail Server Options

The screenshot shows the 'Configure Alerts' window with the 'Mail Server' tab selected. The 'Sender email address' field contains 'monitor@server.com'. The 'SMTP Server' field contains '127.0.0.1'. The checkbox 'This server requires authentication' is checked. Below it are empty fields for 'User name' and 'Password'. At the bottom are buttons for 'Save Backup...', 'Load Backup...', 'OK', 'Cancel', and 'Help'.

2. Enter a new sender email address in the **Sender email address** field or edit the existing sender email address.
3. Enter your SMTP server name/IP Address in the **SMTP Server** field or edit the existing details.
4. Click **OK**.

8.2.6 Authenticating an SMTP Server

You can use the Configure Alerts screen to authenticate the SMTP server, providing an extra level of security. The authentication check box enables the **User name** and **Password** fields when selected by default. Clearing the check box disables these fields.

To enter or modify the SMTP server authentication information, perform the following steps:

1. On the Configure Alerts screen, click the **Mail Server** tab.

The **Mail Server** options appears, as shown in [Figure 8.8](#).

2. If on your SMTP server, the authentication mechanism is enabled and if you want to enable this feature on the MegaRAID Storage Manager software, then you need to select the **This Server requires authentication** check box and enter the authentication details in the corresponding fields (**User name** and **Password**).

If you do not want to enable this feature on the MegaRAID Storage Manager software or if you know that your SMTP server does not support the *Login* mechanism, then de-select the **This Server requires authentication** check box.

Note: The **This Server requires authentication** check box is selected by default.

3. Enter a user name in the **User name** field. (Optional - if **This Server requires authentication** check box is selected).
4. Enter the password in the **Password** field. (Optional - if **This Server requires authentication** check box is selected).
5. Click **OK**.

8.2.7 Saving Backup Configurations

You can save an .xml backup file of the entire alert configuration. This includes all the settings on the Alert Settings, Mail Server, and Email tabs.

1. On the Configure Alerts screen, click the **Alert Setting** tab, **Mail Server** tab, or **Email** tab.
2. Click **Save Backup**.
The drive directory appears.
3. Enter a filename with an .xml extension for the backup configuration (in the format `filename.xml`).
4. Click **Save**.
The drive directory disappears.
5. Click **OK**.

The backup configuration is saved and the Configure Alerts screen closes.

8.2.8 Loading Backup Configurations

You can load all of the values from a previously saved backup into the Configure Alert window (all tabs) to edit or save these values as the current alert notification configuration.

Note: If you choose to load a backup configuration and the Configure Alerts dialog currently contains changes that have not yet been sent to the monitor, the changes will be lost. You are prompted to confirm your choice.

1. On the Configure Alerts screen, click the **Alert Setting** tab, **Mail Server** tab, or **Email** tab.
2. Click **Load Backup**.

A message warns that when you load a saved backup file, all unsaved changes made in the current session will be lost. You are prompted to confirm your choice. Then the drive directory appears from which you can select a backup configuration to load.

3. Select the backup configuration file (it should be in **.xml** format).
4. Click **Open**.

The drive directory disappears.

5. Click **OK**.

The backup configuration is saved and the Configure Alerts screen closes.

8.2.9 Adding Email Addresses of Recipients of Alert Notifications

The **Email** tab portion of the Configure Alerts screen shows the email addresses of recipients of the alert notifications. MegaRAID Storage Manager sends alert notifications to those email addresses. Use the screen to add or remove email addresses of recipients, and to send test messages to recipients that you add.

To add email addresses of recipients of the alert notifications, perform the following steps:

1. Click the **E-mail** tab on the Event Notification Configuration screen.

The **E-mail** section of the screen appears, as shown in the following figure.

Figure 8.9 Email Settings



2. Enter the email address you want to add in the **New recipient email address** field.
3. Click **Add**.
The new email address appears in the **Recipient email addresses** field.

8.2.10 Testing Email Addresses of Recipients of Alert Notifications

Use the **Email** tab portion of the Configure Alerts screen to send test messages to the email addresses that you added for the recipients of alert notifications.

1. Click the **E-mail** tab on the Event Notification Configuration screen.
The **E-mail** section of the screen appears, as shown in [Figure 8.9](#).
2. Click an email address in the **Recipient email addresses** field.
3. Click **Test**.

4. Confirm whether the test message was sent to the email address.
A pop-up message indicates if the test message sent to the email address was successful. If MegaRAID Storage Manager cannot send an email message to the email address, an error message appears.

8.2.11 Removing Email Addresses of Recipients of Alert Notifications

Use the **Email** tab portion of the Configure Alerts screen to remove email addresses of the recipients of alert notifications.


1. Click the **E-mail** tab on the Event Notification Configuration screen.
The **E-mail** section of the screen appears, as shown in [Figure 8.9](#).
2. Click an email address in the **Recipient email addresses** field.
The **Remove** button, which was grayed out, is now active.
3. Click **Remove**.
The email address is deleted from the list.

8.3 Monitoring Server Events

The MegaRAID Storage Manager software enables you to monitor the activity of MegaRAID Storage Manager users in the network.

When a user logs on to or logs off from the application, the event message appears in the log displayed at the bottom of the MegaRAID Storage Manager screen. These event message have a severity level, a date and timestamp (User log on / log off time), and a brief description that contains a user name, client IP address, an access mode (full/view only) and a client system time.

8.4 Monitoring Controllers

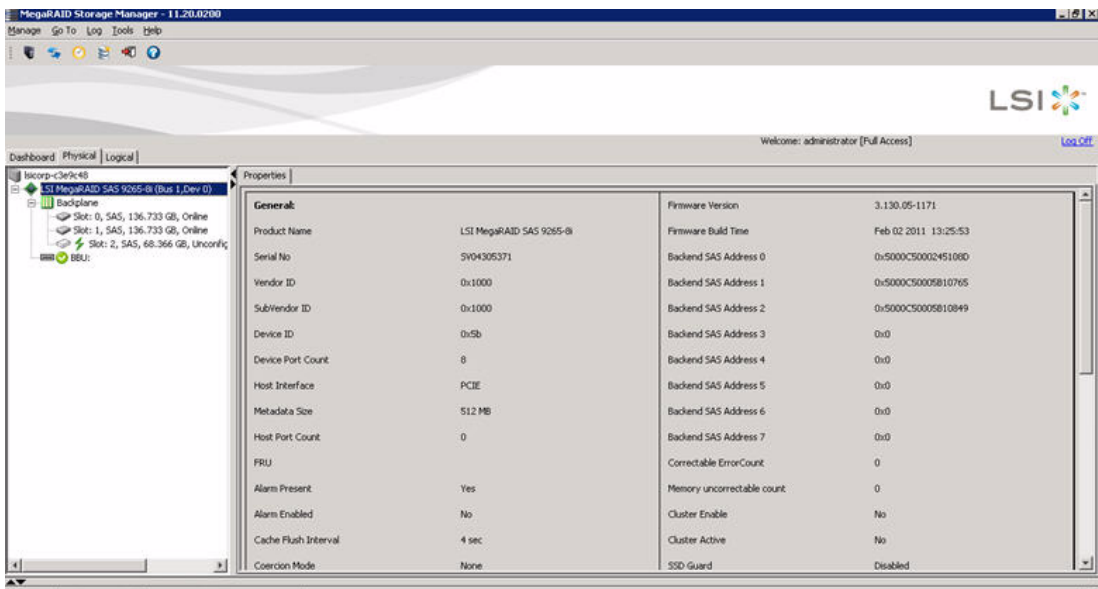
When MegaRAID Storage Manager software is running, you can see the status of all controllers in the left panel of the MegaRAID Storage Manager window. If the controller is operating normally, the controller icon looks like this: . If the controller has failed, a small red circle

appears to the right of the icon. (See [Section 6.5.1, “Dashboard/Physical View/Logical View”](#) for a complete list of device icons.)

To display complete controller information, click a controller icon in the left panel of the MegaRAID Storage Manager window. The controller properties display in the **Properties** tab in the right panel.

The following figure shows the Controller Properties window.

Figure 8.10 Controller Properties





Most of the information on this screen is self-explanatory. Note the following:

- The *Rebuild Rate*, *Patrol Read Rate*, *Reconstruction Rate*, *Consistency Check Rate*, and *BGI Rate* (background initialization) are all user selectable. For more information, see [Section 7.4, “Changing Adjustable Task Rates”](#).
- The *BBU Present* field indicates whether a battery backup unit is installed.
- The *Alarm Present* field and the *Alarm Enabled* field indicate whether the controller has an alarm to alert the user with an audible tone when there is an error or problem on the controller. There are options

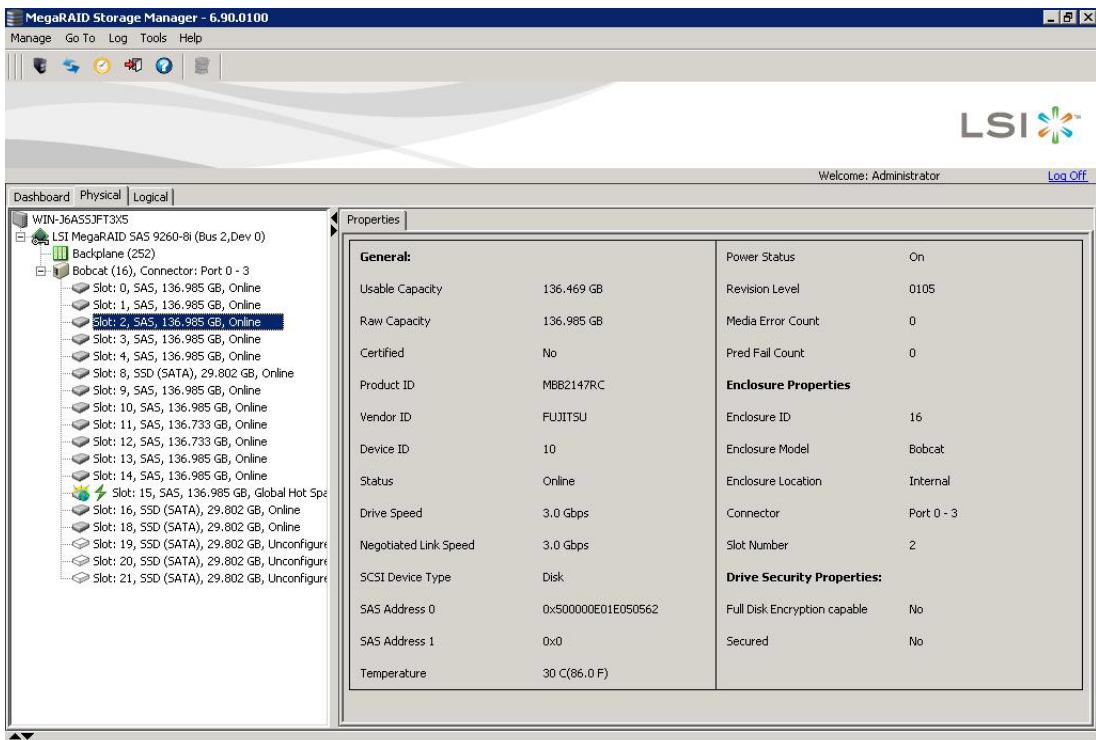
on the controller Properties tab for silencing or disabling the alarm. All controller properties are defined in the [Glossary](#)

8.5 Monitoring Drives

When MegaRAID Storage Manager software is running, you can see the status of all drives in the left panel of the MegaRAID Storage Manager window. If the drive is operating normally, its icon looks like this: . If the drive has failed, a small red circle appears to the right of the icon, like this: . (See [Section 6.5.1, “Dashboard/Physical View/Logical View”](#) for a complete list of device icons.)

To display complete drive information, click a drive icon in the left panel of the MegaRAID Storage Manager window. The drive properties appear in the right panel, as shown in the following figure.

Figure 8.11 Drive Information



The screenshot displays the MegaRAID Storage Manager software interface. The left panel shows a tree view of the storage configuration, including the LSI MegaRAID SAS 9260-8i controller and various drive slots. Slot 2 is selected, showing a 136.985 GB SAS drive in an online state. The right panel displays the properties for this drive, organized into sections: General, Enclosure Properties, and Drive Security Properties.

General:		Power Status	
Usable Capacity	136,469 GB	Power Status	On
Raw Capacity	136,985 GB	Revision Level	0105
Certified	No	Media Error Count	0
Product ID	MBB2147RC	Pred Fail Count	0
Vendor ID	FUJITSU	Enclosure Properties	
Device ID	10	Enclosure ID	16
Status	Online	Enclosure Model	Bobcat
Drive Speed	3.0 gbps	Enclosure Location	Internal
Negotiated Link Speed	3.0 gbps	Connector	Port 0 - 3
SCSI Device Type	Disk	Slot Number	2
SAS Address 0	0x500000E01E050562	Drive Security Properties:	
SAS Address 1	0x0	Full Disk Encryption capable	No
Temperature	30 C(86.0 F)	Secured	No

The information on this panel is self-explanatory. There are no user-selectable properties for physical devices. Icons for other storage devices such as CD-ROM drives and DAT drives can also appear in the left panel.

The **Power Status** property displays the **On** status when a drive is spun up and displays the **Powersave** status when a drive is spun down. Note that SSD drives and other drives that never spin down still show **On**.

If the drives are in a drive enclosure, you can identify which drive is represented by a drive icon on the left. To do this, follow these steps:

1. Click the drive icon in the left panel.
2. Click **Go To >> Physical Drive >> Start Locating Drive**.
3. Select **Locate Physical Drive**, and click **Go**.

The LED on the drive in the enclosure starts blinking to show its location.

Note: LEDs on drives that are global hot spares do not blink.

4. To stop the drive light from blinking, select **Go To >> Physical Drive >> Stop Locating Drive**.

All of the drive properties are defined in the Glossary.

To display a graphical view of a drive, click a drive icon in the left panel of the MegaRAID Storage Manager window, and click the **Graphical View** tab. In Graphical View, the drive's storage capacity is color coded according to the legend shown on the screen: configured space is blue, available space is white, and reserved space is red. When you select a virtual drive from the drop-down menu, the drive space used by that virtual drive is displayed in green.

8.6 Running a Patrol Read

A patrol read periodically verifies all sectors of drives connected to a controller, including the system reserved area in the RAID configured drives. A patrol read can be used for all RAID levels and for all hot spare drives. A patrol read is initiated only when the controller is idle for a defined time period and has no other background activities.

You can set the patrol read properties and start the patrol read operation, or you can start the patrol read without changing the properties.


To start a patrol read, follow these steps:

1. Click a controller icon in the left panel of the MegaRAID Storage Manager window.
2. Select **Go To >> Controller >> Set Patrol Read Properties** or right-click on a controller and select Set Patrol Read Properties from the menu.

To change the patrol read settings, follow these steps:

1. Click the **Logical** tab.
2. Click a controller icon in the left panel of the MegaRAID Storage Manager window.
3. Select the **Operations** tab in the right panel, and select **Set Patrol Read Properties**, as shown in the following figure.

Figure 8.12 Patrol Read Configuration



4. Select an Operation Mode for a patrol read. The options are:
 - **Automatic:** Patrol read runs automatically at the time interval you specify on this screen.
 - **Manual:** Patrol read runs only when you manually start it by selecting **Start Patrol Read** from the controller Options panel.
 - **Disabled:** Patrol read does not run.
5. (Optional) Specify a maximum count of drives to include in the patrol read.

The count must be a number from 1 to 255.

6. Click virtual drives in the list under the heading Virtual Drives to include in the patrol read and click **Add >** or click **Add All >>** to include all of the virtual drives.
7. (Optional) Change the frequency at which the patrol read will run.

The default frequency is weekly (168 hours), which is suitable for most configurations. (You can select hourly, daily, or monthly as the unit of measurement.)

Note: Leave the patrol read frequency and other patrol read settings at the default values to achieve the best system performance. If you decide to change the values, record the original default value here so you can restore them later, if necessary:

Patrol Read Frequency: _____

Continuous Patrolling: Enabled/Disabled

Patrol Read Task Rate: _____

8. (Optional) Set Patrol Read to run at a specific time.

The default is for the patrol read to start when you click **OK** on this screen. To change the default so that the patrol read starts at a specific time, follow these steps (otherwise, skip this step and proceed to the next step):

 - a. Uncheck the box **Perform Patrol Read when I click OK.**
 - b. Select the month, year, day, and time to start patrol read.
9. Click **OK** to enable your patrol read selections.

Note: Patrol read does not report on its progress while it is running. The patrol read status is reported in the event log only.

10. Click **Go** to enable these Patrol Read options.

To start a patrol read without changing the patrol read properties, follow these steps:


1. Click a controller icon in the left panel of the MegaRAID Storage Manager main menu screen.
2. Select **Go To >> Controller >> Start Patrol Read** in the menu bar or right-click a controller and select Start Patrol Read from the menu.
3. When prompted, click **Yes** to confirm that you want to start a patrol read.

8.6.1 Patrol Read Task Rates

You have the option to change the patrol read *task rate*. The task rate determines the amount of system resources that are dedicated to a patrol read when it is running. LSI recommends, however, that you leave the patrol read task rate at its default setting.

If you raise the task rate above the default, foreground tasks will run more slowly and it might seem that the system is not responding. If you lower the task rate below the default, rebuilds and other background tasks might run very slowly and might not complete within a reasonable time. For more information, about the patrol read task rate, see [Section 7.4, “Changing Adjustable Task Rates”](#).

8.7 Monitoring Virtual Drives

When MegaRAID Storage Manager software is running, you can see the status of all virtual drives. If a virtual drive is operating normally, the icon looks like this: . Color-coded circles appear next to the icon to indicate the following statuses:

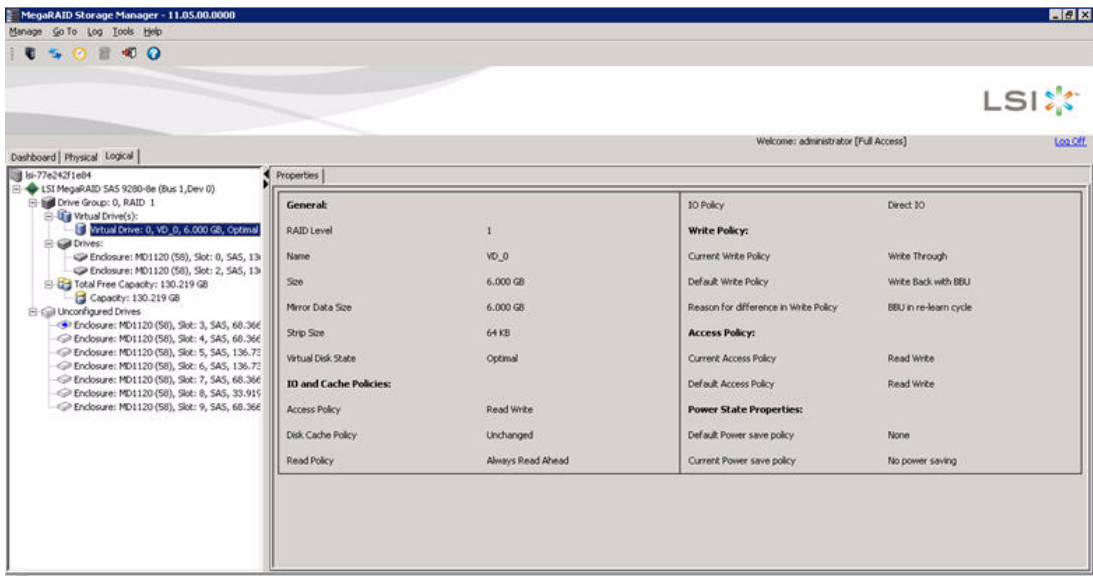
- Green: The server is operating properly.

- Yellow: The server is running in a partially degraded state (for example, if a drive has failed); the data is still safe, but data could be lost if another drive fails.
- Orange: The server is running in a degraded state.
- Red: The server storage configuration has failed.

When the Logical tab is selected, the left panel of the MegaRAID Storage Manager window shows which drives are used by each virtual drive. The same drive can be used by multiple virtual drives.

To display complete virtual drive information, click the **Logical** tab in the left panel and click a virtual drive icon in the left panel. The properties appear in the right panel. All virtual drive properties are defined in [Appendix C, “Glossary”](#). the following figure shows the Properties panel for a virtual drive.

Figure 8.13 Virtual Drive Properties



The RAID level, stripe size, and access policy of the virtual drive are set when it is configured.

Note: You can change the read policy, write policy, and other virtual drive properties. See [Section 7.6, “Changing Virtual](#)

[Drive Properties](#)” for the procedure you can use to change these properties.

Note: The ServeRAID M1015 and M1115 SAS/SATA controllers do not support the cache policies. These policies include access policy, read policy, write policy, IO policy, and drive cache.

If the drives in the virtual drive are in an enclosure, you can identify them by making their LEDs blink. To do this, follow these steps:


1. Click the virtual drive icon in the left panel.
2. Click **Go To >> Virtual Drive >> Start Locating Virtual Drive** or right-click a virtual drive and select **Start Locating Virtual Drive** from the menu.

3. Select **Locate Virtual Drive**, and click **Go**.

The LEDs on the drives in the virtual drive start blinking (except for hot spare drives).

4. To stop the LEDs from blinking, select **Go To >> Virtual Drive >> Stop Locating Virtual Drive** or right-click a virtual drive and select **Stop Locating Virtual Drive** from the menu.

8.8 Monitoring Enclosures

When MegaRAID Storage Manager software is running, you can see the status of all enclosures connected to the server by selecting the **Physical** tab in the left panel. If an enclosure is operating normally, the icon looks like this: . If the enclosure is not functioning normally—for example, if a fan has failed—a small yellow or red circle appears to the right of the icon.

Information about the enclosure appears in the right panel when you select the **Properties** tab on the main menu screen. A graphical display of enclosure information appears when you select the **Graphical View** tab.

The display in the center of the screen shows how many slots of the enclosure are actually populated by the drives and the lights on the drives show the drive status. The information on the right shows you the


status of the temperature sensors, fans, and power supplies in the enclosure.

To view the enclosure properties, in the physical view click on the **Enclosure** node. The **Enclosure Properties** appear, as shown in the following figure.

Figure 8.14 Enclosure Properties

Vendor ID	DELL	FRU Number	41R5133
Enclosure ID	IS	Part Number	CP-111-006-020
Enclosure Type	SES	Component Properties	
Enclosure Model	MD1000U	Number of Temperature Sensors	4
Enclosure Location	External	Number of Fans	4
Firmware Version	A.04	Number of Power Supplies	2
Serial Number	0802V16VTE	Number of Voltage Sensors	0
Connector	Port A		
Number of Slots	15		

8.9 Monitoring Battery Backup Units

When MegaRAID Storage Manager software is running, you can monitor the status of all of the BBUs connected to controllers in the server. If a BBU is operating normally, the icon looks like this: . If it has failed, a red dot appears next to the icon.

To show the properties for a BBU, perform the following steps:

1. On the main menu screen, click the **Physical** tab to open the physical view.
2. Select the BBU icon in the left panel.

The BBU properties, which appear in the right panel, include the following:

- The number of times the BBU has been recharged (Cycle Count)
- The full capacity of the BBU, plus the percentage of its current state of charge, and the estimated time until it is depleted

- The current BBU temperature, voltage, current, and remaining capacity
- If the battery is charging, the estimated time until it is fully charged
- The battery state, which says if it is in operational state.
- If battery replacement is required.
- The BBU retention time, which gives the total number of hours the battery can support the current capacity reserve.

The BBU Properties are displayed, as shown in the following two figures.

Figure 8.15 Battery Backup Unit Properties for iBBU Battery


Properties			
BBU Battery Type	iBBU	Cycle Count	32
Battery State	Operational	Automatic Learn Cycle	Enabled
Battery Replacement	Required / Not required	Auto Learn Period	30 days
Temperature	29.0 C (84.2 F) - Normal	Next Learn Cycle	Aug 10 2010 20:52:13
Voltage	4065 mV	Relative State of Charge	99%
Current	0 mA	Absolute State of Charge	95%
Full Capacity	<value> mAh	Run Time to Empty	Battery is not being discharged
Remaining Capacity	<value> mAh	Average Time to Empty	Battery is not being discharged
BBU Retention Time	 48+ Hours	Average Time to Full	Battery is not being discharged
Estimated Time to Recharge	<value> Mins	Maximum Error Margin	25%
FRU	None		

Figure 8.16 Battery Backup Unit Properties for TMM-C Battery

Properties			
BBU Battery Type	TMM-C (Not activated) 2	Estimated Time to Recharge	<value> Mins
Battery State	Operational	FRU	None
Battery Replacement	Required / Not required	Memory Module FRU	<value>
Temperature	29.0 C (84.2 F) - Normal	Automatic Learn Cycle	Enabled
Voltage	4055 mV	Auto Learn Period	30 days
Current	0 mA	Next Learn Cycle	Aug 10 2010 20:52:13
Full Capacity	<value> Joules		
Remaining Capacity	<value> Joules		
BBU Retention Time	48+ Hours		

8.9.1 Battery Learn Cycle

Learn Cycle is a battery calibration operation that the controller performs periodically to determine the condition of the battery. You can start battery learn cycles manually or automatically. To choose automatic battery learn cycles, enable automatic learn cycles. To choose manual battery learn cycles, disable automatic learn cycles.

If you enable automatic learn cycles, you can delay the start of the learn cycles for up to 168 hours (7 days). If you disable automatic learn cycles, you can start the learn cycles manually, and you can choose to receive a reminder to start a manual learn cycle.

8.9.1.1 Setting Learn Cycle Properties

To set the learn cycle properties, perform the following steps:

1. Click the **Physical** tab to open the physical view.
2. Select the BBU icon in the left panel.
3. Click the **Go To >> BBU >> Set Learn Cycle Properties**.

The BBU operations screen appears.

4. Select one of the two automatic learn cycles:
 - Select the **Enable** radio button to enable an automatic learn cycle.

- Select the **Disable** radio button to disable an automatic learn cycle.
- 5. You can delay the start of the next learn cycle up to 7 days (168 hours) by specifying the hours in the **Delay scheduled learn cycle by** field.
- 6. Select the **Remind me when to start a learn cycle** check box to receive a reminder to start a manual learn cycle.

Note: After you select **Disable**, if you then select **Enable**, the controller firmware resets the battery module properties to initiate an immediate battery learn cycle. The **Next Learn cycle** field is not updated until the battery relearn is completed. Once the relearning cycle is completed, the value in the **Next Learn cycle** field displays the new date and the time of the next battery learning cycle.

8.9.1.2 Starting a Learn Cycle Manually

To start the learn cycle properties manually, perform the following steps:

1. Click the **Physical** tab to open the Physical view.
2. Select the BBU icon in the left panel.
3. Click the **Go To >> BBU >> Start Learn Cycle**.

Alternatively, right-click the BBU icon to open the operations menu and select **Start Learn Cycle**.

8.10 Monitoring Rebuilds and Other Processes

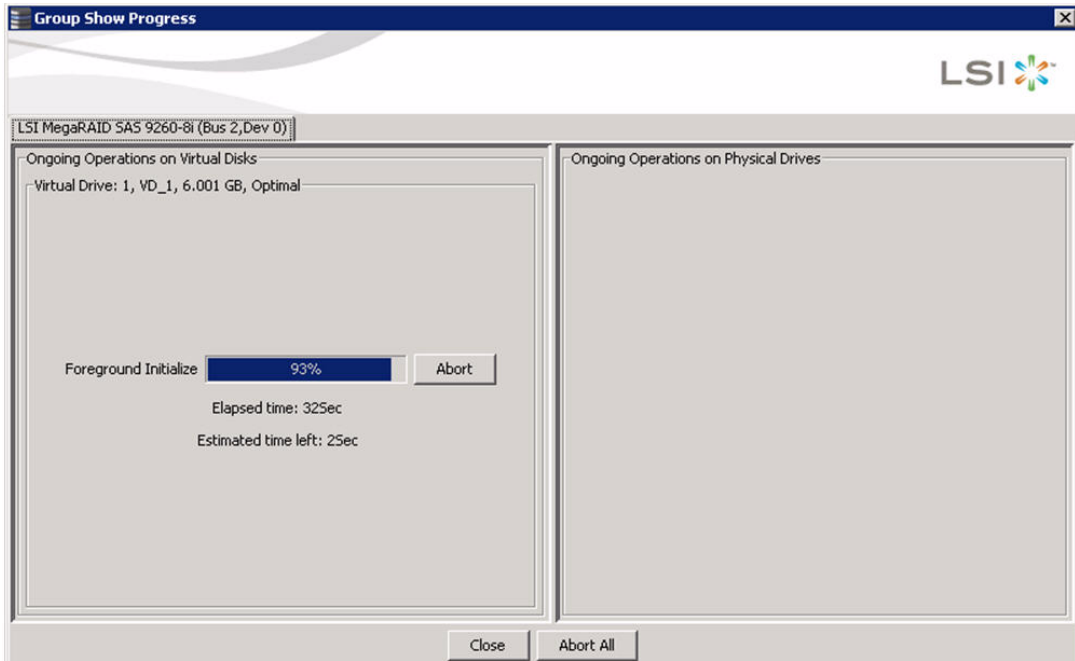
MegaRAID Storage Manager software allows you to monitor the progress of rebuilds and other lengthy processes in the Group Show Progress window.

Follow these steps to monitor the progress of these operations.

1. Select **Manage >> Show Progress** on the menu bar.

The Group Show Progress window appears, as shown in the following figure.

Figure 8.17 Group Show Progress Window



The Group Show Progress window displays a percent-complete indicator for drive rebuilds. Rebuilds can take a long time to complete and they cannot be aborted. An up-arrow appears above the drive icon while it is being rebuilt.

Operations on virtual drives appear in the left panel of the Group Show Progress window, and operations on drives appear in the right panel. The following kinds of operations appear in this window:

- Background or foreground initialization of a virtual drive (see [Section 9.1, “Initializing a Virtual Drive”](#))
- Rebuild (see [Section 9.5, “Rebuilding a Drive”](#))
- Drive group modification (see [Section 7.7.1, “Accessing the Modify Drive Group Wizard”](#))
- Consistency check (see [Section 9.3, “Running a Consistency Check”](#))
- Physical or Virtual Drive Erase (see [Section 6.5.8, “Non-SED Secure Erase Support”](#))

A Modify Drive Group process cannot be aborted. To abort any other ongoing process, click the **Abort** button next to the status indicator. Click **Abort All** to abort all ongoing processes. Click **Close** to close the window.

Chapter 9

Maintaining and Managing Storage Configurations

This section explains how to use MegaRAID Storage Manager software to maintain and manage storage configurations. You must log on to the server in Full Access mode to perform maintenance and management tasks. This chapter explains how to perform the following tasks:

- [Section 9.1, “Initializing a Virtual Drive”](#)
 - [Section 9.2, “Running a Group Initialization”](#)
 - [Section 9.3, “Running a Consistency Check”](#)
 - [Section 9.4, “Scanning for New Drives”](#)
 - [Section 9.5, “Rebuilding a Drive”](#)
 - [Section 9.6, “Making a Drive Offline or Missing”](#)
 - [Section 9.7, “Removing a Drive”](#)
 - [Section 9.8, “Upgrading the Firmware”](#)
-

9.1 Initializing a Virtual Drive

When you create a new virtual drive with the Configuration Wizard, you can select the **Quick Init** or **Full Init** option to initialize the disk immediately. However, you can select **No Init** if you want to initialize the virtual drive later.

To initialize a virtual drive after completing the configuration process, follow these steps:

1. Select the **Logical** tab in the left panel of the MegaRAID Storage Manager window, and click the icon of the virtual drive that you want to initialize.
2. Select **Go To >> Virtual Drive >> Start Initialization**.

The initialize dialog appears.

3. Select the virtual drive(s) to initialize.

Attention: Initialization erases all data on the virtual drive. Be sure to back up any data you want to keep before you initialize. Be sure the operating system is not installed on the virtual drive you are initializing.

4. Select the **Fast Initialization** check box if you want to use this option.

If you leave the box unchecked, MegaRAID Storage Manager software will run a Full Initialization on the virtual drive. (For more information, see [Section 7.1.1, “Selecting Virtual Drive Settings.”](#))

5. Click **Start** to begin the initialization.

You can monitor the progress of the initialization. See [Section 8.10, “Monitoring Rebuilds and Other Processes”](#) for more information.


9.2 Running a Group Initialization

Initialization prepares the storage medium for use. You can run an initialization on multiple drives at one time. Follow these steps to run a group initialize.

1. Select **Manage >> Initialize**.

The Group Initialization dialog appears, as shown in the following figure.

Figure 9.1 Group Initialization Dialog Box



2. Either check the virtual drives on which to run the initialization, or click **Select All** to select all of the virtual drives.
3. Click **Start**.

You can monitor the progress of the group initialization. See [Section 8.10, “Monitoring Rebuilds and Other Processes”](#) for more information.

9.3 Running a Consistency Check

You should periodically run a consistency check on fault-tolerant virtual drives (RAID 1, 5, 6, 10, 50, or 60 configurations; RAID 0 does not provide data redundancy). A consistency check scans the virtual drive to determine whether the data has become corrupted and needs to be restored.

For example, in a system with parity, checking consistency means computing the data on one drive and comparing the results to the contents of the parity drive. You must run a consistency check if you suspect that the data on the virtual drive might be corrupted.

Note: Make sure to back up the data before running a consistency check if you think the data might be corrupted.


To run a consistency check, first set the consistency check properties, and then schedule the consistency check. This section explains how to set the properties, schedule the check, and run the consistency check.

9.3.1 Setting the Consistency Check Properties

Follow these steps to set the properties for a consistency check:

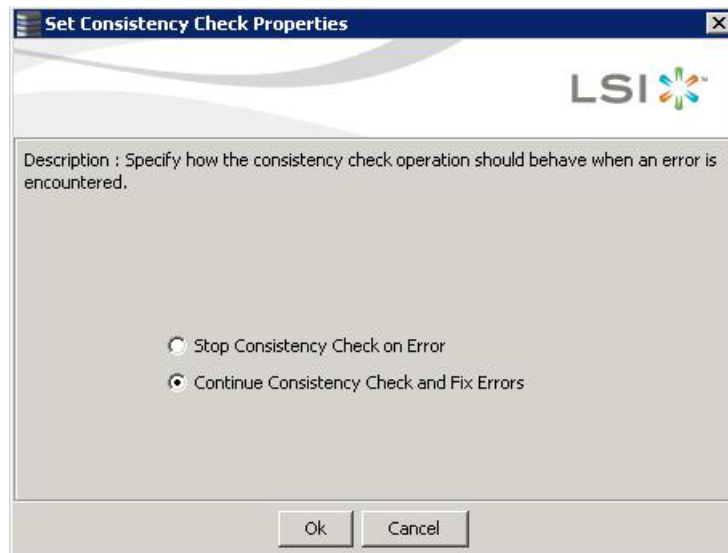
1. Click the **Physical** tab or the **Logical** tab and select a controller.
2. Click **Go To >> Controller >> Set Consistency Check Properties**, as shown in the following figure.

Figure 9.2 Set Consistency Check Properties Option



The Set Consistency Check Properties dialog appears, as shown in the following figure.

Figure 9.3 Set Consistency Check Properties Dialog Box



3. Choose one of the two options:
 - **Stop Consistency Check on Error:** The RAID controller stops the consistency check operation if the utility finds an error.

- **Continue Consistency Check and Fix Errors:** The RAID controller continues the consistency check if the utility finds an error, and then fixes the errors.

4. Click **Ok**.

9.3.2 Scheduling a Consistency Check

Follow these steps to set the properties for a consistency check:

1. Click the **Physical** tab or the **Logical** tab and select the controller.
2. Select **Go To >> Controller >> Schedule Consistency Check**.

The Schedule Consistency Check dialog appears, as shown in the following figure.

Figure 9.4 Schedule Consistency Check Dialog

Schedule Consistency Check

LSI

Description : Establish schedule for consistency check operation.

Run consistency check:

Weekly

☐ Run consistency check continuously

Start on:

March 20 2010

Start time:

03:00 AM

3. Select how often to run the consistency check from the drop-down list.
Click **Advanced** for more detailed date options.
4. (Optional) Select the **Run consistency check continuously** check box.
5. Select the month, day, and year on which to start the consistency check.
6. Select the time of day to start the consistency check.
7. Click **Ok**.

You can monitor the progress of the consistency check. See [Section 8.10, "Monitoring Rebuilds and Other Processes"](#) for more information.


9.3.3 Running a Group Consistency Check

You can run a consistency check on multiple drives at one time. Follow these steps to run a group consistency check.

1. Select **Manage >> Check Consistency**.

The Group Consistency Check dialog appears, as shown in the following figure.

Figure 9.5 Group Consistency Check Dialog Box



2. Either check the virtual drives on which to run the consistency check, or click **Select All** to select all of the virtual drives.
3. Click **Start**.

You can monitor the progress of the group consistency check. See [Section 8.10, “Monitoring Rebuilds and Other Processes”](#) for more information.

9.4 Scanning for New Drives

You can use the Scan for Foreign Configuration option to find drives with foreign configurations. A foreign configuration is a RAID configuration that already exists on a replacement set of physical disks that you install in a computer system. In addition, if one or more drives are removed from a configuration, by a cable pull or drive removal, for example, the configuration on those drives is considered a foreign configuration by the RAID controller. Drives that are foreign are listed on the physical drives list with a special symbol in the MegaRAID Storage Manager software.

The utility allows you to import the existing configuration to the RAID controller or clear the configuration so you can create a new configuration using these drives. You can preview the foreign configuration before you decide whether to import it.

The MegaRAID Storage Manager software usually detects newly installed drives and displays icons for them in the MegaRAID Storage Manager window. If for some reason the MegaRAID Storage Manager software does not detect a new drive (or drives), you can use the Scan for Foreign Configuration command to find it.

Follow these steps to scan for a foreign configuration:

1. Select a controller icon in the left panel of the MegaRAID Storage Manager window.
2. Select **Go To >> Controller >> Scan for Foreign Configuration**.
If MegaRAID Storage Manager software detects any new drives, it displays a list of them on the screen. If not, it notifies you that no foreign configuration is found.
3. Follow the instructions on the screen to complete the drive detection.



9.5 Rebuilding a Drive

If a drive in a redundant virtual drive (RAID 1, 5, 6, 10, 50, or 60) fails, the MegaRAID Storage Manager software automatically rebuilds the data on a hot spare drive to prevent data loss. *The rebuild is a fully automatic process*, so it is not necessary to issue a Rebuild command. You can

monitor the progress of drive rebuilds in the Group Show Progress window. To open this window, select **Group Operations >> Show Progress**.

If a single drive in a RAID 1, RAID 5, RAID 10, or RAID 50 virtual drive fails, the system is protected from data loss. A RAID 6 virtual drive can survive two failed drives. A RAID 60 virtual drive can survive two failed drives in each span in the drive group. Data loss is prevented by using parity data in RAID 5, RAID 6, RAID 50, and RAID 60, and data redundancy in RAID 1 and RAID 10.

The failed drive must be replaced, and the data on the drive must be rebuilt on a new drive to restore the system to fault tolerance. You can choose to rebuild the data on the failed drive if the drive is still operational. If dedicated hot spares or global hot spare disks are available, the failed drive is rebuilt automatically without any user intervention.

If a drive has failed, a red circle appears to the right of the drive icon: . A small yellow circle appears to the right of the icon of the virtual drive that uses this drive: . This icon indicates that the virtual drive is in a degraded state; the data is still safe, but data could be lost if another drive fails.

Follow these steps to rebuild a drive:

1. Right-click the icon of the failed drive, and select **Rebuild**.
2. Click **Yes** when the warning message appears. If the drive is still good, a rebuild will start.

You can monitor the progress of the rebuild in the Group Show Progress window by selecting **Manage >> Show Progress**. If the drive cannot be rebuilt, an error message appears. Continue with the next step.

3. Shut down the system, disconnect the power cord, and open the computer case.
4. Replace the failed drive with a new drive of equal capacity.
5. Close the computer case, reconnect the power cord, and restart the computer.
6. Restart the MegaRAID Storage Manager software.

When the new drive spins up, the drive icon changes back to normal status, and the rebuild process begins automatically. You can monitor the progress of the rebuild in the Group Show Progress window by selecting **Manage >> Show Progress**.

If you want to force a drive into Fail status to trigger a rebuild, right-click the drive icon, and select **Make Drive Offline**. A red circle appears next to the drive icon. Right-click the icon, and select **Rebuild** from the pop-up menu. A drive rebuild cannot be aborted.

Note: A drive rebuild is also started if you select **Make Drive Online** from the pop-up menu.

9.6 Making a Drive Offline or Missing

If a drive is currently part of a redundant configuration and you want to use it in another configuration, you can use MegaRAID Storage Manager commands to remove the drive from the first configuration. When you perform this action, *all data on that drive is lost*.

To remove the drive from the configuration without harming the data on the virtual drive, follow these steps:

1. In the MegaRAID Storage Manager window, select **Go To >> Physical Drive >> Make Drive Offline**.

The drive status changes to Offline.

2. Select **Go To >> Physical Drive >> Mark Drive as Missing**.

The drive status changes to Unconfigured Good.

Attention: After you perform this step, the data on this drive is no longer valid.

3. If necessary, create a hot spare drive for the virtual drive from which you have removed the drive.

When a hot spare is available, the data on the virtual drive will be rebuilt. You can now use the removed drive for another configuration.

Attention: If MegaRAID Storage Manager software detects that a drive in a virtual drive has failed, it makes the drive offline. If this happens, you must remove the drive and replace it.

You cannot make the drive usable for another configuration by using the **Mark physical disk as missing** command and the **Rescan** command.

9.7 Removing a Drive

You might need to remove a non-failed drive that is connected to the controller. For example, you might need to replace the drive with a larger drive. Follow these steps to remove a drive safely:

1. Click the icon of the drive in the left panel, and click the **Operations** tab in the right panel.
2. Select **Prepare for Removal**, and click **Go**.
3. Wait until the drive spins down and remove it.

If you change your mind, select **Undo Prepare for Removal** and click **Go**.

9.8 Upgrading the Firmware

MegaRAID Storage Manager software enables you to easily upgrade the controller firmware.

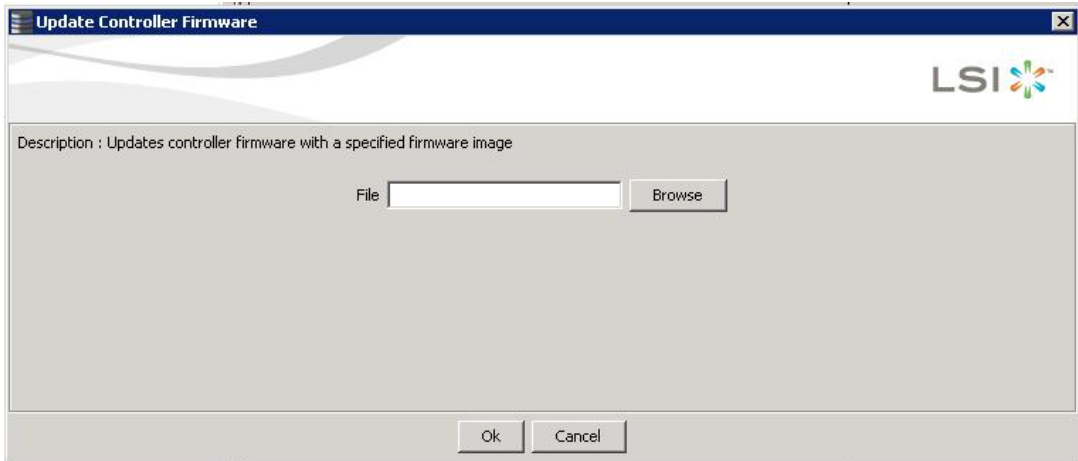
To avoid data loss because of dirty cache on the controller, the utility forces the virtual disks into Write Through mode after a firmware upgrade. It remains in this mode until the server reboots. In Write Through mode, the controller sends a data transfer completion signal to the host when the disk subsystem has received all of the data in a transaction. This ensures that if a power outage occurs, the controller does not discard the dirty cache.

Follow these steps to upgrade the firmware:

1. In the left panel of the MegaRAID Storage Manager window, click the icon of the controller you need to upgrade.
2. In the MegaRAID Storage Manager screen, select **Go To >> Controller >> Update Controller Firmware**.

3. Click **Browse** to locate the .rom update file, as shown in the following figure.

Figure 9.6 Update Controller Firmware Dialog



4. After you locate the file, click **Ok**.
The MegaRAID Storage Manager software displays the version of the existing firmware and the version of the new firmware file.
5. When you are prompted to indicate whether you want to upgrade the firmware, click **Yes**.
The controller is updated with the new firmware code contained in the .rom file.
6. Reboot the system after the new firmware is flashed.
The new firmware does not take effect until reboot.

Chapter 10

CacheCade 2.0

Software

MegaRAID CacheCade Pro 2.0 read/write software eliminates the need for manually configured hybrid arrays by intelligently and dynamically managing frequently-accessed data and copying it from HDD volumes to a higher performance layer of SSD cache. Copying the most accessed data to flash cache relieves the primary HDD array from time-consuming transactions, which allows for more efficient hard disk operation, reduced latency, and accelerated read and write speeds.

This feature provides significant improvements to overall system performance – two to twelve times that of HDD-only configurations – for a wide variety of server applications including web, file, online transaction processing (OLTP) database, data mining and other transaction-intensive applications.

10.1 Logical Drive Property Settings Required for CacheCade

For a logical drive to be valid to use as a CacheCade drive, the logical drive must be set to Write Through (WT) for write policy and Cached IO (CIO) for IO policy. The following screen shows the logical drive properties menu.

Figure 10.1 Virtual Drive Properties Menu

Set Virtual Drive Properties

LSI

Description : Defines virtual disk operation parameters

Name:

Read Policy:

Write Policy:

IO Policy:

Access Policy:

Disk Cache Policy:

Background Initialization:

10.2 Viewing a Logical Drive with CacheCade

If the logical drive properties are correctly set to use as the CacheCade logical drive (Write Through and Cached IO), then MSM correctly shows the associated logical drive when you click on the CacheCade drive.

Note: If no logical drive exists with properties sufficient for the CacheCade virtual drive, this “Associated Virtual Drives” property does not appear. You can refresh the screen after you update the logical drive properties to make sure the properties are correctly updated.

10.3 WebBIOS Configuration for CacheCade

This section contains the procedures for creating CacheCade virtual drives for the CacheCade advanced software feature.

Note: This procedure does not create a RAID configuration. It creates an CacheCade software virtual drive that functions as a secondary tier of cache.

Using CacheCade software as controller cache allows for very large data sets to be present in cache, delivering performance up to 50 times greater than regular cache in read-intensive applications, such as online transaction processing (OLTP), and file and Web server workloads. The solution accelerates the IO performance of HDD-based drive groups while requiring only a small investment in CacheCade software technology.


To support full-throughput for multiple direct-attached CacheCade software, this feature reduces I/O-processing overhead in the 2108-chip-based MegaRAID controllers. CacheCade offers performance equivalent to flash-based controllers and better performance for RAID 5 and RAID 6 when compared to Fusion I/O.

Follow these steps to create a CacheCade drive group.

1. Click **Configuration Wizard** on the WebBIOS main screen.

The first Configuration Wizard screen appears, as shown in the following figure.


Figure 10.2 WebBIOS Configuration Wizard Screen



2. Select **CacheCade(TM) Configuration** and click **Next**.

The Span Definition screen appears, as shown in the following figure.

Figure 10.3 CacheCade Array Selection Screen



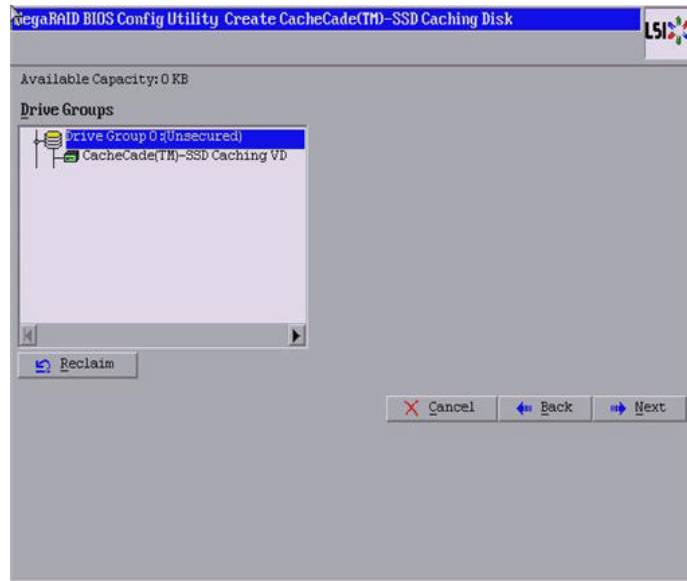
3. Select an array with free space from the drop-down list and click **Select Array**.

The selected array moves to the right frame under the heading **Selected Array**.

4. Click **Next**.

The Create CacheCade Disk screen appears, as shown in the following figure.

Figure 10.4 CacheCade Disk Screen

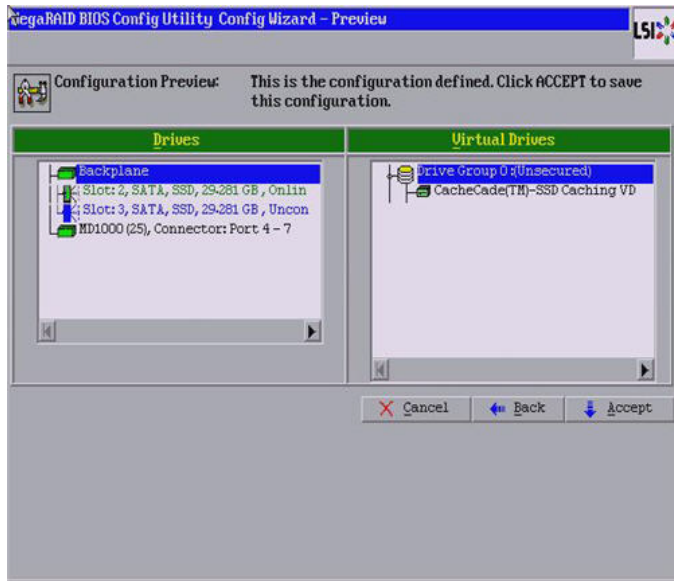


5. Click **Next** to accept the drive group.

If you need to undo the changes, click **Reclaim**.

The Config Wizard Preview screen appears, as shown in the following figure.

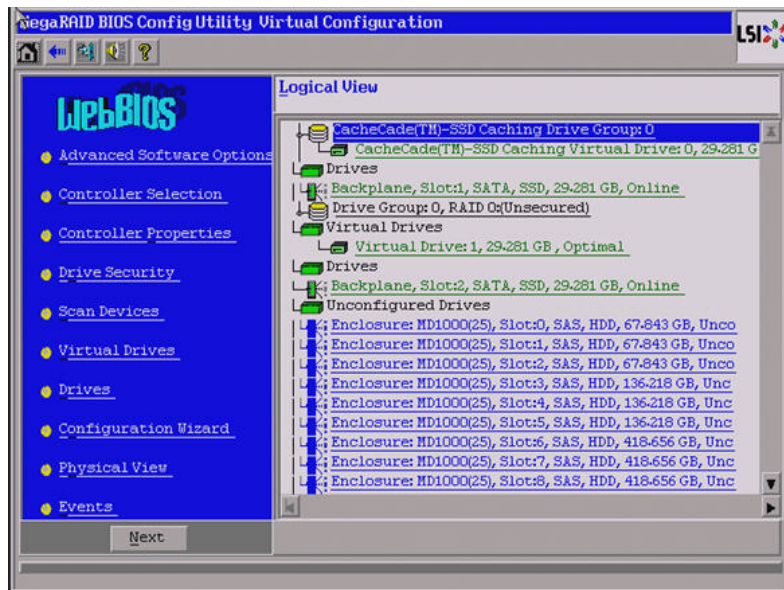
Figure 10.5 CacheCade Configuration Preview



6. Click **Accept** if the configuration is OK. Otherwise, or click **Back** to return to the previous screens and change the configuration.
7. If you accept the configuration, click **Yes** at the prompt to save the configuration.

The WebBIOS main menu screen appears, as shown in the following figure. It shows the CacheCade virtual drive.

Figure 10.6 WebBIOS Main Menu with a CacheCade Virtual Drive



10.4 MegaRAID Storage Manager Configuration for CacheCade

This section contains the procedures for creating CacheCadeRAID virtual drives for the CacheCade advanced software feature.

Note: The MegaRAID firmware has the provision to monitor IO performance; changes have been made to accommodate the CacheCade Pro 2.0 software statistics. The CacheCade Pro 2.0 software metrics are captured for each logical drive that has CacheCade enabled. The CacheCade Pro 2.0 software gathers information about the cache windows allocated for a logical drive, the number of new windows allocated in this metrics collection period, the number of windows that are actively used, and the window hit rates.

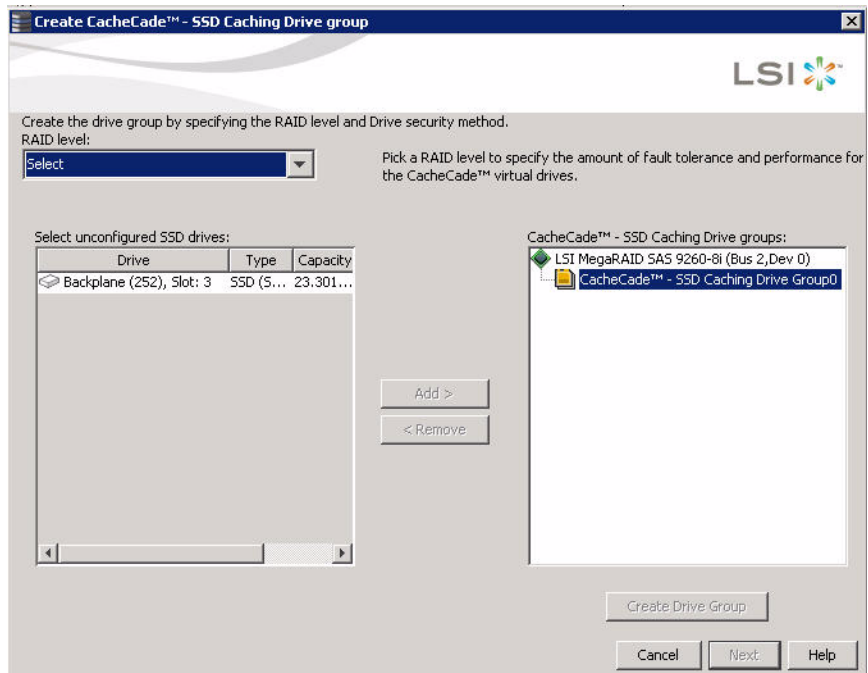
Perform the following steps to use the CacheCade Pro 2.0 software:

1. Perform one of these actions:

- Right-click a controller in the device tree in the left frame of the MegaRAID Storage Manager window and select Create CacheCade SSD Caching.
- Select a controller, and select **Go To >> Controller >> Create CacheCade SSD Caching** in the menu bar.

The CacheCade SSD Caching Wizard appears, as shown in the following figure.

Figure 10.7 CacheCade SSD Caching Wizard - First Screen



1. Select a RAID level for the CacheCade virtual drive in the **RAID level** field.
2. Select an unconfigured SSD drive, for the selected RAID level, from **Select unconfigured SSD Drives** in the left frame.

After you select an unconfigured SSD Drive, the **Add** button is enabled.

3. Click **Add** to add the selected drive to the CacheCade - SSD Caching Drive groups in the right frame.

After you click **Add**, the Create Drive Group button is enabled.

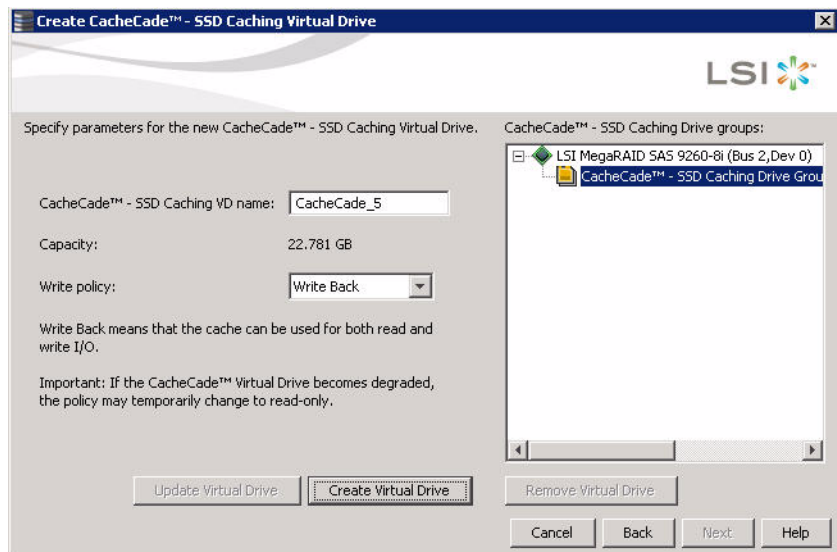
4. Click Create Drive Group.

The newly created drive group appears in CacheCade SSD Caching Drive groups in the right frame.

5. Click Next.

The next wizard screen appears.

Figure 10.8 Parameters for CacheCade SSD Caching Virtual Drive



6. Enter a name for the CacheCade virtual drive in the **CacheCade - SSD caching VD** name field.

7. Select a write policy from the **Write policy** drop-down list.

A description of the selected write policy appears below.

8. Click **Create Virtual Drive**.

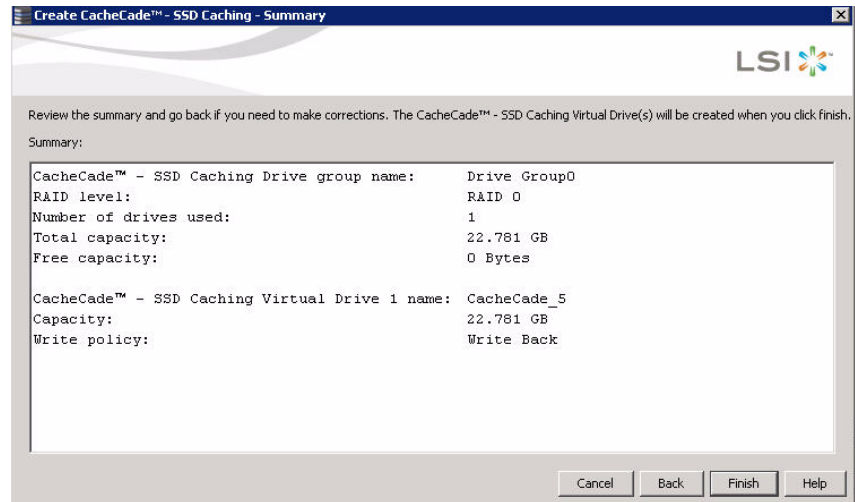
9. The newly created virtual drive appears in the CacheCade SSD Caching Drive groups in the right frame.

The **Remove Virtual Drive** button is enabled. You can select the newly created virtual drive and click **Remove Virtual Drive** to delete the virtual drive.

10. Click **Next**.

The summary screen appears.

Figure 10.9 Create CacheCade - SSD Caching - Summary



This screen displays the drive group name, the RAID level, the number of drives, the total capacity, the free capacity, the CacheCade virtual drive name, the capacity being used, and the write policy.

11. Click Finish.

A confirmation message displays after the CacheCade virtual drive is successfully created. The CacheCade drive icon appears next to the RAID controller in the left frame in the MegaRAID Storage Manager window.

10.5 Modifying the CacheCade Virtual Drive Properties

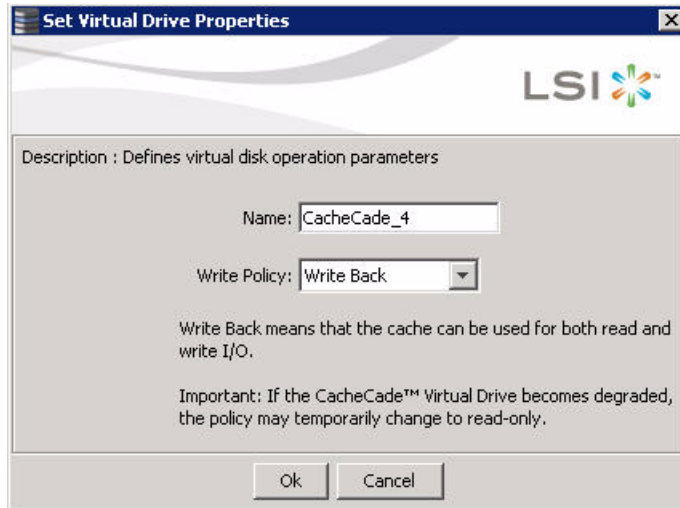
You can modify the name and the write policy of a CacheCade virtual drive any time after a CacheCade virtual drive is created. Perform the following steps to change the virtual drive properties:

1. Perform one of these actions:
 - Right-click a controller in the device tree in the left frame of the MegaRAID Storage Manager window and select **Set Virtual Drive Properties**.

- Select a controller and select **Go To >> Virtual Drive >> Set Virtual Drive Properties**.

The Set Virtual Drive Properties dialog appears, as shown in the following figure.

Figure 10.10 Set Virtual Drive Properties Window



2. Edit the name of a CacheCade virtual drive in the **Name** field.
3. Select a write policy from the **Write Policy** drop down list.
4. Click **OK**.
A confirmation dialog appears with a warning note.
5. Select the **Confirm** check box, and click **OK**.

10.5.1 Enabling SSD Caching on a Virtual Drive

You can enable SSD caching on a virtual drive. When you enable SSD caching on a virtual drive, that virtual drive becomes associated with an existing or with a future CacheCade SSD Caching virtual drive. This option is available only when the virtual drive's caching is currently disabled.


Perform the following steps to enable SSD caching on a virtual drive.

1. Perform one of these actions:

- Right-click a virtual drive in the left frame of the MegaRAID Storage Manager window and select **Enable SSD Caching**.
- Select a virtual drive, and select **Go To >> Virtual Drive >> Enable SSD Caching**.

The **Enable SSD Caching** dialog appears, as shown in the following figure.

Figure 10.11 Enable SSD Caching



2. Click **OK** to enable caching for that virtual drive.

10.5.2 Disabling SSD Caching on a Virtual Drive


You can disable caching on a virtual drive. When you disable SSD caching on a virtual drive, any associations that the selected virtual drive has with a CacheCade SSD Caching virtual drive is removed. This option is only available when the virtual drive's caching is currently enabled.

Perform the following steps to enable SSD Caching on a virtual drive:

1. Perform one of these actions:
 - Right-click a virtual drive in the left frame of the MegaRAID Storage Manager window and select **Disable SSD Caching**.
 - Select a virtual drive, and select **Go To >> Virtual Drive >> Disable SSD Caching**.

The Disable SSD Caching dialog appears, as shown in the following figure.

Figure 10.12 Disable SSD Caching



2. Select the Confirm check box, and click OK to disable caching for that virtual drive.

10.5.3 Enabling or Disabling SSD Caching on Multiple Virtual Drives

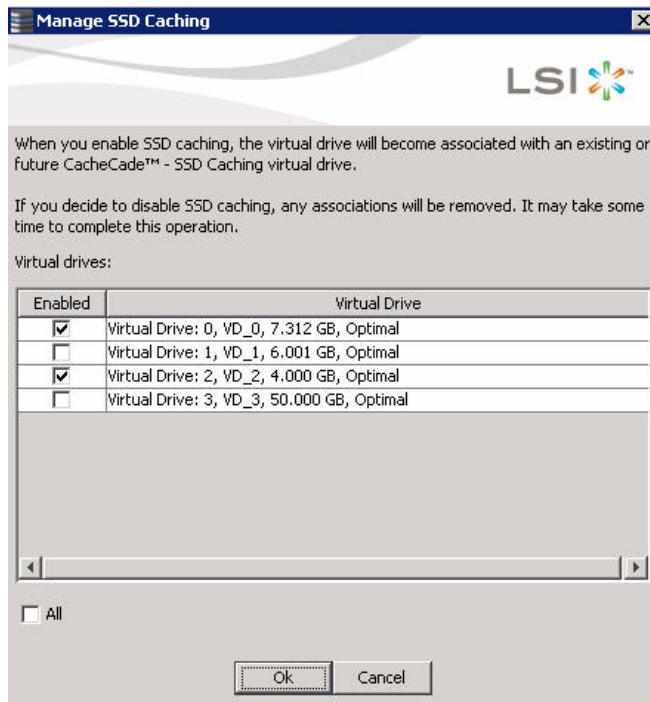
You can enable or disable SSD caching on multiple virtual drives at one go.

Perform the follow steps to enable or disable SSD caching on multiple drives:

1. Perform one of these actions:
 - Right-click a controller in the left frame of the MegaRAID Storage Manager window, and select **Manage SSD Caching**.
 - Select a controller, and select **Go To >> Controller >> Manage SSD Caching**.

The **Manage SSD Caching** dialog appears, as shown in the following figure.

Figure 10.13 Manage SSD Caching



The virtual drives that have SSD caching enabled, have the check boxes next to them selected. The virtual drives that have SSD caching disabled, have deselected check boxes.

2. Select or deselect a check box to change the current setting of a virtual drive.
3. Click **OK**.

If you select the All check box, all the virtual drives are enabled. If you deselect the All check box, all the virtual drives are disabled.

If you disable SSD caching on a virtual drive, the **Disable SSD Caching** dialog appears.

4. Select the **Confirm** check box, and click **OK** to enable or disable SSD caching on the selected virtual drives.

10.5.4 Modifying a CacheCade Drive Group

To modify an existing CacheCade SSD caching drive group, you need to first delete the drive group and then create a new CacheCade drive group.

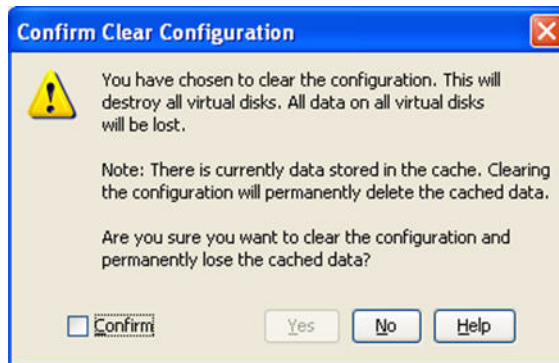
10.5.5 Clearing Configuration on CacheCade Pro 2.0 Virtual Drives

Perform the following steps to clear all existing configurations on a selected controller that has CacheCade Pro 2.0 virtual drives.

1. Perform one of these actions:
 - Right-click a controller in the left frame of the MegaRAID Storage Manager window, and select **Clear Configuration**.
 - Select a controller and select **Go To >> Controller >> Clear Configuration**.

The **Confirm Clear Configuration** dialog appears as shown, in the following figure.

Figure 10.14 Confirm Clear Configuration



2. Select the **Confirm** check box, and click **Yes**.

If the cache becomes inconsistent before the clear configuration operation is performed, the firmware returns an error code. The **Confirm Loss of Cache** dialog appears as a follow-up dialog to the **Confirm Clear Configuration** dialog.
3. Select the **Confirm** check box, and click **Yes**.

10.5.6 Removing Blocked Access

At times, an error may occur in the CacheCade virtual drive and this causes a blocked access to the associated virtual drive.

An icon appears in front of the affected virtual drive, next to the Optimal status.

It is advisable to wait for sometime for the error in the CacheCade virtual drive to get sorted. You can also try to solve the error in the CacheCade virtual drive and bring it back to an optimal status. Once the CacheCade virtual drive is in an optimal status, the blocked virtual drive returns to its former access policy automatically.

If it is not possible to bring the CacheCade virtual drive to its optimal status, follow these steps to remove the blocked access from the virtual drive:

1. Right-click the icon on the virtual drive with the blocked access and select **Remove Blocked Access**.

The **Confirm Remove Blocked Access** dialog appears, as shown in the following figure.

Figure 10.15 Confirm Remove Blocked Access



2. Select the **Confirm** check box, and click **Yes**.

10.5.7 Deleting a Virtual Drive With SSD Caching Enabled

You can delete a virtual drive that has SSD caching enabled on it.

Perform the following steps to delete the virtual drive:

1. Perform one of these actions:
 - Right-click on a **CacheCade** virtual drive, and select **Delete Virtual Drive**.
 - Select a CacheCade virtual drive, and select **Go To >> Virtual Drive >> Delete Virtual Drive**.

The **Confirm Delete Virtual Disk** dialog appears, as shown in the following figure.

Figure 10.16 Confirm Delete Virtual Disk



2. Select the **Confirm** check box, and click **Yes**.

Note: If you select the **Force the delete to complete quickly** check box to delete the virtual drive, the data is not flushed before deleting the virtual drive. In this scenario, if you create this virtual drive after deleting it, no data will be available.

10.6 FastPath Advanced Software

MegaRAID Fast Path is a high-performance IO accelerator for the CacheCade software drive groups connected to a MegaRAID controller card. CacheCade software has a read performance advantage over HDDs and uses less power. This feature dramatically boosts storage subsystem bandwidth and overall transactional application performance when used with a 6Gb/s MegaRAID SATA+SAS controller.

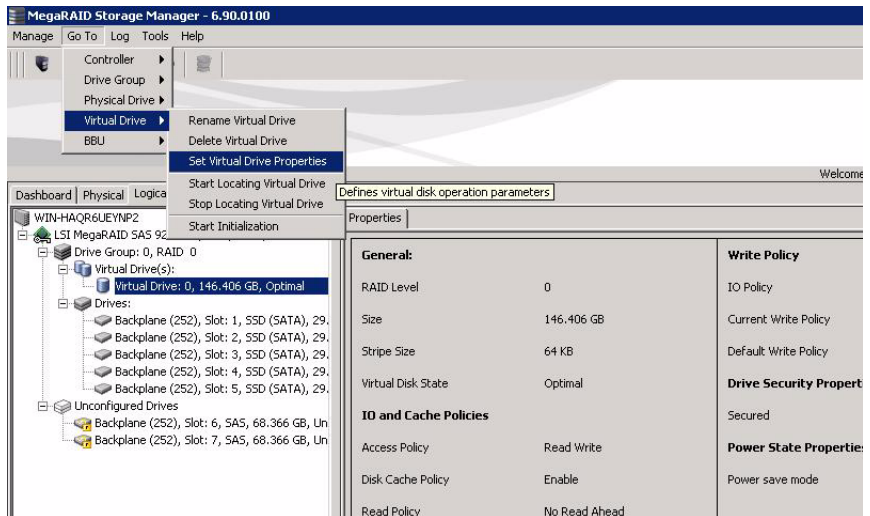
The Fast Path feature supports full optimization of the CacheCade software and hard disk drive (HDD) virtual disk groups to deliver an improvement in read and write IOPS that is three times greater than MegaRAID controllers not using FastPath technology. Also, the Fast Path advanced software is faster and more cost-effective than current flash-based adapter card solutions.

10.6.1 Setting Fast Path Options

Perform the following steps to use the FastPath advanced software:

1. Select the **Logical** tab on the MegaRAID Storage Manager main menu screen for the Logical view.
2. Select a virtual drive icon in the left frame.
3. Select **Virtual Drive >> Set Virtual Drive Properties** on the menu bar, as shown in the following figure.

Figure 10.17 Set Virtual Drive Properties Menu



The Set Virtual Drive Properties dialog appears and shows the default settings for the Fast Path advanced software:

- Write Policy: Write Thru
- IO Policy: Direct IO
- Read Policy: No Read Ahead
- Disk Cache Policy: Disabled
- Strip Size: 64KB

4. Click **OK**.

The confirmation dialog appears.

5. Select the **Confirm** check box, and click **Yes** to confirm that you want to set the virtual drive properties.

Appendix A

Events and Messages

This appendix lists the MegaRAID Storage Manager events that might appear in the event log.

MegaRAID Storage Manager software monitors the activity and performance of all controllers in the workstation and the devices attached to them. When an event occurs, such as the start of an initialization, an event message appears in the log at the bottom of the MegaRAID Storage Manager window.

Each message that appears in the event log has an error level that indicates the severity of the event, as shown in [Table A.1](#).

Table A.1 Event Error Levels

Error Level	Meaning
Information	Informational message. No user action is necessary.
Warning	Some component may be close to a failure point.
Critical	A component has failed, but the system has not lost data.
Fatal	A component has failed, and data loss has occurred or will occur.

[Table A.2](#) lists all of the MegaRAID Storage Manager event messages. The event message descriptions include placeholders for specific values that are determined when the event is generated. For example, in message No. 1 in the Event Messages table, “%s” is replaced by the

firmware version, which is read from the firmware when the event is generated.

Table A.2 Event Messages

Number	Type	Event Text
0x0000	Information	MegaRAID firmware initialization started (PCI ID %04x/%04x/%04x/%04x)
0x0001	Information	MegaRAID firmware version %s
0x0002	Fatal	Unable to recover cache data from TBBU
0x0003	Information	Cache data recovered from TBBU successfully
0x0004	Information	Configuration cleared
0x0005	Warning	Cluster down; communication with peer lost
0x0006	Information	Virtual drive %s ownership changed from %02x to %02x
0x0007	Information	Alarm disabled by user
0x0008	Information	Alarm enabled by user
0x0009	Information	Background initialization rate changed to %d%%
0x000a	Fatal	Controller cache discarded due to memory/battery problems
0x000b	Fatal	Unable to recover cache data due to configuration mismatch
0x000c	Information	Cache data recovered successfully
0x000d	Fatal	Controller cache discarded due to firmware version incompatibility
0x000e	Information	Consistency Check rate changed to %d%%
0x000f	Fatal	Fatal firmware error: %s
0x0010	Information	Factory defaults restored
0x0011	Information	Flash downloaded image corrupt
0x0012	Critical	Flash erase error
0x0013	Critical	Flash timeout during erase
0x0014	Critical	Flash error
0x0015	Information	Flashing image: %s
0x0016	Information	Flash of new firmware image(s) complete
0x0017	Critical	Flash programming error
0x0018	Critical	Flash timeout during programming
0x0019	Critical	Flash chip type unknown
(Sheet 1 of 17)		

Table A.2 Event Messages (Cont.)

Number	Type	Event Text
0x001a	Critical	Flash command set unknown
0x001b	Critical	Flash verify failure
0x001c	Information	Flush rate changed to %d seconds
0x001d	Information	Hibernate command received from host
0x001e	Information	Event log cleared
0x001f	Information	Event log wrapped
0x0020	Fatal	Multi-bit ECC error: ECAR=%x, ELOG=%x, (%s)
0x0021	Warning	Single-bit ECC error: ECAR=%x, ELOG=%x, (%s)
0x0022	Fatal	Not enough controller memory
0x0023	Information	Patrol Read complete
0x0024	Information	Patrol Read paused
0x0025	Information	Patrol Read Rate changed to %d%%
0x0026	Information	Patrol Read resumed
0x0027	Information	Patrol Read started
0x0028	Information	Rebuild rate changed to %d%%
0x0029	Information	Drive group modification rate changed to %d%%
0x002a	Information	Shutdown command received from host
0x002b	Information	Test event: %s
0x002c	Information	Time established as %s; (%d seconds since power on)
0x002d	Information	User entered firmware debugger
0x002e	Warning	Background Initialization aborted on %s
0x002f	Warning	Background Initialization corrected medium error (%s at %lx
0x0030	Information	Background Initialization completed on %s
0x0031	Fatal	Background Initialization completed with uncorrectable errors on %s
0x0032	Fatal	Background Initialization detected uncorrectable double medium errors (%s at %lx on %s)
0x0033	Critical	Background Initialization failed on %s
0x0034	Progress	Background Initialization progress on %s is %s
0x0035	Information	Background Initialization started on %s
0x0036	Information	Policy change on %s from %s to %s
(Sheet 2 of 17)		

Table A.2 Event Messages (Cont.)

Number	Type	Event Text
0x0038	Warning	Consistency Check aborted on %s
0x0039	Warning	Consistency Check corrected medium error (%s at %lx
0x003a	Information	Consistency Check done on %s
0x003b	Information	Consistency Check done with corrections on %s
0x003c	Fatal	Consistency Check detected uncorrectable double medium errors (%s at %lx on %s)
0x003d	Critical	Consistency Check failed on %s
0x003e	Fatal	Consistency Check completed with uncorrectable data on %s
0x003f	Warning	Consistency Check found inconsistent parity on %s at strip %lx
0x0040	Warning	Consistency Check inconsistency logging disabled on %s (too many inconsistencies)
0x0041	Progress	Consistency Check progress on %s is %s
0x0042	Information	Consistency Check started on %s
0x0043	Warning	Initialization aborted on %s
0x0044	Critical	Initialization failed on %s
0x0045	Progress	Initialization progress on %s is %s
0x0046	Information	Fast initialization started on %s
0x0047	Information	Full initialization started on %s
0x0048	Information	Initialization complete on %s
0x0049	Information	LD Properties updated to %s (from %s)
0x004a	Information	Reconstruction complete on %s
0x004b	Fatal	Reconstruction of %s stopped due to unrecoverable errors
0x004c	Fatal	Reconstruct detected uncorrectable double medium errors (%s at %lx on %s at %lx)
0x004d	Progress	Reconstruction progress on %s is %s
0x004e	Information	Reconstruction resumed on %s
0x004f	Fatal	Reconstruction resume of %s failed due to configuration mismatch
0x0050	Information	Reconstruction started on %s
0x0051	Information	State change on %s from %s to %s
(Sheet 3 of 17)		

Table A.2 Event Messages (Cont.)

Number	Type	Event Text
0x0052	Information	Drive Clear aborted on %s
0x0053	Critical	Drive Clear failed on %s (Error %02x)
0x0054	Progress	Drive Clear progress on %s is %s
0x0055	Information	Drive Clear started on %s
0x0056	Information	Drive Clear completed on %s
0x0057	Warning	Error on %s (Error %02x)
0x0058	Information	Format complete on %s
0x0059	Information	Format started on %s
0x005a	Critical	Hot Spare SMART polling failed on %s (Error %02x)
0x005b	Information	Drive inserted: %s
0x005c	Warning	Drive %s is not supported
0x005d	Warning	Patrol Read corrected medium error on %s at %lx
0x005e	Progress	Patrol Read progress on %s is %s
0x005f	Fatal	Patrol Read found an uncorrectable medium error on %s at %lx
0x0060	Critical	Predictive failure: CDB: %s
0x0061	Fatal	Patrol Read puncturing bad block on %s at %lx
0x0062	Information	Rebuild aborted by user on %s
0x0063	Information	Rebuild complete on %s
0x0064	Information	Rebuild complete on %s
0x0065	Critical	Rebuild failed on %s due to source drive error
0x0066	Critical	Rebuild failed on %s due to target drive error
0x0067	Progress	Rebuild progress on %s is %s
0x0068	Information	Rebuild resumed on %s
0x0069	Information	Rebuild started on %s
0x006a	Information	Rebuild automatically started on %s
0x006b	Critical	Rebuild stopped on %s due to loss of cluster ownership
0x006c	Fatal	Reassign write operation failed on %s at %lx
0x006d	Fatal	Unrecoverable medium error during rebuild on %s at %lx
0x006e	Information	Corrected medium error during recovery on %s at %lx
(Sheet 4 of 17)		

Table A.2 Event Messages (Cont.)

Number	Type	Event Text
0x006f	Fatal	Unrecoverable medium error during recovery on %s at %lx
0x0070	Information	Drive removed: %s
0x0071	Information	Unexpected sense: %s, CDB%s, Sense: %s
0x0072	Information	State change on %s from %s to %s
0x0073	Information	State change by user on %s from %s to %s
0x0074	Warning	Redundant path to %s broken
0x0075	Information	Redundant path to %s restored
0x0076	Information	Dedicated Hot Spare Drive %s no longer useful due to deleted drive group
0x0077	Critical	SAS topology error: Loop detected
0x0078	Critical	SAS topology error: Unaddressable device
0x0079	Critical	SAS topology error: Multiple ports to the same SAS address
0x007a	Critical	SAS topology error: Expander error
0x007b	Critical	SAS topology error: SMP timeout
0x007c	Critical	SAS topology error: Out of route entries
0x007d	Critical	SAS topology error: Index not found
0x007e	Critical	SAS topology error: SMP function failed
0x007f	Critical	SAS topology error: SMP CRC error
0x0080	Critical	SAS topology error: Multiple subtractive
0x0081	Critical	SAS topology error: Table to table
0x0082	Critical	SAS topology error: Multiple paths
0x0083	Fatal	Unable to access device %s
0x0084	Information	Dedicated Hot Spare created on %s (%s)
0x0085	Information	Dedicated Hot Spare %s disabled
0x0086	Critical	Dedicated Hot Spare %s no longer useful for all drive groups
0x0087	Information	Global Hot Spare created on %s (%s)
0x0088	Information	Global Hot Spare %s disabled
0x0089	Critical	Global Hot Spare does not cover all drive groups
0x008a	Information	Created %s}
(Sheet 5 of 17)		

Table A.2 Event Messages (Cont.)

Number	Type	Event Text
0x008b	Information	Deleted %s}
0x008c	Information	Marking LD %s inconsistent due to active writes at shutdown
0x008d	Information	Battery Present
0x008e	Warning	Battery Not Present
0x008f	Information	New Battery Detected
0x0090	Information	Battery has been replaced
0x0091	Critical	Battery temperature is high
0x0092	Warning	Battery voltage low
0x0093	Information	Battery started charging
0x0094	Information	Battery is discharging
0x0095	Information	Battery temperature is normal
0x0096	Fatal	Battery has failed and cannot support data retention. Please replace the battery.
0x0097	Information	Battery relearn started
0x0098	Information	Battery relearn in progress
0x0099	Information	Battery relearn completed
0x009a	Critical	Battery relearn timed out
0x009b	Information	Battery relearn pending: Battery is under charge
0x009c	Information	Battery relearn postponed
0x009d	Information	Battery relearn will start in 4 days
0x009e	Information	Battery relearn will start in 2 day
0x009f	Information	Battery relearn will start in 1 day
0x00a0	Information	Battery relearn will start in 5 hours
0x00a1	Information	Battery removed
0x00a2	Information	Current capacity of the battery is below threshold
0x00a3	Information	Current capacity of the battery is above threshold
0x00a4	Information	Enclosure (SES) discovered on %s
0x00a5	Information	Enclosure (SAFTE) discovered on %s
0x00a6	Critical	Enclosure %s communication lost
0x00a7	Information	Enclosure %s communication restored
0x00a8	Critical	Enclosure %s fan %d failed
(Sheet 6 of 17)		

Table A.2 Event Messages (Cont.)

Number	Type	Event Text
0x00a9	Information	Enclosure %s fan %d inserted
0x00aa	Critical	Enclosure %s fan %d removed
0x00ab	Critical	Enclosure %s power supply %d failed
0x00ac	Information	Enclosure %s power supply %d inserted
0x00ad	Critical	Enclosure %s power supply %d removed
0x00ae	Critical	Enclosure %s SIM %d failed
0x00af	Information	Enclosure %s SIM %d inserted
0x00b0	Critical	Enclosure %s SIM %d removed
0x00b1	Warning	Enclosure %s temperature sensor %d below warning threshold
0x00b2	Critical	Enclosure %s temperature sensor %d below error threshold
0x00b3	Warning	Enclosure %s temperature sensor %d above warning threshold
0x00b4	Critical	Enclosure %s temperature sensor %d above error threshold
0x00b5	Critical	Enclosure %s shutdown
0x00b6	Warning	Enclosure %s not supported; too many enclosures connected to port
0x00b7	Critical	Enclosure %s firmware mismatch
0x00b8	Warning	Enclosure %s sensor %d bad
0x00b9	Critical	Enclosure %s phy %d bad
0x00ba	Critical	Enclosure %s is unstable
0x00bb	Critical	Enclosure %s hardware error
0x00bc	Critical	Enclosure %s not responding
0x00bd	Information	SAS/SATA mixing not supported in enclosure; Drive %s disabled
0x00be	Information	Enclosure (SES) hotplug on %s was detected, but is not supported
0x00bf	Information	Clustering enabled
0x00c0	Information	Clustering disabled
0x00c1	Information	Drive too small to be used for auto-rebuild on %s
0x00c2	Information	BBU enabled; changing WT virtual drives to WB
(Sheet 7 of 17)		

Table A.2 Event Messages (Cont.)

Number	Type	Event Text
0x00c3	Warning	BBU disabled; changing WB virtual drives to WT
0x00c4	Warning	Bad block table on drive %s is 80% full
0x00c5	Fatal	Bad block table on drive %s is full; unable to log block %lx
0x00c6	Information	Consistency Check Aborted due to ownership loss on %s
0x00c7	Information	Background Initialization (BGI) Aborted Due to Ownership Loss on %s
0x00c8	Critical	Battery/charger problems detected; SOH Bad
0x00c9	Warning	Single-bit ECC error: ECAR=%x, ELOG=%x, (%s); warning threshold exceeded
0x00ca	Critical	Single-bit ECC error: ECAR=%x, ELOG=%x, (%s); critical threshold exceeded
0x00cb	Critical	Single-bit ECC error: ECAR=%x, ELOG=%x, (%s); further reporting disabled
0x00cc	Critical	Enclosure %s Power supply %d switched off
0x00cd	Information	Enclosure %s Power supply %d switched on
0x00ce	Critical	Enclosure %s Power supply %d cable removed
0x00cf	Information	Enclosure %s Power supply %d cable inserted
0x00d0	Information	Enclosure %s Fan %d returned to normal
0x00d1	Information	BBU Retention test was initiated on previous boot
0x00d2	Information	BBU Retention test passed
0x00d3	Critical	BBU Retention test failed!
0x00d4	Information	NVRAM Retention test was initiated on previous boot
0x00d5	Information	NVRAM Retention test passed
0x00d6	Critical	NVRAM Retention test failed!
0x00d7	Information	%s test completed %d passes successfully
0x00d8	Critical	%s test FAILED on %d pass. Fail data: errorOffset=%x goodData=%x badData=%x
0x00d9	Information	Self check diagnostics completed
0x00da	Information	Foreign Configuration detected
0x00db	Information	Foreign Configuration imported
0x00dc	Information	Foreign Configuration cleared
(Sheet 8 of 17)		

Table A.2 Event Messages (Cont.)

Number	Type	Event Text
0x00dd	Warning	NVRAM is corrupt; reinitializing
0x00de	Warning	NVRAM mismatch occurred
0x00df	Warning	SAS wide port %d lost link on PHY %d
0x00e0	Information	SAS wide port %d restored link on PHY %d
0x00e1	Warning	SAS port %d, PHY %d has exceeded the allowed error rate
0x00e2	Warning	Bad block reassigned on %s at %lx to %lx
0x00e3	Information	Controller Hot Plug detected
0x00e4	Warning	Enclosure %s temperature sensor %d differential detected
0x00e5	Information	Drive test cannot start. No qualifying drives found
0x00e6	Information	Time duration provided by host is not sufficient for self check
0x00e7	Information	Marked Missing for %s on drive group %d row %d
0x00e8	Information	Replaced Missing as %s on drive group %d row %d
0x00e9	Information	Enclosure %s Temperature %d returned to normal
0x00ea	Information	Enclosure %s Firmware download in progress
0x00eb	Warning	Enclosure %s Firmware download failed
0x00ec	Warning	%s is not a certified drive
0x00ed	Information	Dirty cache data discarded by user
0x00ee	Information	Drives missing from configuration at boot
0x00ef	Information	Virtual drives (VDs) missing drives and will go offline at boot: %s
0x00f0	Information	VDs missing at boot: %s
0x00f1	Information	Previous configuration completely missing at boot
0x00f2	Information	Battery charge complete
0x00f3	Information	Enclosure %s fan %d speed changed
0x00f4	Information	Dedicated spare %s imported as global due to missing arrays
0x00f5	Information	%s rebuild not possible as SAS/SATA is not supported in an array
0x00f6	Information	SEP %s has been rebooted as a part of enclosure firmware download. SEP will be unavailable until this process completes.
(Sheet 9 of 17)		

Table A.2 Event Messages (Cont.)

Number	Type	Event Text
0x00f7	Information	Inserted PD: %s Info: %s
0x00f8	Information	Removed PD: %s Info: %s
0x00f9	Information	VD %s is now OPTIMAL
0x00fa	Warning	VD %s is now PARTIALLY DEGRADED
0x00fb	Critical	VD %s is now DEGRADED
0x00fc	Fatal	VD %s is now OFFLINE
0x00fd	Warning	Battery requires reconditioning; please initiate a LEARN cycle
0x00fe	Warning	VD %s disabled because RAID-5 is not supported by this RAID key
0x00ff	Warning	VD %s disabled because RAID-6 is not supported by this controller
0x0100	Warning	VD %s disabled because SAS drives are not supported by this RAID key
0x0101	Warning	PD missing: %s
0x0102	Warning	Puncturing of LBAs enabled
0x0103	Warning	Puncturing of LBAs disabled
0x0104	Critical	Enclosure %s EMM %d not installed
0x0105	Information	Package version %s
0x0106	Warning	Global affinity Hot Spare %s commissioned in a different enclosure
0x0107	Warning	Foreign configuration table overflow
0x0108	Warning	Partial foreign configuration imported, PDs not imported:%s
0x0109	Information	Connector %s is active
0x010a	Information	Board Revision %s
0x010b	Warning	Command timeout on PD %s, CDB:%s
0x010c	Warning	PD %s reset (Type %02x)
0x010d	Warning	VD bad block table on %s is 80% full
0x010e	Fatal	VD bad block table on %s is full; unable to log block %lx (on %s at %lx)
0x010f	Fatal	Uncorrectable medium error logged for %s at %lx (on %s at %lx)
0x0110	Information	VD medium error corrected on %s at %lx
(Sheet 10 of 17)		

Table A.2 Event Messages (Cont.)

Number	Type	Event Text
0x0111	Warning	Bad block table on PD %s is 100% full
0x0112	Warning	VD bad block table on PD %s is 100% full
0x0113	Fatal	Controller needs replacement, IOP is faulty
0x0114	Information	CopyBack started on PD %s from PD %s
0x0115	Information	CopyBack aborted on PD %s and src is PD %s
0x0116	Information	CopyBack complete on PD %s from PD %s
0x0117	Progress	CopyBack progress on PD %s is %s
0x0118	Information	CopyBack resumed on PD %s from %s
0x0119	Information	CopyBack automatically started on PD %s from %s
0x011a	Critical	CopyBack failed on PD %s due to source %s error
0x011b	Warning	Early Power off warning was unsuccessful
0x011c	Information	BBU FRU is %s
0x011d	Information	%s FRU is %s
0x011e	Information	Controller hardware revision ID %s
0x011f	Warning	Foreign import shall result in a backward incompatible upgrade of configuration metadata
0x0120	Information	Redundant path restored for PD %s
0x0121	Warning	Redundant path broken for PD %s
0x0122	Information	Redundant enclosure EMM %s inserted for EMM %s
0x0123	Information	Redundant enclosure EMM %s removed for EMM %s
0x0124	Warning	Patrol Read can't be started, as PDs are either not ONLINE, or are in a VD with an active process, or are in an excluded VD
0x0125	Information	Copyback aborted by user on PD %s and src is PD %s
0x0126	Critical	Copyback aborted on hot spare %s from %s, as hot spare needed for rebuild
0x0127	Warning	Copyback aborted on PD %s from PD %s, as rebuild required in the array
0x0128	Fatal	Controller cache discarded for missing or offline VD %s When a VD with cached data goes offline or missing during runtime, the cache for the VD is discarded. Because the VD is offline, the cache cannot be saved.
0x0129	Information	Copyback cannot be started as PD %s is too small for src PD %s
(Sheet 11 of 17)		

Table A.2 Event Messages (Cont.)

Number	Type	Event Text
0x012a	Information	Copyback cannot be started on PD %s from PD %s, as SAS/SATA is not supported in an array
0x012b	Information	Microcode update started on PD %s
0x012c	Information	Microcode update completed on PD %s
0x012d	Warning	Microcode update timeout on PD %s
0x012e	Warning	Microcode update failed on PD %s
0x012f	Information	Controller properties changed
0x0130	Information	Patrol Read properties changed
0x0131	Information	CC Schedule properties changed
0x0132	Information	Battery properties changed
0x0133	Warning	Periodic Battery Relearn is pending. Please initiate manual learn cycle as Automatic learn is not enabled
0x0134	Information	Drive security key created
0x0135	Information	Drive security key backed up
0x0136	Information	Drive security key from escrow, verified
0x0137	Information	Drive security key changed
0x0138	Warning	Drive security key, re-key operation failed
0x0139	Warning	Drive security key is invalid
0x013a	Information	Drive security key destroyed
0x013b	Warning	Drive security key from escrow is invalid
0x013c	Information	VD %s is now secured
0x013d	Warning	VD %s is partially secured
0x013e	Information	PD %s security activated
0x013f	Information	PD %s security disabled
0x0140	Information	PD %s is reprovisioned
0x0141	Information	PD %s security key changed
0x0142	Fatal	Security subsystem problems detected for PD %s
0x0143	Fatal	Controller cache pinned for missing or offline VD %s
0x0144	Fatal	Controller cache pinned for missing or offline VDs: %s
0x0145	Information	Controller cache discarded by user for VDs: %s
0x0146	Information	Controller cache destaged for VD %s
0x0147	Warning	Consistency Check started on an inconsistent VD %s
(Sheet 12 of 17)		

Table A.2 Event Messages (Cont.)

Number	Type	Event Text
0x0148	Warning	Drive security key failure, cannot access secured configuration
0x0149	Warning	Drive security pass phrase from user is invalid
0x014a	Warning	Detected error with the remote battery connector cable
0x014b	Information	Power state change on PD %s from %s to %s
0x014c	Information	Enclosure %s element (SES code 0x%x) status changed
0x014d	Information	PD %s rebuild not possible as HDD/SSD mix is not supported in a drive group
0x014e	Information	Copyback cannot be started on PD %s from %s, as HDD/SSD mix is not supported in a drive group
0x014f	Information	VD bad block table on %s is cleared
0x0150	Caution	SAS topology error: 0x%lx
0x0151	Information	VD cluster of medium errors corrected for %s at %lx (on %s at %lx)
0x0152	Information	Controller requests a host bus rescan
0x0153	Information	Controller repurposed and factory defaults restored
0x0154	Information	Drive security key binding updated
0x0155	Information	Drive security is in EKM mode
0x0156	Warning	Drive security failed to communicate with EKMS
0x0157	Information	%s needs key to be %s %s
0x0158	Warning	%s secure failed
0x0159	Critical	Controller encountered a fatal error and was reset
0x015a	Information	Snapshots enabled on %s (Repository %s)
0x015b	Information	Snapshots disabled on %s (Repository %s) by the user
0x015c	Critical	Snapshots disabled on %s (Repository %s), due to a fatal error
0x015d	Information	Snapshot created on %s at %s
0x015e	Information	Snapshot deleted on %s at %s
0x015f	Information	View created at %s to a snapshot at %s for %s
0x0160	Information	View at %s is deleted, to snapshot at %s for %s
0x0161	Information	Snapshot rollback started on %s from snapshot at %s
(Sheet 13 of 17)		

Table A.2 Event Messages (Cont.)

Number	Type	Event Text
0x0162	Fatal	Snapshot rollback on %s internally aborted for snapshot at %s
0x0163	Information	Snapshot rollback on %s completed for snapshot at %s
0x0164	Information	Snapshot rollback progress for snapshot at %s, on %s is %s
0x0165	Warning	Snapshot space for %s in snapshot repository %s, is 80%% full
0x0166	Critical	Snapshot space for %s in snapshot repository %s, is full
0x0167	Warning	View at %s to snapshot at %s, is 80%% full on snapshot repository %s
0x0168	Critical	View at %s to snapshot at %s, is full on snapshot repository %s
0x0169	Critical	Snapshot repository lost for %s
0x016a	Warning	Snapshot repository restored for %s
0x016b	Critical	Snapshot encountered an unexpected internal error: 0x%lx
0x016c	Information	Auto Snapshot enabled on %s (snapshot repository %s)
0x016d	Information	Auto Snapshot disabled on %s (snapshot repository %s)
0x016e	Critical	Configuration command could not be committed to disk, please retry
0x016f	Information	COD on %s updated as it was stale
0x0170	Warning	Power state change failed on %s (from %s to %s)
0x0171	Warning	%s is not available
0x0172	Information	%s is available
0x0173	Information	%s is used for CacheCade with capacity 0x%lx logical blocks
0x0174	Information	%s is using CacheCade %s
0x0175	Information	%s is no longer using CacheCade %s
0x0176	Critical	Snapshot deleted due to resource constraints for %s in snapshot repository %s
0x0177	Warning	Auto Snapshot failed for %s in snapshot repository %s
0x0178	Warning	Controller reset on-board expander
(Sheet 14 of 17)		

Table A.2 Event Messages (Cont.)

Number	Type	Event Text
0x0179	Warning	CacheCade (%s) capacity changed and is now 0x%lx logical blocks
0x017a	Warning	Battery cannot initiate transparent learn cycles
0x017b	Information	Premium feature %s key was applied for - %s
0x017c	Information	Snapshot schedule properties changed on %s
0x017d	Information	Snapshot scheduled action is due on %s
0x017e	Information	Performance Metrics: collection command 0x%lx
0x017f	Information	Premium feature %s key was transferred - %s
0x0180	Information	Premium feature serial number %s
0x0181	Warning	Premium feature serial number mismatched. Key-vault serial num - %s
0x0182	Warning	Battery cannot support data retention for more than %d hours. Please replace the battery
0x0183	Information	%s power policy changed to %s (from %s)
0x0184	Warning	%s cannot transition to max power savings
0x0185	Information	Host driver is loaded and operational
0x0186	Information	%s mirror broken
0x0187	Information	%s mirror joined
0x0188	Warning	%s link %d failure in wide port
0x0189	Information	%s link %d restored in wide port
0x018a	Information	Memory module FRU is %s
0x018b	Warning	Cache-vault power pack is sub-optimal. Please replace the pack
0x018c	Warning	Foreign configuration auto-import did not import any drives
0x018d	Warning	Cache-vault microcode update required
0x018e	Warning	CacheCade (%s) capacity exceeds maximum allowed size, extra capacity is not used
0x018f	Warning	LD (%s) protection information lost
0x0190	Information	Diagnostics passed for %s
0x0191	Critical	Diagnostics failed for %s
0x0192	Information	Server Power capability Diagnostic Test Started
0x0193	Information	Drive Cache settings enabled during rebuild for %s
(Sheet 15 of 17)		

Table A.2 Event Messages (Cont.)

Number	Type	Event Text
0x0194	Information	Drive Cache settings restored after rebuild for %s
0x0195	Information	Drive %s commissioned as Emergency spare
0x0196	Warning	Reminder: Potential non-optimal configuration due to drive %s commissioned as emergency spare
0x0197	Information	Consistency Check suspended on %s
0x0198	Information	Consistency Check resumed on %s
0x0199	Information	Background Initialization suspended on %s
0x019a	Information	Background Initialization resumed on %
0x019b	Information	Reconstruction suspended on %s
0x019c	Information	Rebuild suspended on %
0x019d	Information	Copyback suspended on %s
0x019e	Information	Reminder: Consistency Check suspended on %
0x019f	Information	Reminder: Background Initialization suspended on %s
0x01a0	Information	Reminder: Reconstruction suspended on %s
0x01a1	Information	Reminder: Rebuild suspended on %s
0x01a2	Information	Reminder: Copyback suspended on %s
0x01a3	Information	Reminder: Patrol Read suspended
0x01a4	Information	Erase aborted on %s
0x01a5	Critical	Erase failed on %s (Error %02x)
0x01a6	Progress	Erase progress on %s is %s
0x01a7	Information	Erase started on %s
0x01a8	Information	Erase completed on %s
0x01a9	Information	Erase aborted on %s
0x01aa	Critical	Erase failed on %s
0x01ab	Progress	Erase progress on %s is %s
0x01ac	Information	Erase started on %s
0x01ad	Information	Erase complete on %s
0x01ae	Warning	Potential leakage during erase on %s
0x01af	Warning	Battery charging was suspended due to high battery temperature
0x01b0	Information	NVCache firmware update was successful
0x01b1	Warning	NVCache firmware update failed
(Sheet 16 of 17)		

Table A.2 Event Messages (Cont.)

Number	Type	Event Text
0x01b2	Fatal	%s access blocked as cached data in CacheCade is unavailable
0x01b3	Information	CacheCade disassociate started on %s
0x01b4	Information	CacheCade disassociate completed on %s
0x01b5	Critical	CacheCade disassociate failed on %s
0x01b6	Progress	CacheCade disassociate progress on %s is %s
0x01b7	Information	CacheCade disassociate aborted by user on %s
0x01b8	Information	Link speed changed on SAS port %d and PHY %d
0x01b9	Warning	Advanced Software Options was deactivated for - %s
0x01ba	Information	%s is now accessible
0x01bb	Information	%s is using CacheCade
0x01bc	Information	%s is no longer using CacheCade
0x01bd	Warning	Patrol Read aborted on %s
(Sheet 17 of 17)		

Appendix B

SNMP Extension Agent

Trap Definitions

This appendix describes the trap definitions that an SNMP-based management application uses to send information about system events.

An SNMP-based management application (also known as an SNMP manager) can monitor and manage devices through SNMP extension agents. The MegaRAID SNMP subagent reports the information about the RAID controller, virtual drives, physical devices, enclosures, and other items per SNMP request. The MegaRAID SNMP agent supports a set of traps for any state change in the controller and its attached devices. The traps that depend on the notification from the firmware.

There are no traps specific to Logical Unit Numbers (LUNs) as the MegaRAID firmware does not allow online removal or addition of LUNs. Hence, there is no notification from the firmware.

This guide contains the following sections:

- [Section B.1, “SNMP Traps for RAID Controllers”](#)
- [Section B.2, “SNMP Traps for Virtual Drives”](#)
- [Section B.3, “SNMP Traps for Physical Drives”](#)

B.1 SNMP Traps for RAID Controllers

The following table lists the supported SNMP traps and their definitions for events related to the RAID controllers.

Table B.1 Supported SNMP Traps and Definitions for RAID Controllers

Trap Number	Trap Definition
8001	Alarm has been enabled by user for Adapter %d
8002	Background Initialization Rate changed to %d for Adapter %d
8003	Controller %d cache discarded due to memory/battery problems
8004	Unable to recover Cache Data due to configuration mismatch for Adapter %d
8005	Cache Data Recovered for Adapter %d
8006	Controller Cache Discarded due to Firmware version incompatibility for Adapter %d
8007	Consistency Check Rate changed to %d for Adapter %d
8008	Flash download image corrupted for Adapter %d
8009	Flash erase error for Adapter %d
8010	Flash Timeout during Erase for Adapter %d
8011	General Flash Error occurred for Adapter %d
8012	Flashing Image %s for Adapter %d
8013	Flash of Firmware Image complete for Adapter %d
8014	Flash programming error for Adapter %d
8015	Flash timeout during programming for Adapter %d
8016	Event log cleared for Adapter %d
8017	Event log wrapped for Adapter %d
8018	Patrol Read complete for Adapter %d
8019	Patrol Read paused for Adapter %d

Trap Number	Trap Definition
8020	Patrol Read resumed for Adapter %d
8021	Patrol Read started for Adapter %d
8022	Shutdown command received for Adapter %d
8023	Hibernate command received for Adapter %d
8024	Fatal error received for Adapter %d
8025	Rebuild Rate changed to %d for Adapter %d
8026	Patrol Read Rate changed to %d for Adapter %d
8027	Alarm disabled by user for Adapter %d
8028	Configuration cleared for Adapter %d
8029	Reconstruction Rate changed to %d for Adapter %d
8030	Factory defaults restored for Adapter %d
8031	Battery Present for Adapter %d
8032	Battery not Present for Adapter %d
8033	New Battery Detected for Adapter %d
8034	Battery has been replaced for Adapter %d
8035	Battery Temperature is high for Adapter %d
8036	Battery voltage low for Adapter %d
8036	Battery started charging for Adapter %d
8037	Battery started charging for Adapter %d
8038	Battery is discharging for Adapter %d
8039	Battery temperature is normal for Adapter %d
8040	Battery needs replacement for Adapter %d
8041	Battery Relearn started for Adapter %d
8042	Battery Relearn in progress for Adapter %d
8043	Battery Relearn completed for Adapter %d
8044	Battery Relearn timed out for Adapter %d

Trap Number	Trap Definition
8045	Battery Relearn pending; Battery is under charge for Adapter %d
8046	Battery Relearn postponed for Adapter %d
8047	Battery removed for Adapter %d
8048	Current capacity of the battery is below threshold for Adapter %d
8049	Current capacity of the battery is above threshold for Adapter %d
8050	Bbu FRU changed for Adapter %d
8051	Revision Identifier changed for Adapter %d
8052	Drive security key created for Adapter %d
8053	Drive security key backed up for Adapter %d
8054	Drive security key from escrow, verified for Adapter %d
8055	Drive security key changed for Adapter %d
8056	Drive security key, re-key operation failed for Adapter %d
8057	Drive security key is invalid for Adapter %d
8058	Drive security key destroyed for Adapter %d
8059	Drive security key from escrow is invalid for Adapter %d
8060	Drive security key failure, cannot access secured configuration for Adapter %d

B.2 SNMP Traps for Virtual Drives

The following table lists the supported SNMP traps and their definitions for events related to the virtual drives.

Table B.2 Supported SNMP Traps and Definitions for Virtual Drives

Trap Number	Trap Definition
8101	BGI aborted on Adapter -%d VD Target -%d
8102	BGI completed on Adapter -%d VD Target -%d
8103	BGI completed with uncorrectable errors on Adapter -%d VD Target -%d
8104	BGI failed on Adapter -%d VD Target -%d
8105	BGI started on Adapter -%d VD Target -%d
8107	CC started on Adapter -%d VD Target -%d
8108	CC completed on Adapter -%d VD Target -%d
8109	CC failed on Adapter -%d VD Target -%d
8110	CC aborted on Adapter -%d VD Target -%d
8111	CC completed with correction on Adapter -%d VD Target -%d
8112	Initialization aborted on Adapter -%d VD Target -%d
8113	Initialization failed on Adapter -%d VD Target -%d
8114	Initialization completed on Adapter -%d VD Target -%d
8115	Fast Initialization started on Adapter -%d VD Target -%d
8116	Full Initialization started on Adapter -%d VD Target -%d

Trap Number	Trap Definition
8117	Reconstruction started on Adapter -%d VD Target -%d
8118	Reconstruction completed on Adapter -%d VD Target -%d
8119	Reconstruction resumed on Adapter -%d VD Target -%d
8120	Reconstruction stopped on Adapter -%d VD Target -%d
8121	VD state changed on Adapter -%d VD Target -%d from %d to %d
8122	A new VD created on Adapter -%d VD Target -%d

B.3 SNMP Traps for Physical Drives

The following table lists the supported SNMP traps and their definitions for events related to the physical drives.

Table B.3 MegaRAID Trap Definitions

Trap Number	Trap Definition
8201	Physical Drive Clear aborted on Adapter -%d Dev -%d Enc -%d Slot-%d
8202	Physical Drive Clear failed on Adapter -%d Dev -%d Enc -%d Slot -%d
8203	Physical Drive Clear started on Adapter -%d Dev -%d Enc -%d Slot -%d
8204	Physical Drive Clear completed on Adapter -%d Dev -%d Enc -%d Slot -%d
8205	Error occurred on Adapter -%d Dev -%d Enc -%d Slot -%d
8206	Physical Drive Format started on Adapter -%d Dev -%d Enc -%d Slot -%d
8207	Physical Drive Format completed on Adapter -%d Dev -%d Enc -%d Slot -%d
8208	PD inserted on Adapter -%d Dev -%d Enc -%d Slot -%d
8209	Rebuild aborted on Adapter -%d Dev -%d Enc -%d Slot -%d

Table B.3 MegaRAID Trap Definitions

Trap Number	Trap Definition
8210	Rebuild completed on Adapter -%d Dev -%d Enc -%d Slot -%d
8211	Rebuild failed(bad source) on Adapter -%d Dev -%d Enc -%d Slot -%d
8212	Rebuild failed(bad target) on Adapter -%d Dev -%d Enc -%d Slot -%d
8213	Rebuild started on Adapter -%d Dev -%d Enc -%d Slot -%d
8214	Rebuild resumed on Adapter -%d Dev -%d Enc -%d Slot -%d
8215	Rebuild started(auto) on Adapter -%d Dev -%d Enc -%d Slot -%d
8216	Physical Drive removed on Adapter -%d Dev -%d Enc -%d Slot -%d
8217	PD state changed on Adapter -%d Dev -%d Enc -%d Slot -%d from %d to %d
8218	Redundant Path For Physical Drive Broken on Adapter -%d Dev -%d Enc -%d Slot -%d EncDevId -%d pdSAS -%s
8219	Redundant Path Restored For Physical Drive on Adapter -%d Dev -%d Enc -%d Slot -%d EncDevId -%d pdSAS -%s
8220	Redundant Encl Modul Inserted on Adapter -%d Dev -%d Enc -%d Slot -%d EncDevId -%d pdSAS -%s pdpath -%d
8221	Redundant Encl Modul Removed on Adapter -%d Dev -%d Enc -%d Slot -%d EncDevId -%d pdSAS -%s pdpath -%d
8222	FRU changed on Adapter -%d Dev -%d Enc -%d Slot -%d EncDevId -%d pdFRU -%s
8223	PD security activated on Adapter -%d Dev -%d Enc -%d Slot -%d
8224	PD security disabled on Adapter -%d Dev -%d Enc -%d Slot -%d

Table B.3 MegaRAID Trap Definitions

Trap Number	Trap Definition
8225	PD is reprovisioned on Adapter -%d Dev -%d Enc -%d Slot -%d
8226	PD security key changed on Adapter -%d Dev -%d Enc -%d Slot -%d
8227	Security subsystem problems detected on Adapter -%d Dev -%d Enc -%d Slot -%d
8228	Power State Change on Adapter -%d Dev -%d Enc -%d Slot -%d NewPowerState -%d
8229	PD Not Supported on Adapter -%d Dev -%d Enc -%d Slot -%d
8230	PD Not Certified on Adapter -%d Dev -%d Enc -%d Slot -%d

Appendix C

Glossary

access policy	A virtual drive property indicating what kind of access is allowed for a particular virtual drive. The possible values are <i>Read/Write</i> , <i>Read Only</i> , or <i>Blocked</i> .
alarm enabled	A controller property that indicates whether the controller's onboard alarm is enabled.
alarm present	A controller property that indicates whether the controller has an onboard alarm. If present and enabled, the alarm is sounded for certain error conditions.
array	See <i>drive group</i> .
BBU present	A controller property that indicates whether the controller has an onboard battery backup unit to provide power in case of a power failure.
BGI rate	A controller property indicating the rate at which the background initialization of virtual drives will be carried out.
BIOS	Basic Input/Output System. The computer BIOS is stored on a flash memory chip. The BIOS controls communications between the microprocessor and peripheral devices, such as the keyboard and the video controller, and miscellaneous functions, such as system messages.
cache	Fast memory that holds recently accessed data. Use of cache memory speeds subsequent access to the same data. When data is read from or written to main memory, a copy is also saved in cache memory with the associated main memory address. The cache memory software monitors the addresses of subsequent reads to see if the required data is already stored in cache memory. If it is already in cache memory (a cache hit), it is read from cache memory immediately and the main memory read is aborted (or not started). If the data is not cached (a cache miss), it is fetched from main memory and saved in cache memory.

cache flush interval	A controller property that indicates how often the data cache is flushed.
caching	The process of using a high speed memory buffer to speed up a computer system's overall read/write performance. The cache can be accessed at a higher speed than a drive subsystem. To improve read performance, the cache usually contains the most recently accessed data, as well as data from adjacent drive sectors. To improve write performance, the cache may temporarily store data in accordance with its write back policies.
capacity	A property that indicates the amount of storage space on a drive or virtual drive.
coerced capacity	A drive property indicating the capacity to which a drive has been coerced (forced) to make it compatible with other drives that are nominally the same capacity. For example, a 4 Gbyte drive from one manufacturer may be 4,196 Mbytes, and a 4 Gbyte from another manufacturer may be 4,128 Mbytes. These drives could be coerced to a usable capacity of 4,088 Mbytes each for use in a drive group in a storage configuration.
coercion mode	A controller property indicating the capacity to which drives of nominally identical capacity are coerced (forced) to make them usable in a storage configuration.
consistency check	An operation that verifies that all stripes in a virtual drive with a redundant RAID level are consistent and that automatically fixes any errors. For RAID 1 drive groups, this operation verifies correct mirrored data for each stripe.
consistency check rate	The rate at which consistency check operations are run on a computer system.
controller	A chip that controls the transfer of data between the microprocessor and memory or between the microprocessor and a peripheral device such as a drive. RAID controllers perform RAID functions such as striping and mirroring to provide data protection. MegaRAID Storage Manager software runs on ServeRAID-M controllers.
copyback	The procedure used to copy data from a source drive of a virtual drive to a destination drive that is not a part of the virtual drive. The copyback operation is often used to create or restore a specific physical

configuration for a drive group (for example, a specific arrangement of drive group members on the device I/O buses). The copyback operation can be run automatically or manually.

Typically, a drive fails or is expected to fail, and the data is rebuilt on a hot spare. The failed drive is replaced with a new drive. Then the data is copied from the hot spare to the new drive, and the hot spare reverts from a rebuild drive to its original hot spare status. The copyback operation runs as a background activity, and the virtual drive is still available online to the host.

current write policy

A virtual drive property that indicates whether the virtual drive currently supports Write Back mode or Write Through mode.

- In Write Back mode the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction.
- In Write Through mode the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction.

default write policy

A virtual drive property indicating whether the default write policy is Write Through or Write Back. In Write Back mode the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. In Write Through mode the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction.

device ID

A controller or drive property indicating the manufacturer-assigned device ID.

device port count

A controller property indicating the number of ports on the controller.

drive cache policy

A virtual drive property indicating whether the virtual drive cache is enabled, disabled, or unchanged from its previous setting.

drive group

A group of drives attached to a RAID controller on which one or more virtual drives can be created. All virtual drives in the drive group use all of the drives in the drive group.

drive state

A drive property indicating the status of the drive. A drive can be in one of the following states:

- **Unconfigured Good:** A drive accessible to the RAID controller but not configured as a part of a virtual drive or as a hot spare.
- **Hot Spare:** A drive that is configured as a hot spare.
- **Online:** A drive that can be accessed by the RAID controller and will be part of the virtual drive.
- **Rebuild:** A drive to which data is being written to restore full redundancy for a virtual drive.
- **Failed:** A drive that was originally configured as Online or Hot Spare, but on which the firmware detects an unrecoverable error.
- **Unconfigured Bad:** A drive on which the firmware detects an unrecoverable error; the drive was Unconfigured Good or the drive could not be initialized.
- **Missing:** A drive that was Online, but which has been removed from its location.
- **Offline:** A drive that is part of a virtual drive but which has invalid data as far as the RAID configuration is concerned.
- **None:** A drive with an unsupported flag set. An Unconfigured Good or Offline drive that has completed the prepare for removal operation.

drive subsystem	A collection of drives and the hardware that controls them and connects them to one or more controllers. The hardware can include an intelligent controller, or the drives can attach directly to a system I/O bus controller.
drive type	A drive property indicating the characteristics of the drive.
fast initialization	A mode of initialization that quickly writes zeroes to the first and last sectors of the virtual drive. This allows you to immediately start writing data to the virtual drive while the initialization is running in the background.
fault tolerance	The capability of the drive subsystem to undergo a single drive failure per drive group without compromising data integrity and processing capability. ServeRAID-M controllers provides fault tolerance through redundant drive groups in RAID levels 1, 5, 6, 10, 50, and 60. They also support hot spare drives and the auto-rebuild feature.
firmware	Software stored in read-only memory (ROM) or programmable ROM (PROM). Firmware is often responsible for the behavior of a system when it is first turned on. A typical example would be a monitor program

in a system that loads the full operating system from drive or from a network and then passes control to the operating system.

foreign configuration	A RAID configuration that already exists on a replacement set of drives that you install in a computer system. MegaRAID Storage Manager software allows you to import the existing configuration to the RAID controller or clear the configuration so you can create a new one.
formatting	The process of writing a specific value to all data fields on a drive, to map out unreadable or bad sectors. Because most drives are formatted when manufactured, formatting is usually done only if a drive generates many media errors.
hole	In MegaRAID Storage Manager, a <i>hole</i> is a block of empty space in a drive group that can be used to define a virtual drive.
host interface	A controller property indicating the type of interface used by the computer host system: for example, <i>PCIX</i> .
host port count	A controller property indicating the number of host data ports currently in use.
host system	Any computer system on which the controller is installed. Mainframes, workstations, and standalone desktop systems can all be considered host systems.
hot spare	<p>A standby drive that can automatically replace a failed drive in a virtual drive and prevent data from being lost. A hot spare can be dedicated to a single redundant drive group or it can be part of the global hot spare pool for all drive groups controlled by the controller.</p> <p>When a drive fails, MegaRAID Storage Manager software automatically uses a hot spare to replace it and then rebuilds the data from the failed drive to the hot spare. Hot spares can be used in RAID 1, 5, 10, and 50 storage configurations.</p>
initialization	The process of writing zeros to the data fields of a virtual drive and, in fault-tolerant RAID levels, generating the corresponding parity to put the virtual drive in a Ready state. Initialization erases all previous data on the drives. Drive groups will work without initializing, but they can fail a consistency check because the parity fields have not been generated.
IO policy	A virtual drive property indicating whether Cached I/O or Direct I/O is being used. In Cached I/O mode, all reads are buffered in cache

memory. In Direct I/O mode, reads are not buffered in cache memory. Data is transferred to cache and the host concurrently. If the same data block is read again, it comes from cache memory. (The IO Policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.)

load-balancing	A method of spreading work between two or more computers, network links, CPUs, drives, or other resources. Load balancing is used to maximize resource use, throughput, or response time.
media error count	A drive property indicating the number of errors that have been detected on the drive media.
migration	The process of moving virtual drives and hot spare drives from one controller to another by disconnecting the drives from one controller and attaching them to another one. The firmware on the new controller will detect and retain the virtual drive information on the drives.
mirroring	The process of providing complete data redundancy with two drives by maintaining an exact copy of one drive's data on the second drive. If one drive fails, the contents of the other drive can be used to maintain the integrity of the system and to rebuild the failed drive.
multipathing	The firmware provides support for detecting and using multiple paths from the RAID controllers to the SAS devices that are in enclosures. Devices connected to enclosures have multiple paths to them. With redundant paths to the same port of a device, if one path fails, another path can be used to communicate between the controller and the device. Using multiple paths with load balancing, instead of a single path, can increase reliability through redundancy.
name	A virtual drive property indicating the user-assigned name of the virtual drive.
non-redundant configuration	A RAID 0 virtual drive with data striped across two or more drives but without drive mirroring or parity. This provides for high data throughput but offers no protection in case of a drive failure.
NVRAM	Acronym for non-volatile random access memory. A storage system that does not lose the data stored on it when power is removed. NVRAM is used to store firmware and configuration data on the RAID controller.

NVRAM present	A controller property indicating whether an NVRAM is present on the controller.
NVRAM size	A controller property indicating the capacity of the controller's NVRAM.
offline	A drive is offline when it is part of a virtual drive but its data is not accessible to the virtual drive.
patrol read	A process that checks the drives in a storage configuration for drive errors that could lead to drive failure and lost data. The patrol read operation can find and sometimes fix any potential problem with drives prior to host access. This enhances overall system performance because error recovery during a normal I/O operation may not be necessary.
patrol read rate	The user-defined rate at which patrol read operations are run on a computer system.
product info	A drive property indicating the vendor-assigned model number of the drive.
product name	A controller property indicating the manufacturing name of the controller.
RAID	A group of multiple, independent drives that provide high performance by increasing the number of drives used for saving and accessing data. A RAID drive group improves input/output (I/O) performance and data availability. The group of drives appears to the host system as a single storage unit or as multiple virtual drives. Data throughput improves because several drives can be accessed simultaneously. RAID configurations also improve data storage availability and fault tolerance. Redundant RAID levels (RAID levels 1, 5, 6, 10, 50, and 60) provide data protection.
RAID 0	Uses data striping on two or more drives to provide high data throughput, especially for large files in an environment that requires no data redundancy.
RAID 1	Uses data mirroring on pairs of drives so that data written to one drive is simultaneously written to the other drive. RAID 1 works well for small databases or other small applications that require complete data redundancy.
RAID 5	Uses data striping and parity data across three or more drives (distributed parity) to provide high data throughput and data redundancy, especially for applications that require random access.

RAID 6	Uses data striping and parity data across three or more drives (distributed parity) to provide high data throughput and data redundancy, especially for applications that require random access. RAID 6 can survive the failure of two drives.
RAID 10	A combination of RAID 0 and RAID 1 that uses data striping across two mirrored drive groups. It provides high data throughput and complete data redundancy.
RAID 50	A combination of RAID 0 and RAID 5 that uses data striping across two drive groups with parity data. It provides high data throughput and complete data redundancy.
RAID 60	A combination of RAID 0 and RAID 6 that uses data striping across two drive groups with parity data. It provides high data throughput and complete data redundancy. RAID 60 can survive the failure of two drives in each RAID set in the spanned drive group.
RAID level	A virtual drive property indicating the RAID level of the virtual drive. ServeRAID-M controllers support RAID levels 0, 1, 5, 6, 10, 50, and 60.
raw capacity	A drive property indicating the actual full capacity of the drive before any coercion mode is applied to reduce the capacity.
read policy	A controller attribute indicating the current Read Policy mode. In Always Read Ahead mode, the controller reads sequentially ahead of requested data and stores the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data, but there is little improvement when accessing random data. In No Read Ahead mode, read ahead capability is disabled. In Adaptive Read Ahead mode, the controller begins using read ahead if the two most recent drive accesses occurred in sequential sectors. If the read requests are random, the controller reverts to No Read Ahead mode.
rebuild	The regeneration of all data to a replacement drive in a redundant virtual drive after a drive failure. A drive rebuild normally occurs without interrupting normal operations on the affected virtual drive, though some degradation of performance of the drive subsystem can occur.
rebuild rate	The percentage of central processing unit (CPU) resources devoted to rebuilding data onto a new drive after a drive in a storage configuration has failed.

reclaim virtual drive	A method of undoing the configuration of a new virtual drive. If you highlight the virtual drive in the Configuration Wizard and click the Reclaim button, the individual drives are removed from the virtual drive configuration.
reconstruction rate	The user-defined rate at which a reconstruction operation is carried out.
redundancy	A property of a storage configuration that prevents data from being lost when one drive fails in the configuration.
redundant configuration	A virtual drive that has redundant data on drives in the drive group that can be used to rebuild a failed drive. The redundant data can be parity data striped across multiple drives in a drive group, or it can be a complete mirrored copy of the data stored on a second drive. A redundant configuration protects the data in case a drive fails in the configuration.
revertible hot spare	When you use the Replace Member procedure, after data is copied from a hot spare to a new drive, the hot spare reverts from a rebuild drive to its original hot spare status.
revision level	A drive property that indicates the revision level of the drive's firmware.
SAS	Acronym for Serial Attached SCSI. SAS is a serial, point-to-point, enterprise-level device interface that leverages the Small Computer System Interface (SCSI) protocol set. The SAS interface provides improved performance, simplified cabling, smaller connectors, lower pin count, and lower power requirements when compared to parallel SCSI.
SATA	Acronym for Serial Advanced Technology Attachment. A physical storage interface standard. SATA is a serial link that provides point-to-point connections between devices. The thinner serial cables allow for better airflow within the system and permit smaller chassis designs.
SCSI device type	A drive property indicating the type of the device, such as drive.
SSD	Acronym for Solid State Devices. A Solid State Device uses solid-state memory to store data. They have no moving parts, and they are faster and more reliable than hard disk drives (HDDs).
serial no.	A controller property indicating the manufacturer-assigned serial number.

capacity	A virtual drive property indicating the amount of storage space on the virtual drive.
strip size	The portion of a stripe that resides on a single drive in the drive group.
stripe size	A virtual drive property indicating the length of the interleaved data segments that the RAID controller writes across multiple drives, not including parity drives. For example, consider a stripe that contains 64 KB of drive space and has 16 KB of data residing on each drive in the stripe. In this case, the stripe size is 64 KB and the strip size is 16 KB. The user can select the stripe size.
striping	A technique used to write data across all drives in a virtual drive. Each stripe consists of consecutive virtual drive data addresses that are mapped in fixed-size units to each drive in the virtual drive using a sequential pattern. For example, if the virtual drive includes five drives, the stripe writes data to drives one through five without repeating any of the drives. The amount of space consumed by a stripe is the same on each drive. Striping by itself does not provide data redundancy. Striping in combination with parity does provide data redundancy.
subvendor ID	A controller property that lists additional vendor ID information about the controller.
uncorrectable error count	A controller property that lists the number of uncorrectable errors detected on drives connected to the controller. If the error count reaches a certain level, a drive will be marked as failed.
vendor ID	A controller property indicating the vendor-assigned ID number of the controller.
vendor info	A drive property listing the name of the vendor of the drive.
virtual drive	A storage unit created by a RAID controller from one or more drives. Although a virtual drive may be created from several drives, it is seen by the operating system as a single drive. Depending on the RAID level used, the virtual drive may retain redundant data in case of a drive failure.
virtual drive state	A virtual drive property indicating the condition of the virtual drive. Examples include Optimal and Degraded.
write-back	In Write-Back Caching mode, the controller sends a data transfer completion signal to the host when the controller cache has received all

of the data in a drive write transaction. Data is written to the drive subsystem in accordance with policies set up by the controller. These policies include the amount of dirty/clean cache lines, the number of cache lines available, and elapsed time from the last cache flush.

write policy

See *Default Write Policy*.

write-through

In Write-Through Caching mode, the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data and has completed the write transaction to the drive.

Appendix D

Battery Glossary

Absolute state of charge	Predicted remaining battery capacity expressed as a percentage of Design Capacity. Note that the Absolute State of Charge operation can return values greater than 100 percent.
Auto learn mode	<p>The controller performs the learn cycle automatically in this mode. This mode offers the following options:</p> <p>BBU Auto Learn: Firmware tracks the time since the last learn cycle and performs a learn cycle when due.</p> <p>BBU Auto Learn Disabled: Firmware does not monitor or initiate a learn cycle. You can schedule learn cycles manually.</p> <p>BBU Auto Learn Warn: Firmware warns about a pending learn cycle. You can initiate a learn cycle manually. After the learn cycle is complete, the firmware resets the counter and warns you when the next learn cycle time is reached.</p>
Auto learn period	Time between learn cycles. A learn cycle is a battery calibration operation performed periodically by the controller to determine the condition of the battery.
Average time to empty	One-minute rolling average of the predicted remaining battery life.
Average time to full	Predicted time to charge the battery to a fully charged state based on the one minute rolling average of the charge current.
Battery module version	Current revision of the battery pack module.
Battery replacement	Warning issued by firmware that the battery can no longer support the required data retention time.

Battery retention time	Time, in hours, that the battery can maintain the contents of the cache memory.
Battery status	Operating status of the battery. Possible values are Missing, Optimal, Failed, Degraded (need attention), and Unknown.
Battery type	Possible values are intelligent Battery Backup Unit (BBU), intelligent Battery Backup Unit (iBBU), intelligent Transportable Battery Backup Unit (iTBBU), and ZCR Legacy.
Current	Measure of the current flowing to (+) or from (-) the battery, reported in milliamperes.
Cycle count	The count is based on the number of times the near fully charged battery has been discharged to a level below the cycle count threshold.
Design capacity	Designed charge capacity of the battery, measured in milliamperes-hour units (mAh).
Design charge capacity remaining	Amount of the charge capacity remaining, relative to the battery pack design capacity.
Design voltage	Designed voltage capacity of the battery, measured in millivolts (mV).
Device chemistry	Possible values are NiMH (nickel metal hydride) and LiON (lithium ion).
Estimated time to recharge	Estimated time necessary to complete recharge of the battery at the current charge rate.
Expected margin of error	Indicates how accurate the reported battery capacity is in terms of percentage.
Full charge capacity	Amount of charge that can be placed in the battery. This value represents the last measured full discharge of the battery. This value is updated on each learn cycle when the battery undergoes a qualified discharge from nearly full to a low battery level.
Gas gauge status	Hexadecimal value that represents the status flag bits in the gas gauge status register.
Learn cycle	Battery calibration operation performed by the controller periodically to determine the condition of the battery. You can start battery learn cycles manually or automatically.

Learn delay interval	Length of time between automatic learn cycles. You can delay the start of the learn cycles for up to 168 hours (seven days).
Learn mode	Mode for the battery auto learn cycle. Possible values are Auto, Disabled, and Warning.
Learn state	Indicates that a learn cycle is in progress.
Low-power storage mode	Storage mode that causes the battery pack to use less power, which save battery power consumption.
Manufacturing date	Date on which the battery pack assembly was manufactured.
Manufacturing name	Device code that indicates the manufacturer of the components used to make the battery assembly.
Max error	<p>Expected margin of error (percentage) in the state of charge calculation.</p> <p>For example, when Max Error returns 10 percent and Relative State of Charge returns 50 percent, the Relative State of charge is more likely between 50 percent and 60 percent. The gas gauge sets Max Error to 100 percent on a full reset. The gas gauge sets Max Error to 2 percent on completion of a learn cycle, unless the gas gauge limits the learn cycle to the +512/-256-mAh maximum adjustment values. If the learn cycle is limited, the gas gauge sets Max Error to 8 percent unless Max Error was already below 8 percent. In this case Max Error does not change. The gas gauge increments Max Error by 1 percent after four increments of Cycle Count without a learn cycle.</p>
Maximum learn delay from current start time	Maximum length of time between automatic learn cycles. You can delay the start of a learn cycle for a maximum of 168 hours (7 days).
Next learn time	Time at which the next learn cycle starts.
Predicted battery capacity status (hold 24hr charge)	Indicates whether the battery capacity is capable of supporting a 24-hour data retention time.
Relative state of charge	Predicted remaining battery capacity expressed as a percentage of Full Charge Capacity.

Remaining capacity	Amount of remaining charge capacity of the battery as stated in milliamp hours. This value represents the available capacity or energy in the battery at any given time. The gas gauge adjusts this value for charge, self-discharge, and leakage compensation factors.
Run time to empty	Predicted remaining battery life at the present rate of discharge in minutes.
Temperature	Temperature of the battery pack, measured in Celsius.

Appendix E

Notices

This information was developed for products and services offered in the U.S.A.

IBM® may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product, and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

E.1 Trademarks

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

IBM	ServeRAID	System x
IBM (logo)	ServerProven	

Intel, Intel Xeon, Itanium, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

Adobe and PostScript are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Red Hat, the Red Hat “Shadow Man” logo, and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

E.2 Important Notes

IBM makes no representation or warranties regarding non-IBM products and services that are ServerProven[®], including but not limited to the implied warranties of merchantability and fitness for a particular purpose. These products are offered and warranted solely by third parties.

IBM makes no representations or warranties with respect to non-IBM products. Support (if any) for the non-IBM products is provided by the third party, not IBM.

Some software might differ from its retail version (if available) and might not include user manuals or all program functionality.



Part number: 00D2418