



Embedded MegaRAID Software

User Guide

Revision 2.1
March 19, 2018

pub-005085

For a comprehensive list of changes to this document, see the [Revision History](#).

Corporate Headquarters

San Jose, CA

Website

www.broadcom.com

Broadcom, the pulse logo, Connecting everything, Avago Technologies, Avago, and the A logo are among the trademarks of Broadcom and/or its affiliates in the United States, certain other countries and/or the EU.

Copyright © 2018 by Broadcom. All Rights Reserved.

The term "Broadcom" refers to Broadcom Limited and/or its subsidiaries. For more information, please visit www.broadcom.com.

Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design.

Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

Table of Contents

Chapter 1: Overview	6
1.1 Embedded RAID Software Features	6
1.1.1 Device Support	6
1.1.2 RAID Features	7
1.1.3 Error Handling	7
1.1.4 Driver Features	7
1.1.5 BIOS Features	7
1.1.6 Unified Extensible Firmware Interface Features	8
1.1.7 RAID Management Utility Features	8
1.2 RAID Overview	8
1.2.1 RAID 0 Configuration Description	9
1.2.2 RAID 1 Configuration Description	10
1.2.3 RAID 5 Configuration Description	10
1.2.4 RAID 10 Configuration Description	11
Chapter 2: Driver Installation	12
2.1 Windows Driver Installation	12
2.1.1 Windows 7, Windows 2008, Windows Vista, Windows 10, and Windows 2012 Driver Installation	12
2.1.2 Updating the Windows Driver	12
2.1.3 Confirming the Windows Driver Installation	13
2.2 Linux Driver Installation	13
2.2.1 Obtaining the Driver Image File	13
2.2.2 Preparing the Installation Disks for Linux	13
2.2.2.1 Preparing Installation Disks with the Windows Operating System	13
2.2.2.2 Preparing Installation Disks with the Red Hat Linux 4 and 5 and SuSE Linux Operating Systems	14
2.2.3 Installing Red Hat Linux 6.x	14
2.2.4 Installing Red Hat Linux 7.x	18
2.2.4.1 Legacy Installation	19
2.2.5 Installing the Red Hat Linux Driver on a New System	20
2.2.6 Updating the Red Hat Linux Driver (Generic)	21
2.2.7 Enabling RAID Mode during Red Hat Linux 5 Driver Installation	22
2.2.8 Known Restrictions for the Driver Installation Process	23
2.2.8.1 Operating System using Driver Update Diskette Images	23
2.2.9 Installing the SuSE Linux Enterprise Server Operating System Driver	24
2.2.9.1 Installing SuSE Linux Enterprise Server 10, SP-4 Drivers	24
2.2.9.2 Installing the SLES-10, SP-4 XEN Drivers	29
2.2.10 Installing the SuSE Linux Enterprise Server 11 and 12 Driver	50
2.2.10.1 Loading the Xen Kernel RPM on SLES-11	58
2.2.10.2 Updating or Loading Software RAID Driver Kernel RPM on SuSE 12	59
2.2.11 Installing the SLES Operating System on Grantley Servers	60
2.2.11.1 Operating System Installation in RAID Mode	60
2.2.11.2 Operating System Installation in AHCI Mode	67
2.2.11.3 Installing the Software RAID Driver (RPM) Package	70
Chapter 3: Software RAID Configuration Utility	72
3.1 Performing a Quick Configuration	72
3.2 Management Menu	73
3.3 Configuration Menu	73
3.3.1 Configuration Menu Options	73
3.4 Configuring Drive Groups and Virtual Drives	74
3.5 Creating a Storage Configuration	75
3.5.1 Selecting the Configuration Method	75
3.5.2 Using the Easy Configuration Option	75
3.5.3 Using the New Configuration Option	79

3.5.4 Using the View/Add Configuration Option	82
3.6 Clearing a Storage Configuration	85
3.7 Configuring a Bootable Virtual Drive	86
3.8 Initializing Virtual Drives	86
3.8.1 First Initialization Method	87
3.8.2 Second Initialization Method	88
3.9 Rebuilding a Drive	89
3.10 Creating a Global Hot Spare Drive	89
3.11 Checking Data Consistency	90
3.12 Displaying and Changing Controller Properties	91
3.12.1 Displaying and Changing Controller Properties	91
3.12.2 Displaying and Changing Drive Properties	93
3.12.3 Viewing or Changing Virtual Drive Properties	94
3.13 Forcing Drives Online or Offline	95
3.14 Sector Size Not Supported	95
Chapter 4: Human Interface Infrastructure Configuration Utility	96
4.1 Managing Controllers	96
4.1.1 Viewing Controller Properties	96
4.1.2 Changing Controller Properties	100
4.1.3 Clearing Configurations	102
4.2 Managing Virtual Drives	102
4.2.1 Configuring Virtual Drives	103
4.2.2 Managing Virtual Drive Properties	105
4.2.3 Selecting Virtual Drive Operations	107
4.2.4 Viewing Drive Group Properties	110
4.3 Managing Drives	112
4.3.1 Viewing Drive Properties	112
4.3.2 Selecting Drive Operations	115
4.4 UDK2010 Support	116
4.5 4K Sector Size Support	116
Chapter 5: MegaRAID Storage Manager Utility	117
5.1 Overview	117
5.1.1 Creating Storage Configurations	117
5.1.2 Monitoring Storage Devices	117
5.1.3 Maintaining Storage Configurations	118
5.2 Hardware and Software Requirements	118
5.3 Installation	118
5.3.1 Installing the MegaRAID Storage Manager Utility on Microsoft Windows	118
5.3.2 Installing the MegaRAID Storage Manager Utility for Linux	121
5.3.2.1 Linux Installation Messages	121
Chapter 6: Configuration	123
6.1 Creating a New Storage Configuration	123
6.1.1 Selecting Virtual Drive Settings	123
6.1.2 Creating a Virtual Drive Using Simple Configuration	125
6.1.3 Creating a Virtual Drive Using Advanced Configuration	127
6.2 Changing Adjustable Task Rates	131
6.3 Changing Virtual Drive Properties	132
6.4 Deleting a Virtual Drive	134
Chapter 7: System Event and Storage Device Monitors	135
7.1 Monitoring System Events	135
7.2 Configuring Alert Notifications	137
7.2.1 Setting Alert Delivery Methods	139
7.2.2 Changing Alert Delivery Methods for Individual Events	139
7.2.3 Changing the Severity Level for Individual Events	140

7.2.4 Multiple Events Displayed in a Single Pop-Up Window	141
7.2.5 Entering or Editing the Sender Email Address and SMTP Server	142
7.2.6 Authenticating a Server	142
7.2.7 Saving Backup Configurations	143
7.2.8 Loading Backup Configurations	143
7.2.9 Adding Email Addresses of Recipients of Alert Notifications	143
7.2.10 Testing Email Addresses of Recipients of Alert Notifications	144
7.2.11 Removing Email Addresses of Recipients of Alert Notifications	145
7.3 Monitoring Controllers	145
7.4 Monitoring Drives	146
7.5 Running a Patrol Read Operation	147
7.5.1 Patrol Read Task Rates	149
7.6 Monitoring Virtual Drives	149
7.7 Monitoring Enclosures	151
7.8 Monitoring Rebuilds and Other Processes	151
Chapter 8: Storage Configuration Management and Maintenance	153
8.1 Initializing a Virtual Drive	153
8.1.1 Running a Group Initialization	153
8.2 Running a Consistency Check Operation	154
8.2.1 Running a Group Consistency Check	154
8.3 Scanning for New Drives	155
8.4 Rebuilding a Drive	156
8.5 Making a Drive Offline or Missing	156
Chapter 9: MegaRAID Storage Manager Screen and Menus	158
9.1 Starting MegaRAID Storage Manager	158
9.2 MegaRAID Storage Manager Main Menu Screen	160
9.2.1 Dashboard View/Physical View/Logical View	160
9.2.1.1 Dashboard View	160
9.2.1.2 Physical View	161
9.2.1.3 Logical View	162
9.2.2 Event Log Panel	164
9.2.3 Menu Bar	164
9.2.3.1 Manage Menu	164
9.2.3.2 Go To Menu	164
9.2.3.3 Log Menu	165
9.2.3.4 Tools Menu	165
9.2.3.5 Help Menu	165
Appendix A: Events and Messages	166
Glossary	178
Revision History	182

Chapter 1: Overview

This document describes the features of the Embedded MegaRAID® Software. It includes instructions to use the Avago® Software RAID configuration utility, the Human Interface Infrastructure Configuration Utility, and the MegaRAID Storage Manager™ configuration utility.

NOTE Information about the StorCLI command line utility is located in Chapter 6 of the *12Gb/s MegaRAID Software User Guide*. Information on converting MegaCLI command to StorCLI commands is located in Appendix C of the *12Gb/s MegaRAID Software User Guide*. Additional information about using the StorCLI command line utility with the Embedded MegaRAID software is located in Appendix D of the *12Gb/s MegaRAID Software User Guide*.

You can use these utilities to create redundant array of independent disks (RAID) storage configurations on drives controlled by the Embedded RAID utility. The manual also includes instructions to install the Embedded RAID drivers in Microsoft® Windows® systems and Linux® systems.

1.1 Embedded RAID Software Features

The Embedded RAID utility supports up to eight ports, which depends on the hardware platform. This support provides a cost-effective way to achieve higher transfer rates and reliability.

The following sections list the features available for devices, RAID, error handling, drivers, BIOS, Ctrl+M configuration utility, and RAID disk management.

1.1.1 Device Support

The Embedded RAID utility offers the following device support:

- Support for up to eight physical drives
- Support for SATA 6Gb/s drives
- Support for SAS 3Gb/s drives
- Support for solid state drives (SSDs)
- Support for 512e drives
- Support for SATA CD/DVD-ROM¹ drives
- Support for SATA tape devices¹
- Optical device (CD/DVD) hot plug feature used to connect optical devices while the operating system is running
- Hot plug support (online drive insertion and removal)
- Support for drive roaming
- Support for disk coercion (None, 128 MB, and 1 GB)
- Support SAS drives of maximum capacity available in the market

1. AHCI-based chipsets only.

1.1.2 RAID Features

The Embedded RAID utility supports the following RAID features:

- Support for RAID 0, RAID 1, RAID 5¹, and RAID 10 configurations
- Support for up to eight virtual drives
- Support for virtual drives larger than 2 TB
- Stripe size of 64 KB only
- Virtual drive availability immediately after creation
- Support for the random deletion of virtual drives
- Support for array cache setting (a RAID 10 volume is considered a single array, although it might have two, three, or four spans)
- Support for a migration path from Embedded RAID Software to MegaRAID SAS RAID controllers
- Check consistency for RAID 1, RAID 5, and RAID 10 configurations
- Drive group initialization support (fast and full)
- Support for auto or manual *rebuild* operations
- Ability to set the rates for the *background initialization* (BGI), *consistency check*, and *patrol read* operations
- Automatic resumption of *rebuild*, *check consistency*, *full initialization*, and *BGI* processes
- Global hot spare support

1.1.3 Error Handling

The Embedded RAID utility supports the following error handling features:

- Soft bad block management (SBBM) support
- Error/event logging and notification

1.1.4 Driver Features

The Embedded RAID driver supports the following features:

- Error logging and notification
- Support for Microsoft Windows Server® 2008, Microsoft Windows Server 2008R2, Microsoft Windows Vista®, Microsoft Windows Vista workstation, Microsoft Windows 7, Microsoft Windows 7 workstation, Microsoft Windows Blue workstation
- Support for Red Hat® Enterprise Linux (RHEL) 5.x, 6.x, and 7x
- Support for SuSE® Linux for 2.4, 2.6, and 3.0 kernels
- Support for SuSE Linux Enterprise (SLES) 10.x, 11.x and 12.x

1.1.5 BIOS Features

The Embedded RAID BIOS has the following features:

- Support for INT13 and enhanced disk drive specification
- Support for INT19h
- Option ROM size of 64 KB
- Support for BIOS boot specification (BBS) (If available in the system BIOS, this feature lets you select the controller from which to boot.)

1. The RAID 5 configuration is a premium feature.

- Support for power-on self-test (POST)
- Support for post memory management (PMM): Specification v7, July 2010
- Industry-standard extended BIOS data area (EBDA)
- POST and run-time BIOS support for device insertion and removal
- Support for *Stop On Error* during boot-up

The following feature is supported by the BIOS and the Ctrl+M configuration utility:

Automatic resumption of *rebuild*, *check consistency*, *full initialization*, and *BGI* (BGI is used in RAID 5 configurations only) operations

NOTE The BIOS program and the BIOS configuration utility (Ctrl+M) do not start or resume a *BGI* operation. If a *BGI* process is already in progress, you cannot start the *check consistency* operation.

1.1.6 Unified Extensible Firmware Interface Features

The Embedded RAID utility supports the following Unified Extensible Firmware Interface (UEFI) features:

- UEFI integration with American MegaTrends, Inc.® (AMI®) core systems with basic functionality
- Human Interface Infrastructure (HII) compatibility with AMI core system BIOS
- Driver health protocol (DHP) implementation with AMI core system BIOS

1.1.7 RAID Management Utility Features

The following features are available to manage the virtual drives and the physical drives in the system:

- Configuration information display (in the MegaRAID Storage Manager)
- Physical drive properties and virtual drive properties
- Drive group (array) management
- Error logging and notification
- Auto configuration support of newly added drives
- Ability to save and restore a configuration

1.2 RAID Overview

This section provides a brief overview of the types of RAID configurations that the Embedded RAID utility supports.

The first step in creating a RAID storage configuration is to configure drives in drive groups (also known as arrays). As defined for the Embedded RAID utility, a drive group is a group of one to eight physical disks that is seen by the host computer system as one large disk drive, or virtual drive. Only one RAID level can be assigned to a drive group.

- A RAID 0 drive group consists of one to eight drives.
- A RAID 1 drive group consists of two drives.
- A RAID 5 drive group consists of three to eight drives.
- A RAID 10 drive group consists of four, six, or eight drives.

NOTE Some hardware configurations do not support eight drives. Depending on the hardware, the actual maximum number of drives for RAID 0, RAID 5, and RAID 10 drive groups can be fewer than eight.

You can use any of these three strategies when creating RAID drive groups and virtual drives:

- **Maximize Fault Tolerance**

You can maximize fault tolerance to protect against loss of data by creating a RAID 1 drive group with mirroring. All data is written to the primary drive in the drive group and is also written (mirrored) to a second drive.

- **Maximize Virtual Drive Performance**

You can maximize virtual drive performance by creating a RAID 0 drive group with striping. Data is broken into segments and can be simultaneously written to or read from several different stripes on several different drives in the drive group.

RAID 10 drive groups combine both striping and mirroring to provide high data transfer rates and data redundancy.

- **Maximize Storage Capacity**

You can maximize storage capacity when selecting a RAID level. Striping alone (RAID 0) requires less storage space than mirrored data (RAID 1) or distributed parity (RAID 5). A RAID 5 drive group, which provides redundancy for one drive failure without duplicating the contents of entire drives, requires less space than a RAID 1 drive group.

1.2.1 RAID 0 Configuration Description

A RAID 0 configuration provides disk striping across all drives in the drive group. The RAID 0 configuration does not provide any data redundancy, but does offer the best performance of any RAID level. The RAID 0 configuration breaks up data into smaller segments called *strips*, and then *stripes* the data segments across each drive in the drive group. The size of each data segment is determined by the strip size, which is 64 KB.


NOTE It is possible to create each disk as a single-drive RAID 0 drive group. However, spanning across single drive RAID 0 drive groups is not supported.

By breaking up a large file into smaller segments, and writing to or reading from several drives at the same time, the Embedded MegaRAID Software utility can read from or write to the file faster. This feature makes the RAID 0 configuration ideal for applications that require high bandwidth but do not require fault tolerance.

- Uses** Provides high data throughput, especially for large files; use this configuration in 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; all data is lost if any drive fails.
- Drives** One to eight.

The following figure shows a RAID 0 drive group with two drives.

Figure 1 RAID 0 Drive Group Example with Two Drives




1.2.2 RAID 1 Configuration Description

The RAID 1 configuration duplicates all data from one drive to a second drive. The RAID 1 configuration provides complete data redundancy, but at the cost of doubling the required data storage capacity.

Uses	Databases or any other mission critical environment that requires fault tolerance.
Strong Points	Provides complete data redundancy; the RAID 1 configuration is ideal for any application that requires fault tolerance.
Weak Points	Requires twice as many drives; performance is impaired during drive rebuilds.
Drives	Two.

The following figure shows a RAID 1 drive group.

Figure 2 RAID 1 Drive Group



1.2.3 RAID 5 Configuration Description

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


NOTE The RAID 5 configuration is a premium feature. You might need to install a software key to enable a RAID 5 configuration. The key you need depends on your supplier. Contact your supplier for more information.

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

Uses	Provides high data throughput. Use the RAID 5 configuration for transaction processing applications because each drive can perform read and write operations 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 that require a lot of small write operations. Drive performance is reduced if a drive is being rebuilt or a background initialization is in progress. 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	Three to eight.

The following figure shows a RAID 5 drive group with six drives.

Figure 3 RAID 5 Drive Group



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

1.2.4 RAID 10 Configuration Description


The RAID 10 configuration, a combination of RAID 1 and RAID 0 drive groups, has mirrored drives. It breaks up data into smaller blocks, and then stripes the blocks of data to each RAID 1 RAID set. Each RAID 1 RAID set then duplicates its data to its other drive. The size of each block is determined by the strip size parameter, which is 64 KB. A RAID 10 drive group can sustain one drive failure in each drive group and maintain data integrity.

NOTE On a RAID 10 drive group, you can create only one virtual drive, and that virtual drive must occupy the entire space of the RAID 10 drive group.

- Uses** Works best for data storage that must provide 100 percent redundancy of a RAID 1 configuration (mirrored drive groups) and that also requires the enhanced I/O performance of a RAID 0 configuration (striped drive groups). The RAID 10 configuration 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.
- Drives** Four, six, or eight.

The following figure shows a RAID 10 drive group with four drives.

Figure 4 RAID 10 Drive Group



Chapter 2: Driver Installation

This chapter explains how to install the Embedded RAID Software drivers for the following operating systems:

- Microsoft Windows Server 2008, 2008R2, 2012, and 2012R2
- Microsoft Windows 7 Workstation
- Microsoft Windows Vista Workstation
- Microsoft Windows 8 Workstation
- Microsoft Windows 8.1 Workstation
- Microsoft Windows 10 Workstation
- Red Hat Enterprise Linux® (RHEL) 5.x, 6.x, and 7.x
- SuSE Linux Enterprise (SLES) 10, 11, and 12

2.1 Windows Driver Installation

2.1.1 Windows 7, Windows 2008, Windows Vista, Windows 10, and Windows 2012 Driver Installation

Perform the following steps to install the MegaRAID device driver in a new Windows 7, Windows 8, Windows 8.1, Windows 2008, Windows Vista, Windows 2010, or Windows 2012, operating system.

1. Start the Windows 7, Windows 8, Windows 8.1, Windows 2008, Windows Vista, Windows 2010, or Windows 2012 installation by booting from the appropriate Windows DVD.
The system BIOS must support booting from a DVD. The BIOS settings might require changes to allow DVD booting. Refer to your system documentation.
Windows loads the file and the first installation screen appears.
2. Select your language and other settings based on your location and preference.
3. Click **Next**.
The **Windows Install** screen appears.
4. Press **Install Now** to start the installation wizard.
The **Software License** screen appears.
5. Click the check box to accept the software license and click **Next**.
The next installation screen appears.
6. Select the type of installation you want to perform.
7. Follow the prompts to select the location where you want to install Windows and click **Next**.
The program installs the files. Your system will restart several times during the installation process.

2.1.2 Updating the Windows Driver

Perform the following steps to update the Embedded RAID driver for Windows or to install this driver on an existing system booted from a standard IDE (integrated drive electronics) drive.

1. Click **Start**, point to **Settings**, and click **Control Panel**.
2. Double-click **System**, click the **Hardware** tab, and click **Device Manager**.
The Device Manager starts.

3. In the **Device Manager** screen, double-click **SCSI and RAID Controllers**, right-click the device on which you are installing the driver, and click **Properties**.
4. On the **Driver** tab, click **Update Driver** to open the **Update Device Driver** wizard, and follow the wizard instructions to update the driver.

2.1.3 Confirming the Windows Driver Installation

Perform the following steps to confirm that the Embedded RAID driver for Windows is installed correctly.

1. Click **Start**, point to **Settings**, and click **Control Panel**.
2. Double-click **System**, click the **Hardware** tab, and click **Device Manager**.
The Device Manager starts.
3. In the **Device Manager** screen, double-click **SCSI and RAID Controllers**, right-click the device for which you are installing the driver, and click **Properties**.
4. On the **Driver** tab, click **Driver Details** and verify that the driver information is correct.

2.2 Linux Driver Installation

This section describes the steps to install the Embedded RAID device driver in a Red Hat Enterprise Linux (RHEL) installation or a SuSE Linux Enterprise Server (SLES) installation.

2.2.1 Obtaining the Driver Image File

The Linux driver is offered in the form of a driver update disk. The required file is `dud-[driver version].img`, which is the driver update disk for the Embedded RAID stack.

2.2.2 Preparing the Installation Disks for Linux

This section describes how to prepare the Linux installation disks from the driver image files, using either the Windows operating system or the Linux operating system.

2.2.2.1 Preparing Installation Disks with the Windows Operating System

Under Windows, you can use the `RaWrite` floppy image writer utility to create disk images from image files. The image writer can be downloaded from the Internet. Perform the following steps to build installation diskettes.

NOTE The installation instructions in this section are applicable for RHEL 4.x and RHEL 5.x only.

1. Copy the driver update disk (DUD) image `dud-[driver version].img` and the `raw write.exe` file from the DUD to a directory.
2. Confirm that the files were copied into the selected directory.
3. If necessary, use this command to change the filename of the driver update disk to a name with fewer than eight characters:

```
copy dud-[driver version].img dud.img
```

4. Open the DOS® command prompt, and navigate to the directory where the `raw write.exe` file is located.
5. Type the following command to create the installation diskette:

```
raw write
```

6. Press Enter.
You are prompted to enter the name of the boot image file.
7. Type the following:
`dud.img`
8. Press Enter.
You are prompted for the target drive diskette.
9. Insert a floppy diskette into the floppy drive and type:
`A:`
10. Press Enter.
11. Press Enter again to start copying the file to the diskette.
12. After the command prompt returns and the floppy disk drive LED goes out, remove the diskette.
13. Label the diskette with the image name.

2.2.2.2 Preparing Installation Disks with the Red Hat Linux 4 and 5 and SuSE Linux Operating Systems

Under Red Hat Linux 4 and 5 and SuSE Linux, you can use a driver diskette utility to create disk images from image files. Perform the following steps to create the driver update disk:

1. Copy the DUD image `dud-[driver version].img` to a Linux system.
2. Insert a blank floppy diskette into the floppy drive.
3. Confirm that the files are in the selected directory.
4. Create the DUD using the following command:
`dd if=dud-[driver version].img of=/dev/fd0`
5. After the command prompt returns and the floppy disk drive LED goes out, remove the diskette.
6. Label the diskette with the image name.

2.2.3 Installing Red Hat Linux 6.x

This section describes installation of the MegaRAID device driver for the Red Hat Enterprise Linux (RHEL) 6.x operating systems. The installation method is same for both 32-bit and 64-bit installations.

NOTE

- EFI mode of installation is applicable only to 64-bit OS.
- Kernel-based virtual machine (KVM) hypervisor is supported only with 64-bit OS.
- The TBOOT enabled in RHEL 6.x is not supported on UEFI boot.
You need to disable the `tboot` command.

For more information refer the RedHat Knowledgebase article:

<https://access.redhat.com/knowledge/articles/186583>

To disable the `tboot` command.

http://support.lenovo.com/en_IN/downloads/detail.page?DocID=HT075591

1. Copy the MegaSR drive image (`megasr-* .img`) from the DUD folder to a USB drive.
2. Configure the BIOS and OpROM settings.

The following window appears after the OS is loaded from the drive.

NOTE

The following screenshots are from RHEL 6.1 installation. Slight variations between different versions of RHEL 6.x might exist.

Figure 5 Red Hat Enterprise Linux 6.x Welcome Screen



3. Choose the first option (**Install or upgrade an existing system**) and press Tab to edit the boot argument.

Provide the following command as a boot argument.

```
linux dd blacklist=iscsi nodmraid
```

This argument makes sure that the SCU controller is claimed by the MegaSR driver. Avago supports the AHCI controller for some OEMs. On select OEM hardware, for the MegaSR driver to claim both SCU and AHCI controllers, use the following command.

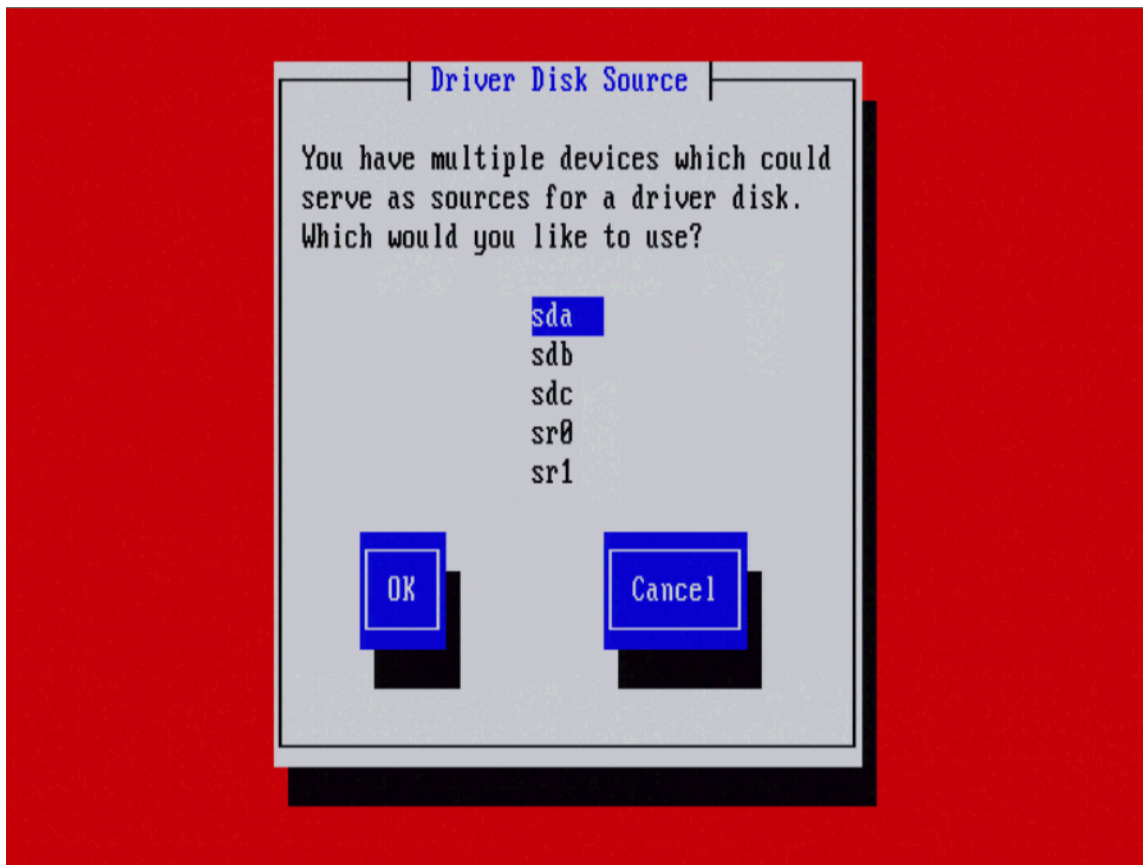
```
linux dd blacklist=ahci,iscsi nodmraid
```

4. For EFI based installation press ESC and use the below commands

```
linux virtEFI dd blacklist=<inbox driver,...> nodmraid
```

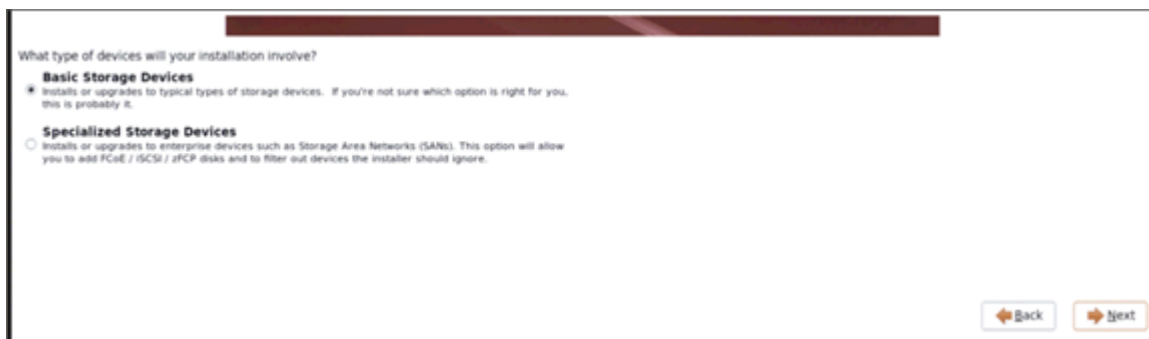
The `virtEFI` argument is not a mandatory argument for some Linux versions during EFI mode installation.

Figure 6 MegaSR Driver Messages Screen



5. Select the driver disk and click **OK**.
6. Browse through the driver disk and select the driver image you want to install.
The CD/DVD-ROM label begins with `sr` (for example, `sr0`, `sr1`).
To confirm whether the driver is loaded, press `Ctrl+Shift+F4`. In the terminal, messages related to loading the MegaSR driver appear.
7. Select which device to use as a source and click **OK** to continue.
The **Device Installation** window appears.

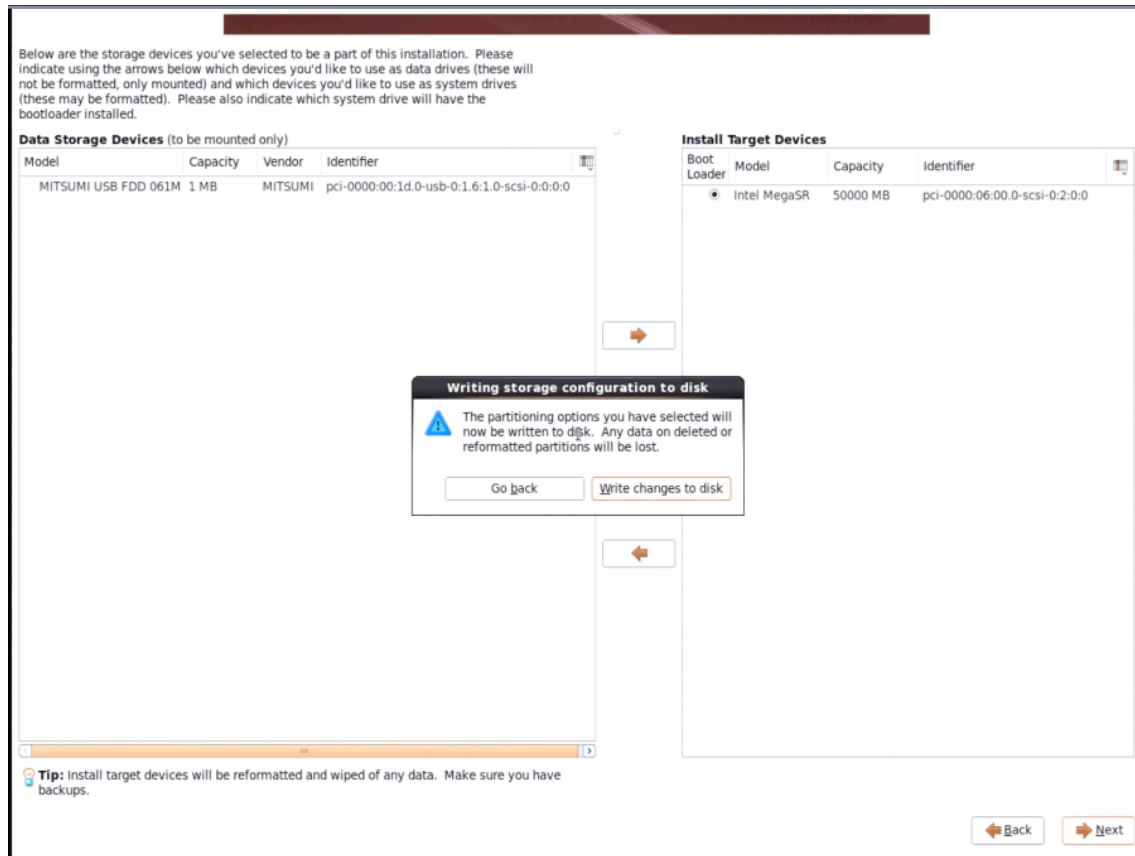
Figure 7 Device Installation Screen



8. Select **Basic storage devices** and click **Next**.

9. Click **Yes** to discard any data.
The screen shows the virtual drive discovered after loading the MegaSR driver.
10. Select the target installation medium from the set of **Target Data Storage** devices in the right pane.
Use the arrow button in the center of the screen to move the selected device to the **Install Target Devices** pane and proceed. In the pop-up window that appears click **Write changes to disk**.

Figure 8 Target Installation Screen




11. For XEN/KVM based installations select **Virtual Host** as the installation type.

NOTE Make sure that the BIOS has Virtualization enabled.

Or you can select any other mode based on the requirement.

To install additional packages, select the **Customize now** radio button at the bottom of the screen. A list of additional software packages appears. Select the appropriate packages.

Figure 9 Red Hat Enterprise Linux Installation Type Screen



12. After finishing installation, re-boot from the HDD to the installed RHEL kernel.

2.2.4 Installing Red Hat Linux 7.x

NOTE The RHEL 7.1 installation works with only .iso images and not with .img images.

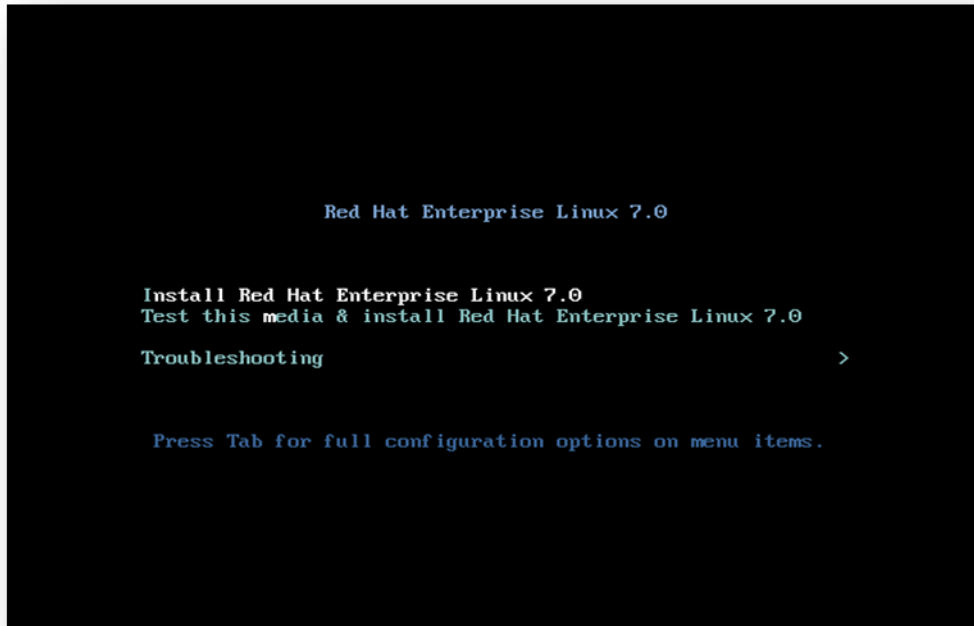
1. Prepare a driver disk using USB disk.
You cannot use a floppy disk because of the size of MegaSR DUD image.
 - a. Copy the DUD image file to any Linux system.
Ex: `megasr-16.01.2014.0611-1-rhel70-ga-x86_64.img` for RHEL 7-64bit
 - b. Create a directory and mount the DUD image to that directory using the following command.
Syntax: `mount -o loop <source_folder> <destination_folder>`
Ex: `mkdir image`
`mount -o loop megasr-16.01.2014.0611-1-rhel70-ga-x86_64.img image`
 - c. Copy the image folder to a USB drive.
You can use a formatted USB drive a driver disk.
 - d. Connect the physical drives (PDs) to a controller that is configured in a RAID mode.
Create a RAID volume for the installation.
2. Configure the BIOS and OpROM settings.

2.2.4.1 Legacy Installation

The following window appears when the system is booted from an Installation disk using one of the following options.

- Install Red Hat Enterprise Linux 7.0
- Test this media & install Red Hat Enterprise Linux 7.0
- Troubleshooting

Figure 10 Red Hat Enterprise Linux 7.0 Installation Screen



1. Choose the **Install Red Hat Enterprise Linux 7.0** option and press Esc to edit the boot argument as described in the following steps.
2. Provide the following boot argument.

```
linux dd modprobe.blacklist=iscsi
```

This argument makes sure that the SCU controller is claimed by the MegaSR driver.
3. Use the following boot argument on select hardware for the MegaSR driver to claim both the SCU and ACHI controllers.

```
linux dd modprobe.blacklist=ahci,iscsi
```

NOTE Avago supports the AHCI controller for certain OEMs.

Figure 11 Linux AHCI Controller Command



```
boot: linux dd modprobe.blacklist=ahci_
```

2.2.5 Installing the Red Hat Linux Driver on a New System

This section describes the fresh installation of the Red Hat Enterprise Linux 4 or 5 device driver on systems with the Embedded RAID stack. After you prepare the installation disks with the driver image, perform the following steps to install the driver on a new system.

1. Boot to CD-ROM (Disk 1).
The **Red Hat Introductory** screen appears.
2. Type the following instruction at the boot prompt:
`linux dd noprobe=<value> (<value> depends on the number of drives)`
For example, to install Red Hat Linux on a RAID 5 configuration with three drives, enter:
`linux dd noprobe=ata1 noprobe=ata2 noprobe=ata3`
3. Press Enter.
The prompt asks whether you have a driver disk.
4. Use the arrow key to select **Yes**, and then press Enter.
5. Select **fd0** to indicate you have a floppy diskette with the driver on it.
6. Insert the floppy diskette in the **A :** / drive and press Enter.
The installer locates and loads the driver for your device. The following message appears:
`Loading megasr driver...`
The prompt at the next screen asks whether you have another driver.
7. Follow the Red Hat Linux installation procedure to complete the installation.
8. Before you reboot, go to a text console and perform the following steps:
 - a. Press **Ctrl+Alt+F2** to access the text console.
 - b. Enter the command `cat /proc/partitions` to get the major and minor number of the floppy drive.
 - c. Execute the `mknod /dev/sd(x) b major minor` command.
 - d. Create a directory, such as a `mkdir swr` directory.
 - e. Mount the floppy drive to that directory with the `mount /dev/sd(x) swr` command.
 - f. Run the `./replaceachi.sh` script.

9. Reboot the system.

2.2.6 Updating the Red Hat Linux Driver (Generic)

Perform the following steps to update the Red Hat Linux driver or to install the Red Hat Linux driver in an existing system booted from a standard SATA drive or systems with the Embedded Software RAID stack.

NOTE The following steps show example code snippets, you must change the values according to your environment.

1. Boot the system with the Red Hat Linux Installation CD from the primary controller or disk. The **Red Hat Introductory** screen appears.
2. Mount the driver update diskette using the following command:
3. Unzip the `modules.cgz` file that is on the DUD to get driver images for different Red Hat operating systems:

```
#mount /dev/fd0 /mnt/floppy
#mkdir -p /home/megasr
#cd /home/megasr
#cp /mnt/floppy/modules.cgz
#gunzip -S .cgz modules.cgz
```

This action generates the `modules` file:

```
#cpio -ivd < modules
```

This action provides the following driver images:

```
<kernel version>/megasr.o
<kernel version>smp/megasr.o
<kernel version>BOOT/megasr.o
```

4. Update the MegaSR driver module for the required kernels using the following commands:

```
#cd /home/megasr
```

If the `/lib/modules/<kernel version>/update/` directory is present, use the following command:

```
# cp <kernel version>/megasr.[o/ko]
/lib/modules/<kernel version>/update/megasr.[o/ko]
```

If the `/lib/modules/<kernel version>/update/` directory is not present, use the following command:

```
# cp <kernel version>/megasr.[o/ko]
/lib/modules/<kernel version>
/kernel/drivers/scsi/megasr.[o/ko]
```

5. Create a MegaSR driver entry in the configuration file.

The Red Hat configuration file is `/etc/modules.conf`. If the MegaSR entry is not present in `/etc/modules.conf`, add the following line:

```
alias scsi_hostadapter megasr
```

If the `ahci` SCSI driver entry (located on the following paragraph) is present in the `/etc/modules.conf` file, remove it. You must remove the entry or the `ahci` driver will take control of the RAID controller without checking the subsystem device or Vendor ID. The `ahci` SCSI driver entry is as follows:

```
alias scsi_hostadapter ahci
```

6. Create a new `initrd` image for the required kernel.

Red Hat installation uses the `mk_initrd` command to create an `initrd` image. The following command creates an `initrd` image for the `<kernel version>smp` kernel in the boot directory. Refer to the `mk_initrd` man page for more information. The command is as follows:

```
#mkinitrd /boot/initrd<kernel version>smp.img.new <kernel version>smp
```

7. Modify the `lilo.conf`/`grub.conf` file by adding any newly created `initrd` images as new entries in the `/etc/lilo.conf` file.

The suggested method is to copy an existing `lilo` entry in the file and paste it as a new one. Then modify its kernel image name, `initrd` image name, and label name.

Sample Lilo Entry

```
image=/boot/vmlinuz-<kernel version>smp label=linuxnew
initrd=/boot/initrd-<kernel version>smp.img.new
read-only appended=root=LABEL=/ "
```

Sample Grub Entry

```
title Red Hat Linux (<kernel version> with Megasr driver)
root (hd0,0)
kernel /vmlinuz-<kernel version> ro root=LABEL=/
initrd /initrd-<kernel version>.img.new
```

8. Update the boot loader.

If the boot loader is Lilo, run the `lilo` command to update the boot loader:

```
#lilo
```

9. Reboot the system to the new boot loader entry.

2.2.7 Enabling RAID Mode during Red Hat Linux 5 Driver Installation

This section documents how to load the RHEL5 operating system drivers so the drivers recognize RAID mode and support RAID 0, RAID 1, RAID 5, and RAID 10 functionality.

On servers using on-board, Advanced Host Controller Interface (AHCI) based SATA controllers, the controllers can be set to host bus adapter (HBA) mode or RAID mode using the system BIOS. The system BIOS changes the PCI sub-system IDs based on the selected mode. The MegaRAID software RAID (MegaSR) driver can be set to RAID mode for the AHCI-based SATA controllers and support RAID level 0, RAID level 1, RAID level 5, and RAID level 10. The MegaSR driver considers all four PCI IDs to decide whether the driver supports the controller.

In Red Hat Enterprise Linux 5 (RHEL 5), however, the native AHCI driver Linux kernel looks only at the vendor ID and the device ID, and ignores the subsystem IDs. As a result, the RHEL 5 operating system does not support RAID mode and ignores the mode. The RHEL 5 installation process loads the native AHCI driver even when the MegaSR DUD is present. Also, the process includes the AHCI driver in the `initrd` images (during which a temporary file system is loaded into memory in the Linux kernel boot process).

Perform the following steps to load the RHEL 5 operating system drivers so they recognize the RAID mode.

1. Use the `noprobe` option at the installation kernel boot prompt, in addition to the `dd` option.

This step prevents the AHCI driver from loading ahead of the MegaSR driver, letting you install the RHEL 5 operating system on the `megasr` virtual drive.

The following string is an example of the `noprobe` option:

```
boot: linux dd noprobe=ata1 noprobe=ata2 noprobe=ata3 noprobe=ata4
```

At the last installation step, after all the packages are installed, the RHEL 5 operating system prompts you to reboot. At this point, the `initrd` image is built with the AHCI driver. You must add `megasr` and delete `ahci` from the `initrd` image before you reboot.

2. Press `Ctrl+Alt+F2`.

This step takes you to a text console prompt.

3. Before you reboot, perform these steps:

- a. Press `Ctrl+Alt+F2` to go to the text console.
- b. Enter the `cat /proc/partitions` command to get the major and minor number of the floppy drive.

- c. Enter the `mknod /dev/sd(x) b major minor` command.
- d. Create a directory (for example, `mkdir swr`).
- e. Mount the floppy drive to that directory with the `mount /dev/sd(x) swr` command.
- f. Run the `./replaceachi.sh` script.

If the script is on the floppy drive, you must mount the device first. In the normal system, use `/dev/sd<x>` device names to mount a device. In the pre-installation environment, these names might not exist. You can find the major and minor numbers of your device by reading the `/proc/partitions` file.

4. Use the major and minor numbers of your device in the `mknod` command to create your own device name.

2.2.8 Known Restrictions for the Driver Installation Process

This section documents known restrictions that you must follow when you install the operating system drivers for various operating systems.

2.2.8.1 Operating System using Driver Update Diskette Images

Follow these steps when you install the operating system using DUD images:

- For SLES11 GA, SP1, SP2, SP3, and SP4 variants, with SLES, (32-bit and 64-bit) platforms:
Enter `brokenmodules=ahci` while installing the driver.
Enter `brokenmodules=ahci brokenmodules=iscsi` while installing the driver.
- For RHEL 6.x general availability (GA) (32-bit and 64-bit) platform operating system installation using the Software RAID DUD images:
 - a. While booting from DVD, press Esc so you can install the third party driver.
 - b. Enter the following command to install the driver:

```
Linux dd blacklist=ahci
```
- For RHEL 7.x general availability (32-bit and 64-bit) platform operating system installation using the Software RAID DUD images:
 - a. While booting from DVD, press Esc so you can install the third party driver.
 - b. Enter the following command to install the driver:

```
Linux dd modprobe.blacklist=ahci
```
- For RHEL 5 U7 to U11 (32-bit and 64-bit) platform operating system installation using the SWR DUD images, follow these steps:
 - a. While booting from DVD, press Esc so you can install the third party driver.
 - b. Enter the following command to install the driver:

```
Linux dd blacklist=iscsi blacklist=ahci
```
- For all SLES platforms (32-bit and 64-bit), platform operating system installation using the SWR DUD images, follow the steps in this example for SLES 11 SP1 (32-bit and 64-bit) platforms, in which the DUD image size exceeds the floppy disk size:
 - a. Copy the image file to any Linux system.
The image files are as follows:

```
megasr-14.00.0722.2010-1-sles11-sp1-x86.img for 32-bit SLES11 SP1  
megasr-14.00.0722.2010-1-sles11-sp1-x86_64.img for 64-bit SLES11 SP1
```
 - b. Create a directory (for example, `mkdir image`).
 - c. Mount the DUD image on the `image` directory using the following command:

```
mount -o loop megasr-14.00.0722.2010-1-sles11-sp1-<arch>.img image.
```
 - d. Change the directory to `image` (for example, `cd image`).
 - e. Copy the contents of the `image` directory to a USB drive.
Contents of the image start with the `/01` directory, so copy the `/01` directory to a USB drive.

- f. Use the USB drive to provide the third-party driver during the operating system installation.
- g. Enter the following command while installing the driver.

```
brokenmodules=ahci
```
- h. Make sure the USB drive is formatted before using it for operating system installation using DUD images.

2.2.9 Installing the SuSE Linux Enterprise Server Operating System Driver

This section describes the fresh installation of the SuSE Linux Enterprise Server (SLES) 10, 11, or 12 driver on a system with the Embedded RAID stack. Prepare installation disks with the driver image, and then perform the following steps to install the driver.

1. Create a RAID drive group using one of the configuration utilities.
2. Boot the system using the SLES disk.
3. When the first screen appears, select **Installation** on the menu.
4. Type the following in the **Boot Options** field:

```
brokenmodules=achi
```
5. Press F6 for the driver and select **Yes**.
6. Insert the driver update diskette in the A : / drive and press Enter.
Yes appears under the **F6 Driver** heading.
7. Click **OK**.
The following message appears:

```
LSI Soft RAID Driver Updates added.
```
8. At the menu, select the driver update medium and click the **Back** button.
9. Continue and complete the installation process.

2.2.9.1 Installing SuSE Linux Enterprise Server 10, SP-4 Drivers

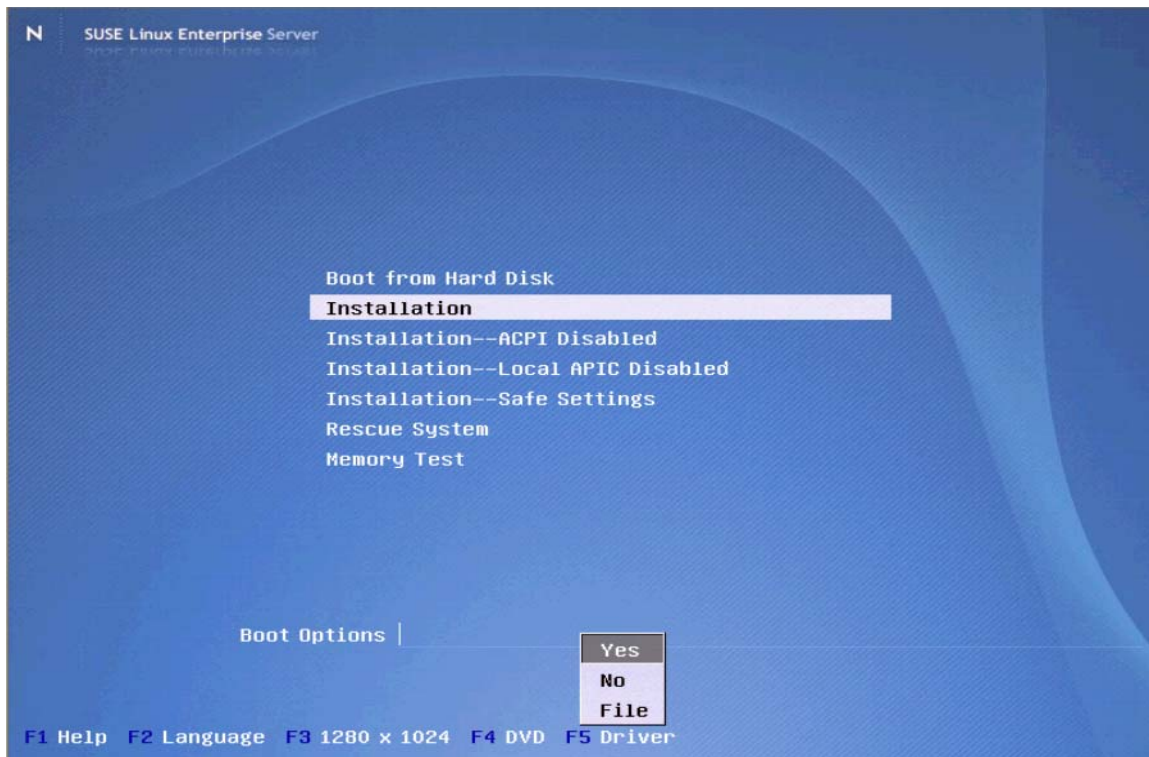
Use the following steps to install the SLES-10-SP4 drivers. The installation is the same for both 32-bit and 64-bit drivers.

1. Prepare a driver disk using a USB disk.

NOTE You cannot use a floppy disk because of the size of the MegaSR DUD image.

- a. Copy the DUD image file to any Linux system.
Ex: `megasr-14.00.0722.2010-1-sles11-sp1-x86_64.img` for SLES11-SP1-64bit
 - b. Create a directory and mount the DUD image to that directory using the following command.
Syntax: `mount -o loop <source> <destination_folder>`
Ex: `mkdir image`
`mount -o loop megasr-14.00.0722.2010-1-sles11-sp1- x86_64.img image`
 - c. Copy the contents of the sub-directories called 01 and 02 to the image directory of your USB drive.
Use a formatted USB drive to make the driver disk.
2. Configure the BIOS and OpROM settings.
 3. Connect the USB drive that contains the MegaSR driver and reboot the system.
The following screen appears.

Figure 12 SuSE Linux Enterprise Server Splash Screen



4. Choose the **Installation** option.
5. Press F5 and click **Yes** to load the driver.
This option appears only if the installation is using NON EFI mode.

NOTE Avago also supports AHCI controllers on select hardware.

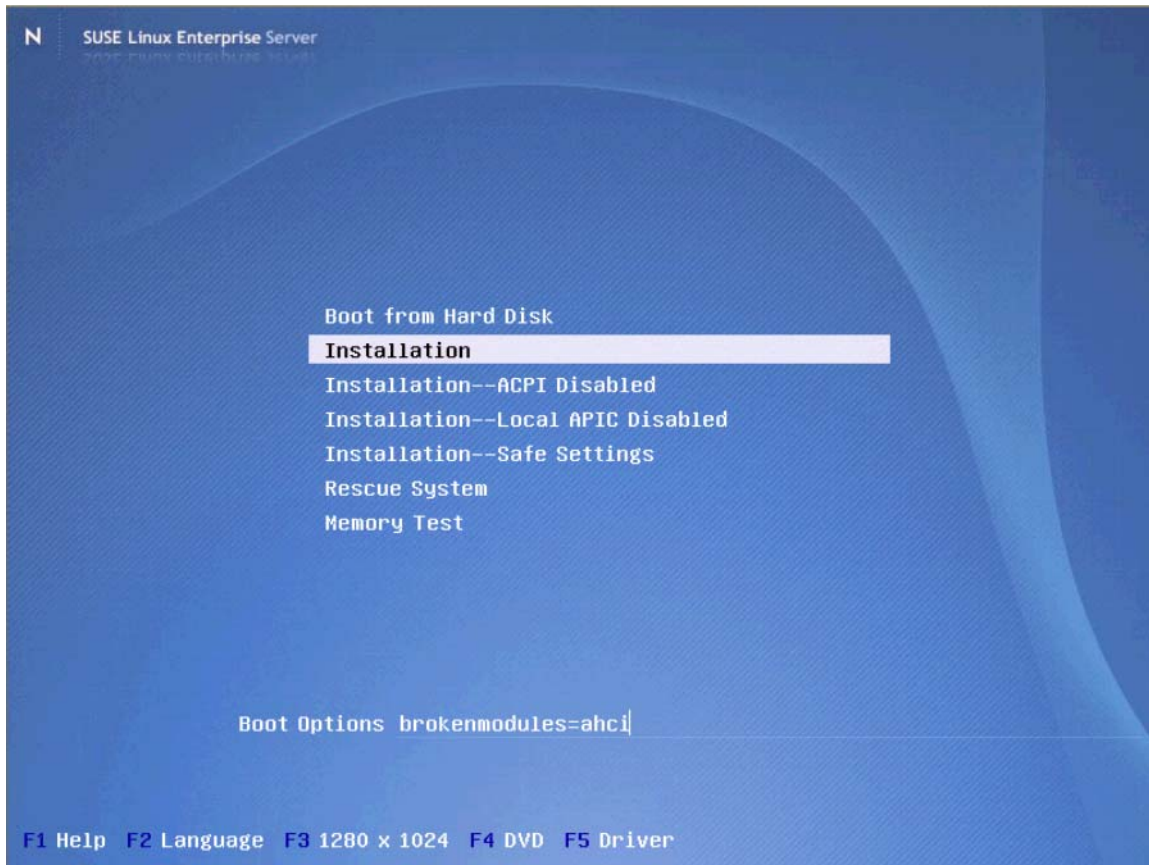
6. Provide one of the following commands as the boot argument.

```
brokenmodules=iscsi  
brokenmodules=ahci
```

These arguments ensure that either the SCU controller or AHCI controller is claimed by the MegaSR driver, respectively. For the MegaSR driver to claim both the SCU and AHCI controllers, use the following command.

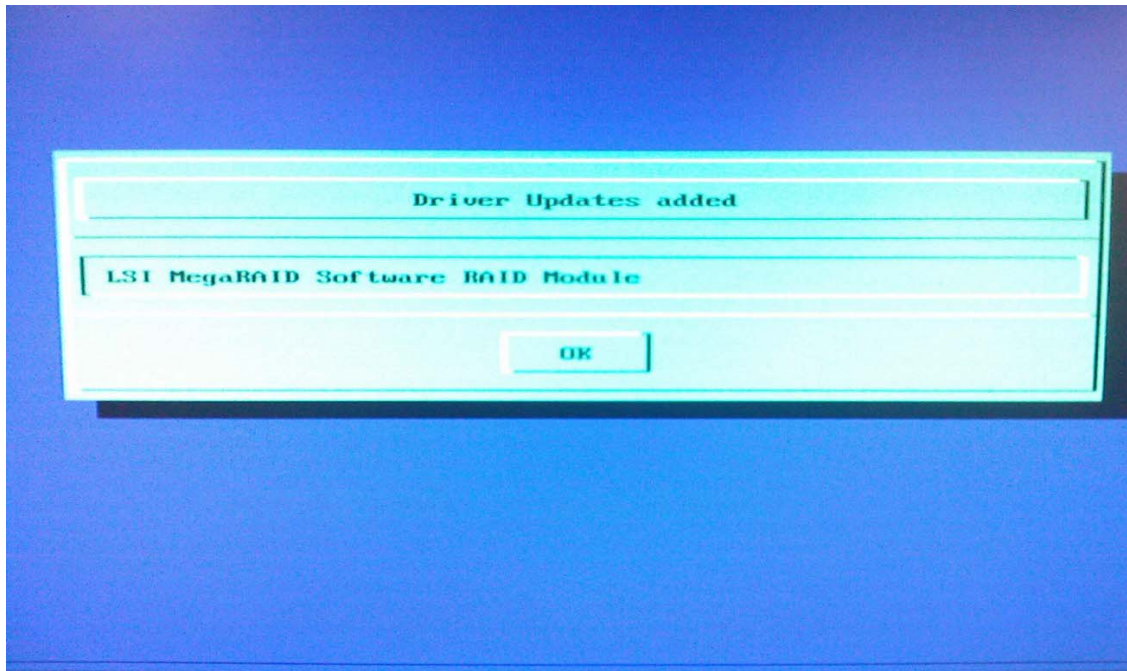
```
brokenmodules=iscsi,ahci
```

Figure 13 SuSE Linux Enterprise Server Boot Options Screen



The following screen appears.

Figure 14 Driver Updates Message Screen




This message shows that the MegaSR driver that is present on the USB drive is recognized and loaded.

7. Click OK to proceed.

The following window appears.

Figure 15 Driver Update Choice Screen

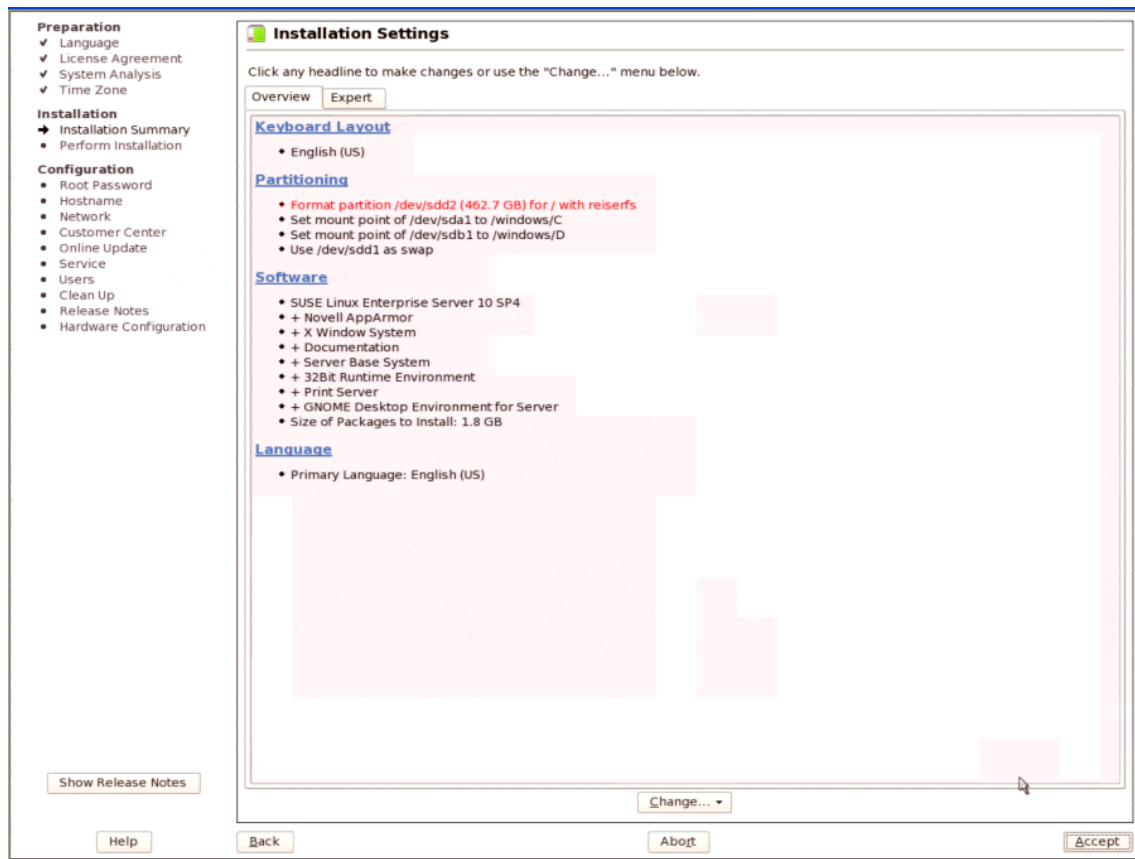


8. Click **Back** to proceed.

The SLES OS switches to the installation GUI.

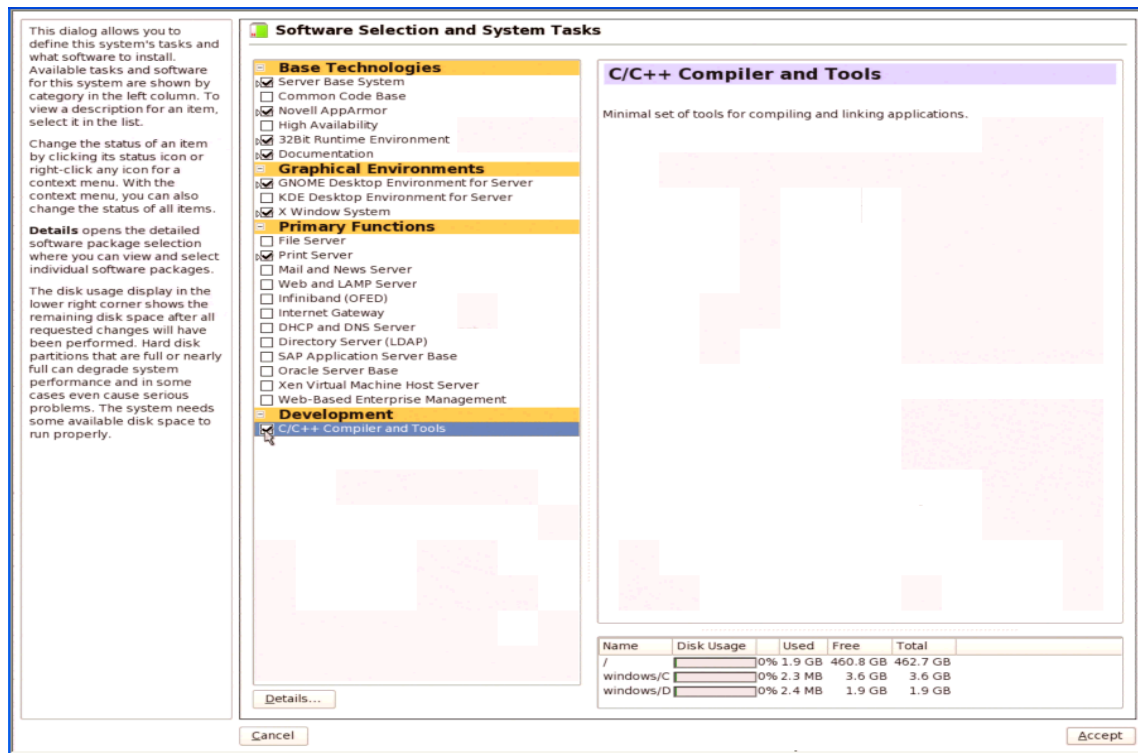
9. Accept the Licensing agreement and choose your language.
10. Click **<button>** to proceed.
11. Choose the installation media through which to start the OS installation.
In this case, you are using a USB-DVD-ROM.
12. Choose the desired installation mode.
Select **New Installation** to begin a fresh OS installation.
13. Select the **Software** option to add the development kit.
Before going to next screen, press CTRL+ALT+F2 to check the command prompt for `lsmod | grep megasr`.
This command indicates whether the Software RAID driver has loaded or not.

Figure 16 Installation Settings Screen



14. Choose the desired software to load.

Figure 17 Software Selection and System Tasks Screen



15. Click **Accept** to proceed with the installation.

2.2.9.2 Installing the SLES-10, SP-4 XEN Drivers

Use the following instructions for a XEN environment installation of the SLES-10, SP4 OS.

1. Prepare a USB-Drive with MegaSR driver.

This step is mandatory because the DUD image for the SLES OS exceeds floppy size.

- Need two USB drives (let USB1 and USB2).
- Make sure that USB drives are formatted before using for OS installation using DUD images.
 - a. Copy the DUD image file to any Linux system.

Ex: `megasr-15.00.0405.2012-1-sles10-sp4-x86_64.img` for SLES10-SP4-64bit

b. Create an image directory.

Ex: `mkdir image`

c. Mount the DUD image on to image directory using the following command.

Syntax: `mount -o loop <source> <destination_folder>`

Ex: `mount -o loop megasr-15.00.0405.2012-1-sles10-sp4-x86_64.img image`

d. Change the directory to image.

Ex: `cd image`

Two folders are generated (01 and 02).

e. Copy both the 01 and 02 folders onto the USB1 drive.

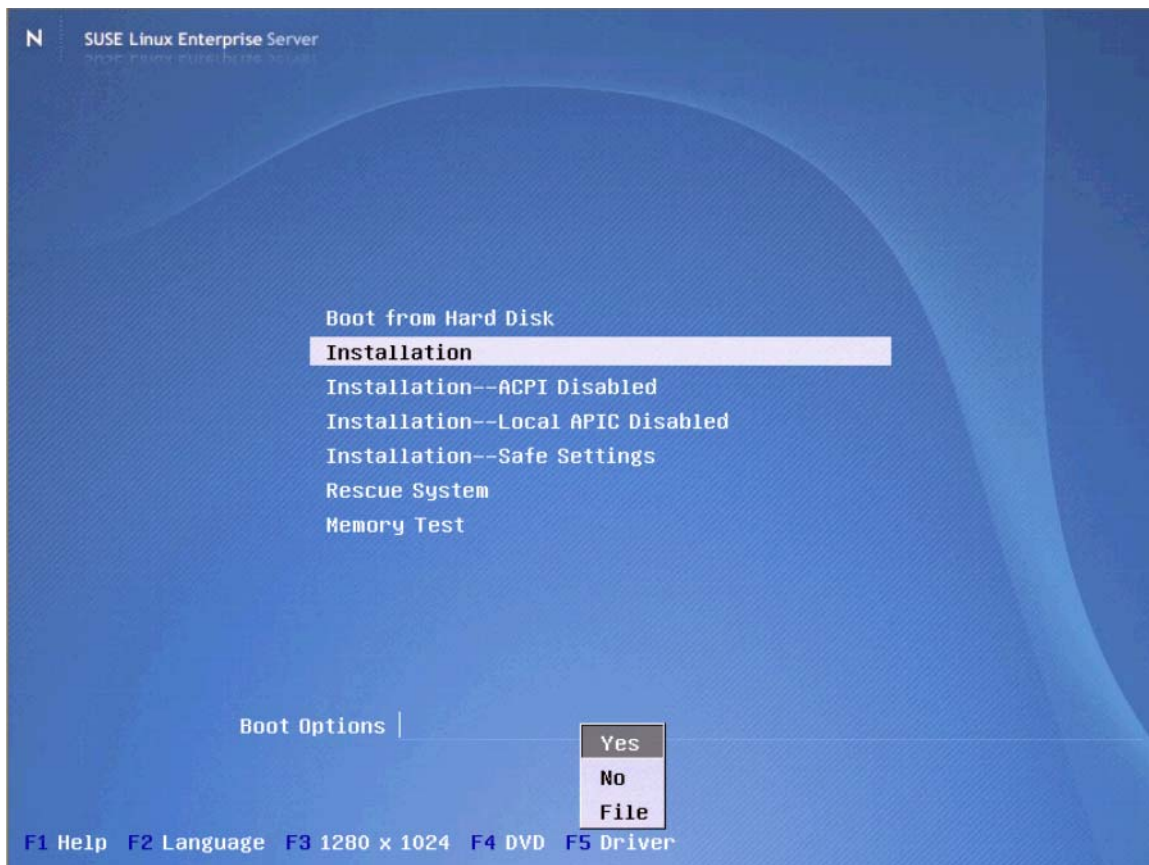
Verify whether the 02 folder contains `megasr.ko` file.

f. Copy the DUD image file to any Linux system.

Ex: `megasr-15.00.0405.2012-1-sles10-ext-sp4-x86_64.img` for SLES10-SP4-64bit

- g. Create an image directory.
Ex: `mkdir xen_image`
 - h. Mount the DUD image onto the image directory using the following command.
Syntax: `mount -o loop <source> <destination_folder>`
Ex: `mount -o loop megasr-15.00.0405.2012-1-sles10-ext-sp4-x86_64.img xen_image`
 - i. Change the directory to the image directory.
Ex: `cd xen_image`
Two folders are generated (01 and 02).
 - j. Change the directory to `02\linux\suse\x86_64-sles10\`
Ex: `cd 02\linux\suse\x86_64-sles10\`
The directory contains a file called `update.tar.gz`.
 - k. Use the `tar -xvzf update.tar.gz` command to extract the files from the GZIP file.
 - l. Go back to the `xen_image` directory.
 - m. Copy both the 01 and 02 folders onto the USB2 drive.
 - n. Insert both the USB drives (USB1 and USB2) into the machine and power ON the machine.
2. Configure the BIOS and OpROM settings.
 3. Connect the USB drive the contains the MegaSR driver and reboot the system.
The following window appears after booting from the DVD.
 4. Choose the Installation option.

Figure 18 Installation Screen – Driver Loading Screen



5. Press F5 and click **Yes** to load the driver.

NOTE Avago supports AHCI controllers on select hardware.

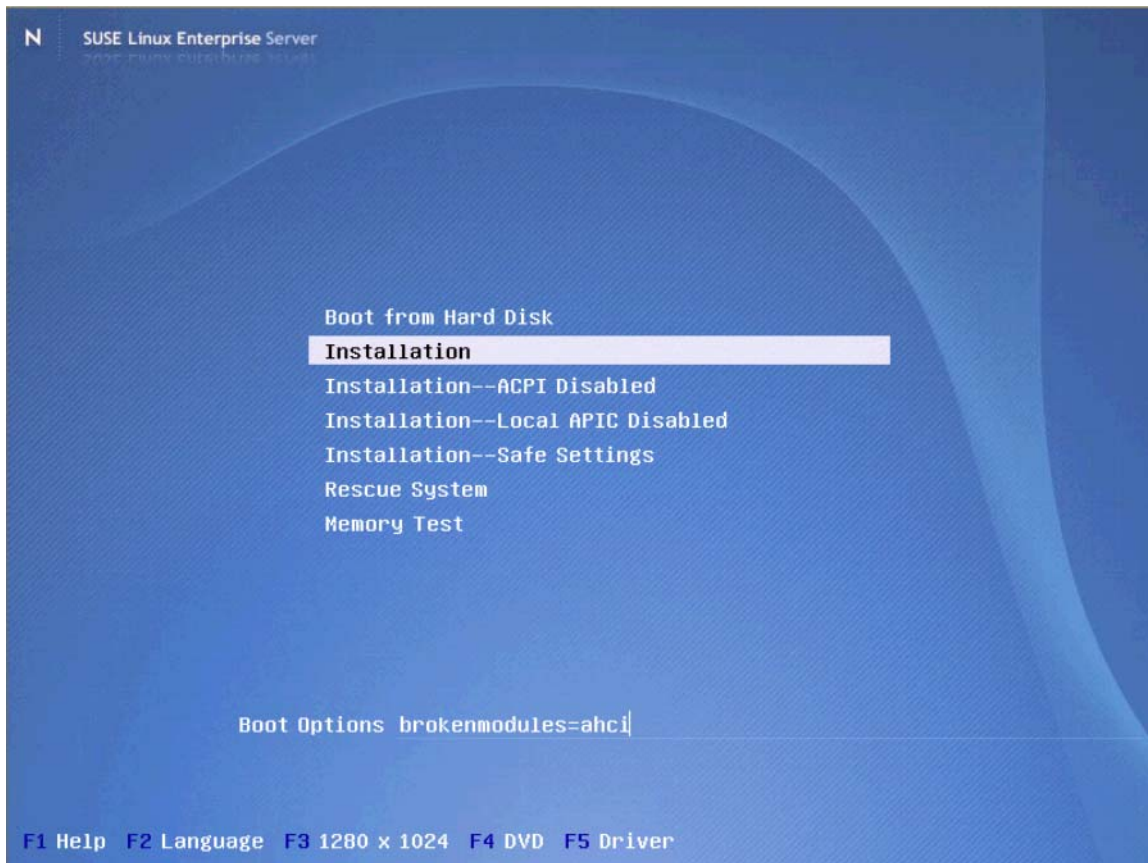
6. Provide one of the following commands as the boot argument.

```
brokenmodules=isci  
brokenmodules=ahci
```

These arguments ensure that either the SCU controller or AHCI controller is claimed by the MegaSR driver, respectively. For the MegaSR driver to claim both the SCU and AHCI controllers, use the following command.


```
brokenmodules=isci,ahci
```

Figure 19 Installation Screen – Dual Controller Mode



The following screen appears.

Figure 20 Driver Updates Added Confirmation Screen

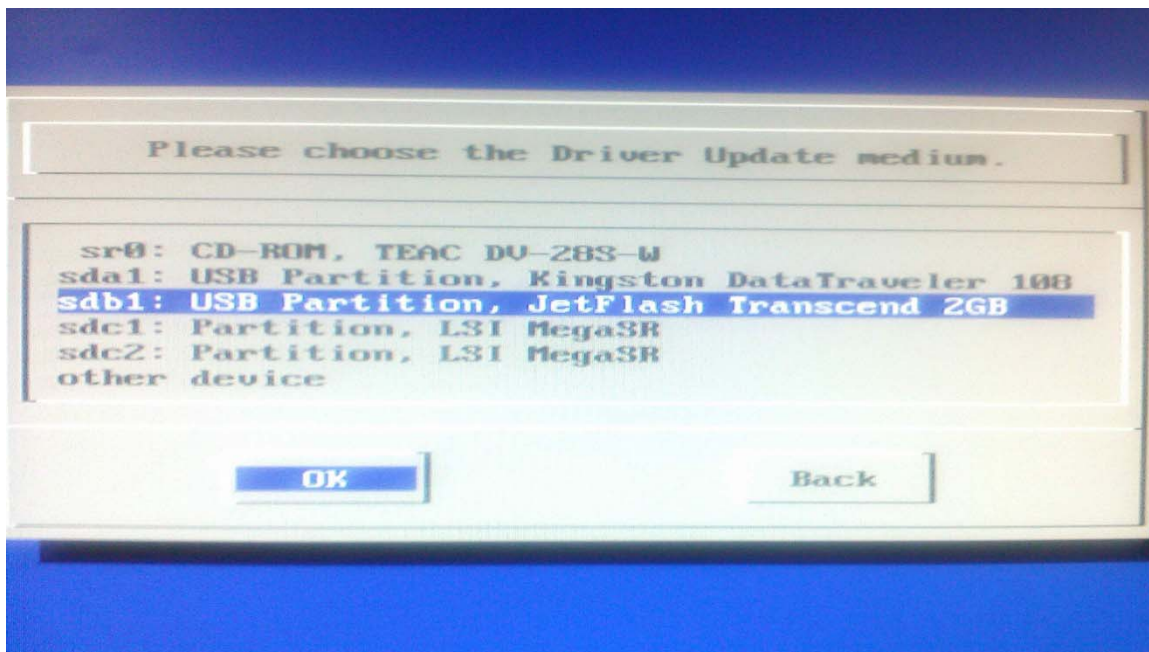


This screen shows the MegaSR driver the is in the pen drive is recognized and loaded.

7. Click **OK** to proceed.

The following screen appears.

Figure 21 Driver Update Medium Screen

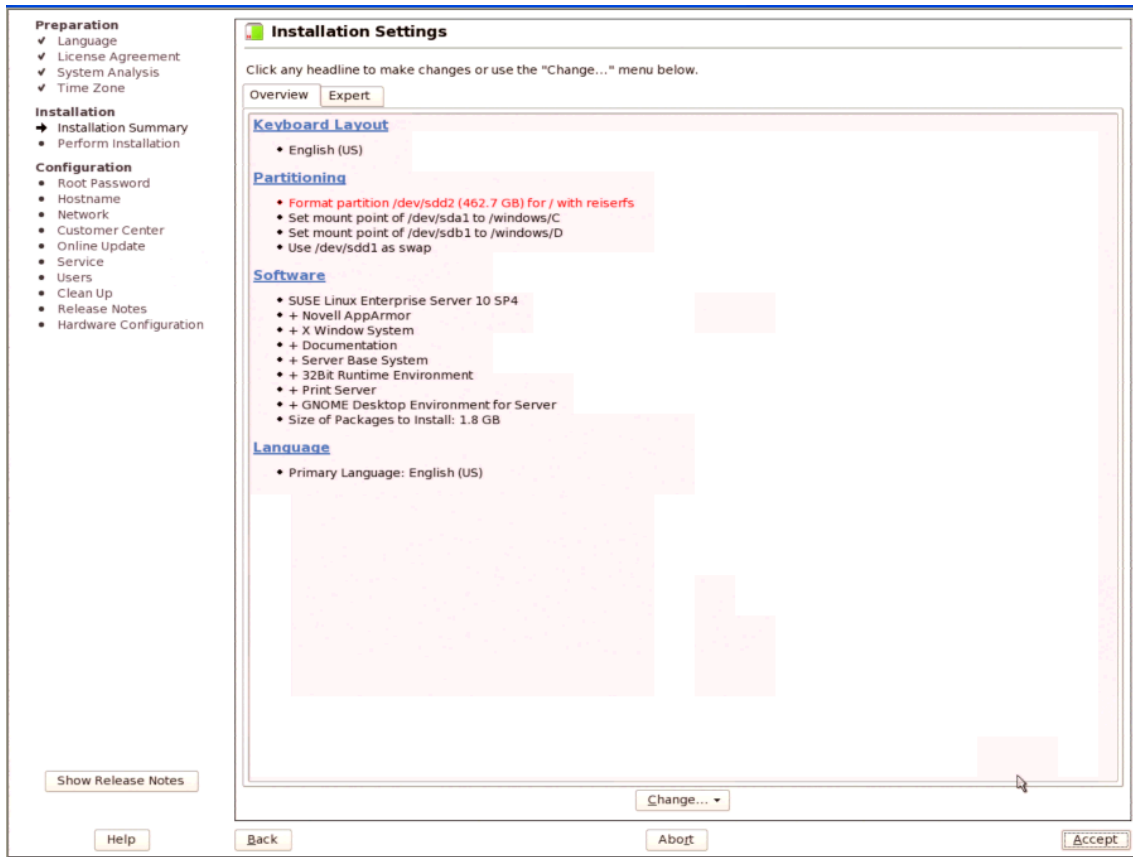


8. Click **Back** to proceed.
The SLES operating system switches to the GUI installation.
9. Accept the Licensing agreement and choose your language, and proceed.

Choose the installation media to use to start your OS installation. In this case, you are using a USB-DVD-ROM.


1. Choose the desired installation mode.
Select **New Installation** to start a fresh installation. The following screen appears.

Figure 22 Installation Settings Screen – Partitioning



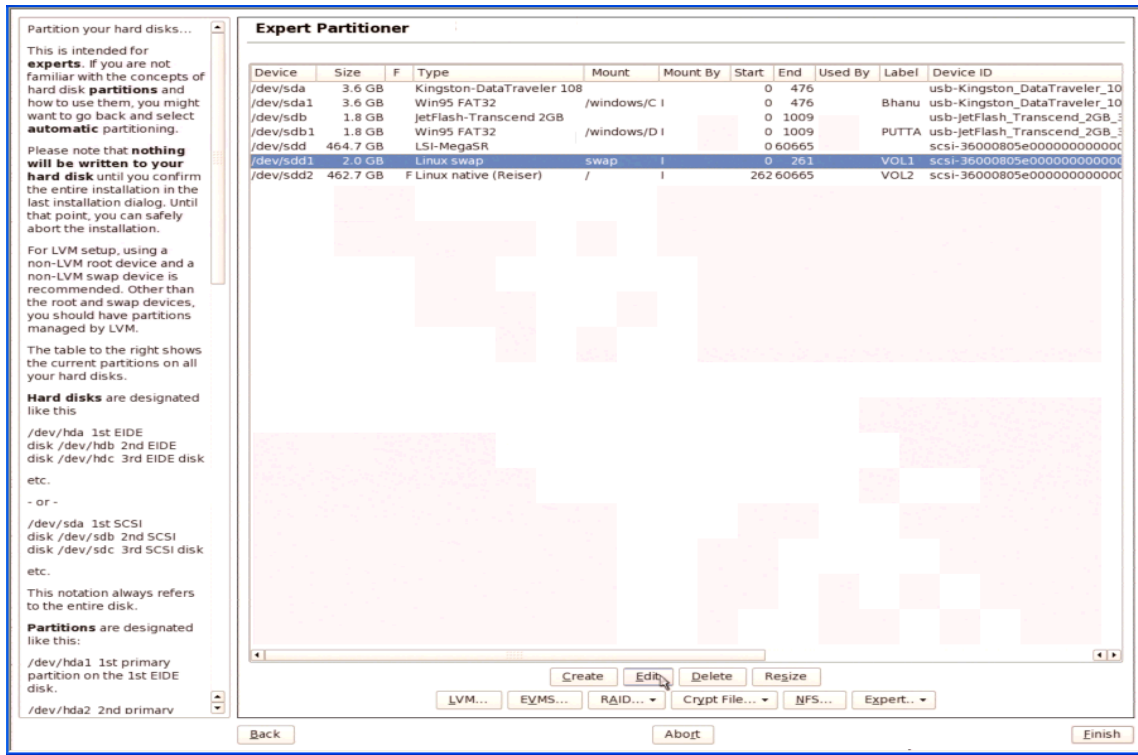
2. Select the **Partitioning** option and use the following procedure.
 - a. Select the **Base partition setup on this proposal** option and click **Next** to proceed.

Figure 23 Suggested Partitioning Screen




- b. Change the **Fstab** options for each partition on the /dev/sdb1 directory.

Figure 24 Expert Partitioner Screen



- c. Click **Edit**.
The following screen appears.


Figure 25 Expert Partitioner Screen – Edit Existing Partition Window



d. Click **Fstab Options**.

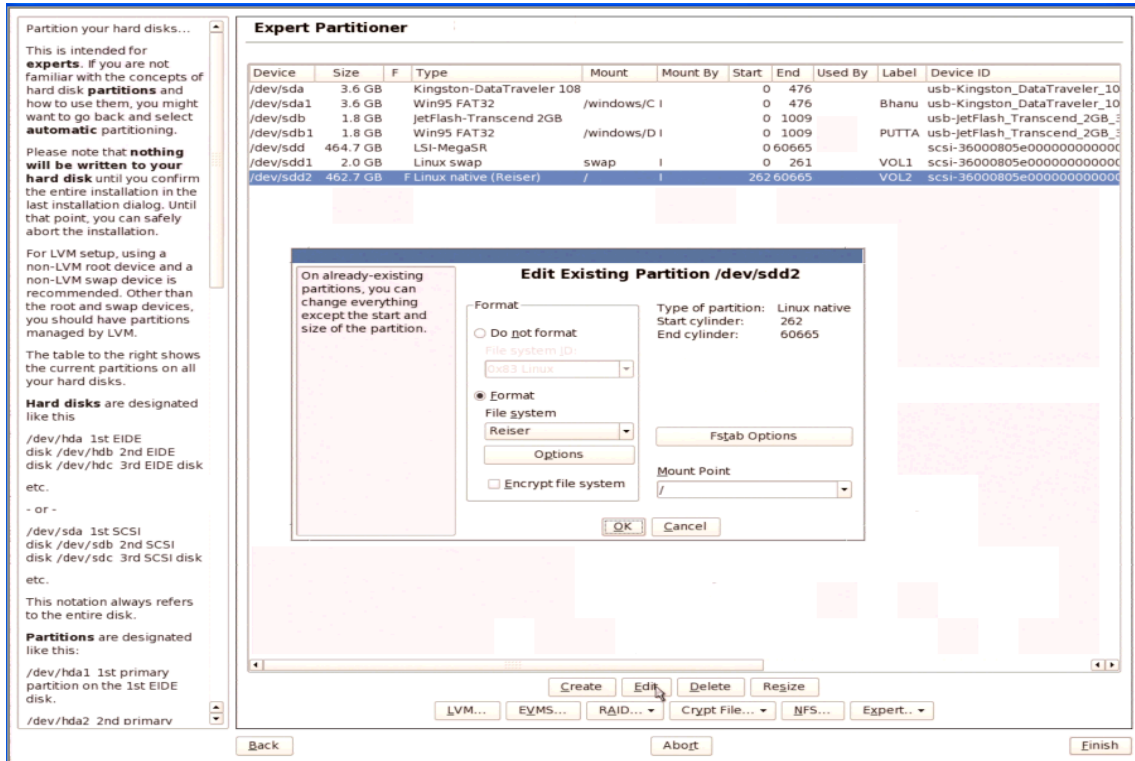
The following screen appears.

Figure 26 Expert Partitioner Screen – Fstab Options Window



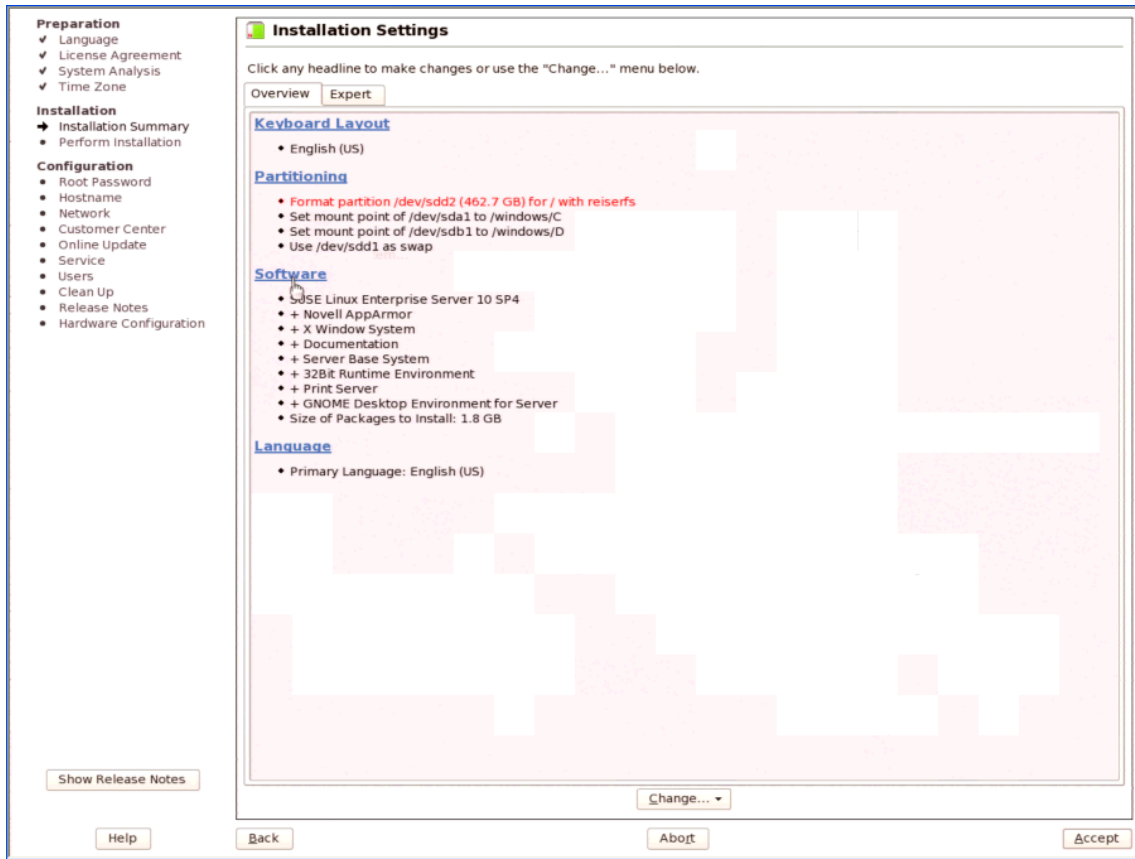
- e. Select a volume label for the **Mount in /etc/fstab by** pane, name the Volume Label VOL1.
 - f. Click **OK** to proceed.
- The following screen appears.

Figure 27 Expert Partitioner Screen – Edit Existing Partition Window



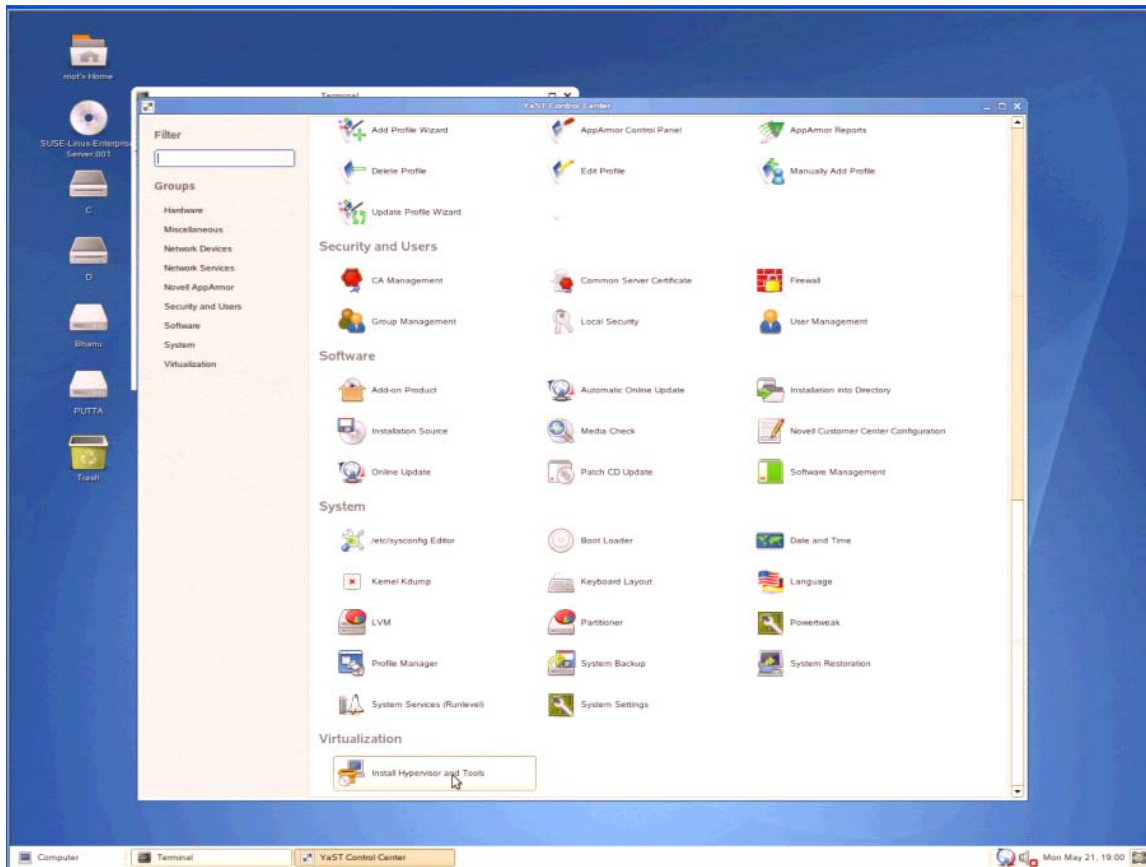
- g. Enter a name for the volume into the **Volume Label** field, for example VOL2, VOL3, and so on.
 - h. Click **Finish** to proceed when you finish naming of all the volumes.
3. Select the **Software** option if you want to install a development kit.

Figure 28 Installation Settings Screen – Software Option



- a. Choose the desired options.
- b. Proceed with the normal installation.
- c. When the machine is up and running, go to the terminal and type `yast2`.
The following screen appears.

Figure 29 YaST Control Center Screen




- d. Select the **Install Hypervisor and Tool** option.

NOTE Before you select this option make sure that the CD-DVD-ROM contains the SLES-10-SP4 OS disk.

The OS loads the selected packages from the OS disk.


- e. Click **Continue** when the desired packages finish loading.

Figure 30 Configuring the Virtual Machine Server Screen



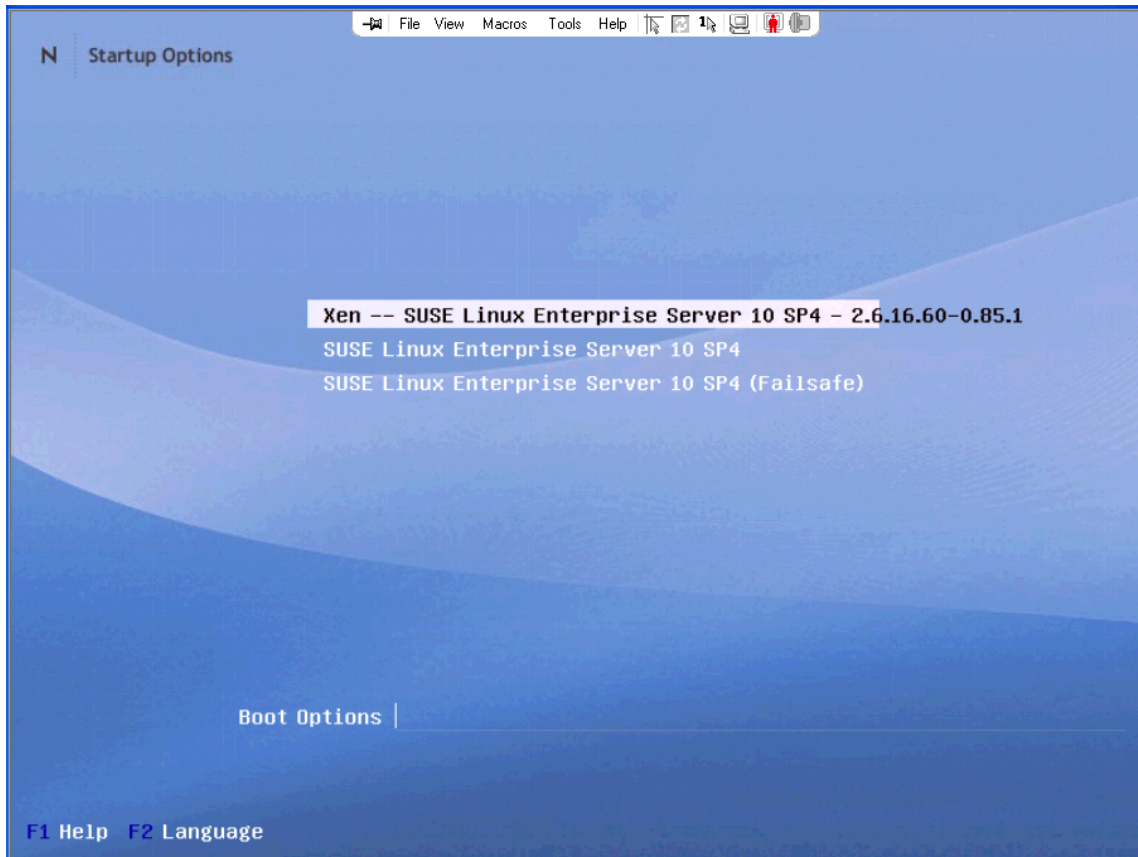
- f. Click **OK** when the reboot window appears.

Figure 31 Reboot Request Window



4. Select the **Boot from Disk** option on the SLES home page.
5. Select the **Xen-SUSE Linux Enterprise server 10 SP4 -2.6.16.60-0.85.1** option.

Figure 32 Startup Options Screen



The Xen Kernel is loaded but Graphical User Interface (GUI) is not yet configured for this Xen Kernel. Perform the following steps to configure the kernel.

1. Login as `root` (super user).
If you do not log in as the root user, you might not have all privileges needed to install some of the software packages.

Figure 33 SuSE Linux SLES Login Screen

```
Welcome to SUSE Linux Enterprise Server 10 SP4 (x86_64) - Kernel 2.6.16.60-0.85
.1-xen (tty2).

linux-1m1e login: root
Password:
Login incorrect

linux-1m1e login: root
Password:
Last login: Mon May 21 18:56:24 EDT 2012 on :0
You have new mail.
linux-1m1e:~ # _
```

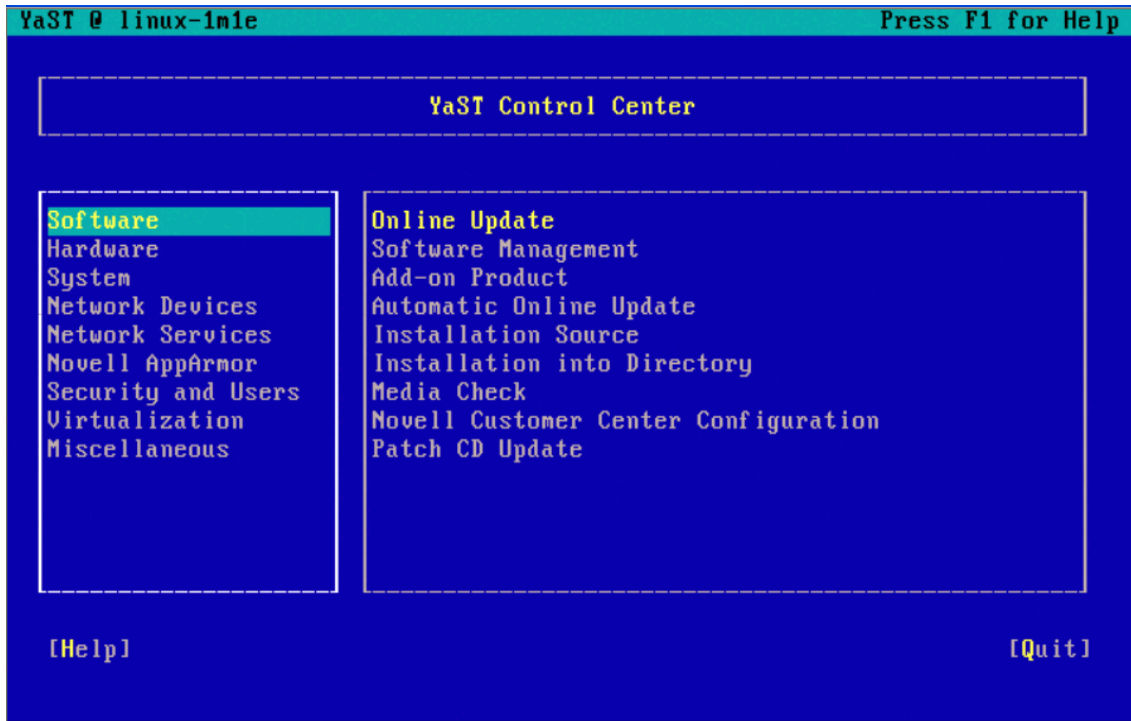
2. Type `yast2` at the terminal prompt.

Figure 34 Terminal Prompt

```
linux-1m1e:~ #
linux-1m1e:~ #
linux-1m1e:~ # yast2
```

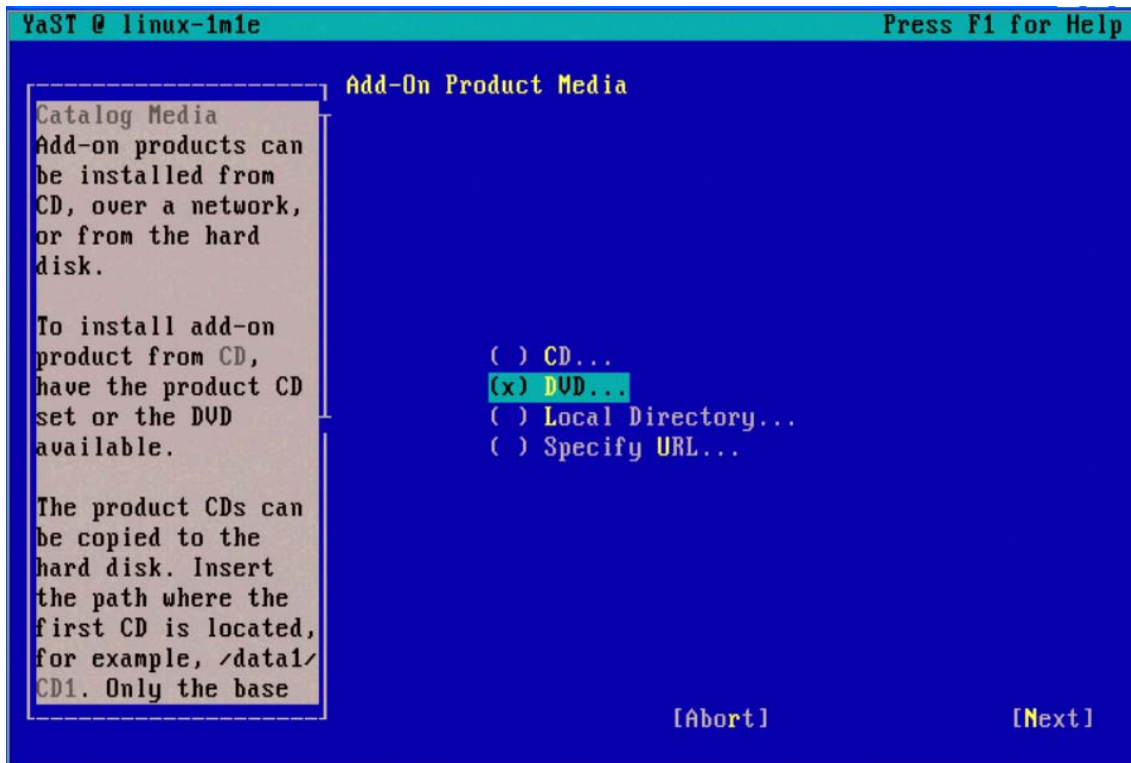
3. Select **Software > Add-on Products** (the OS disk should still be in the CD-DVD-ROM drive).
Use the Tab key in the key board to navigate through the options.

Figure 35 YaST Control Center – Software Online Update Screen



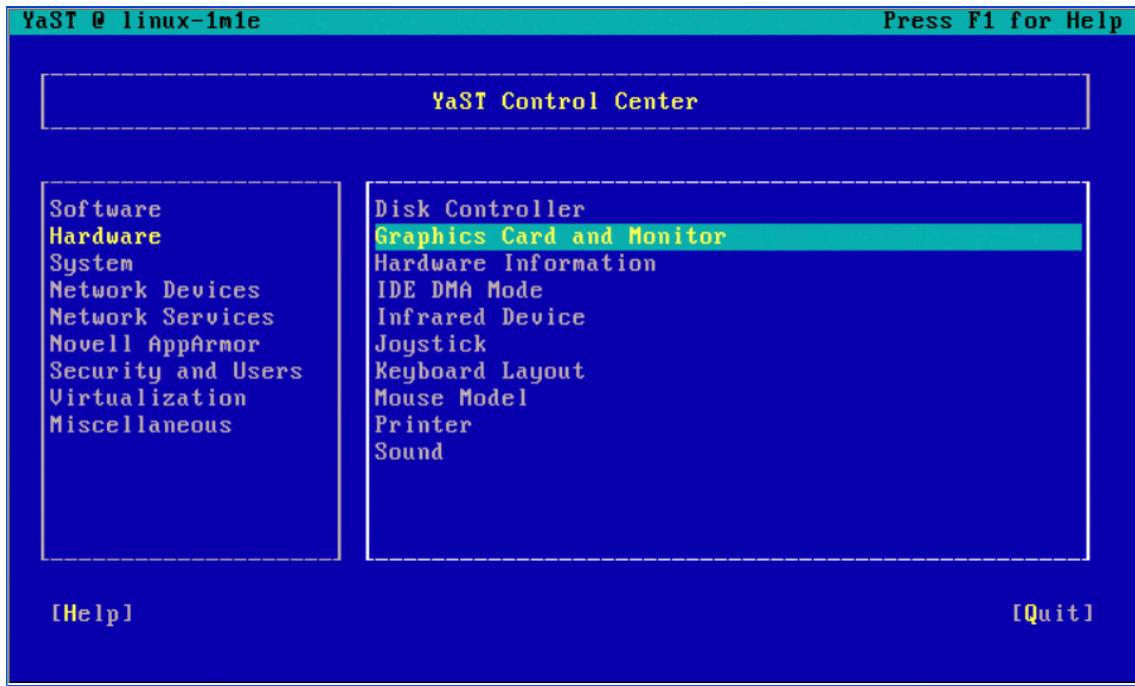
4. Choose the product medium and click **Next**.

Figure 36 Add-On Product Media Screen



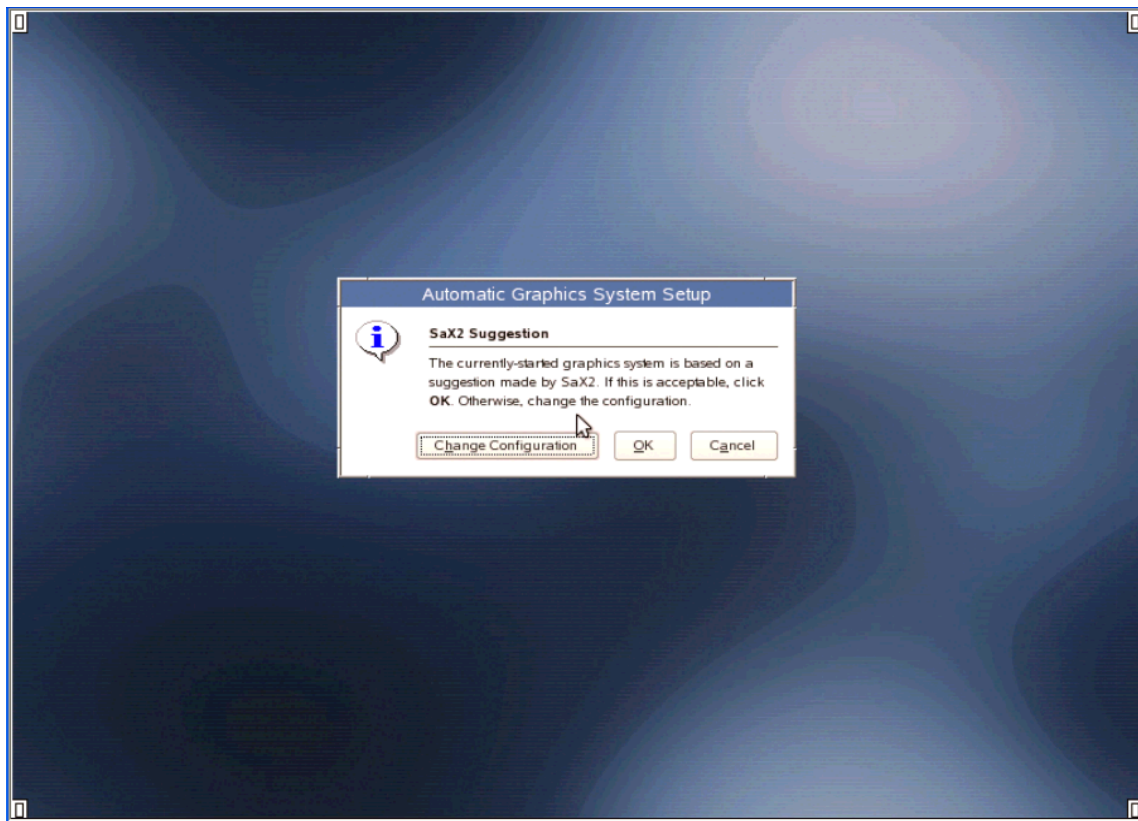
5. Select the **Hardware > Graphics Card and Monitor** option.

Figure 37 YaST Control Center Screen – Graphics Card and Monitor



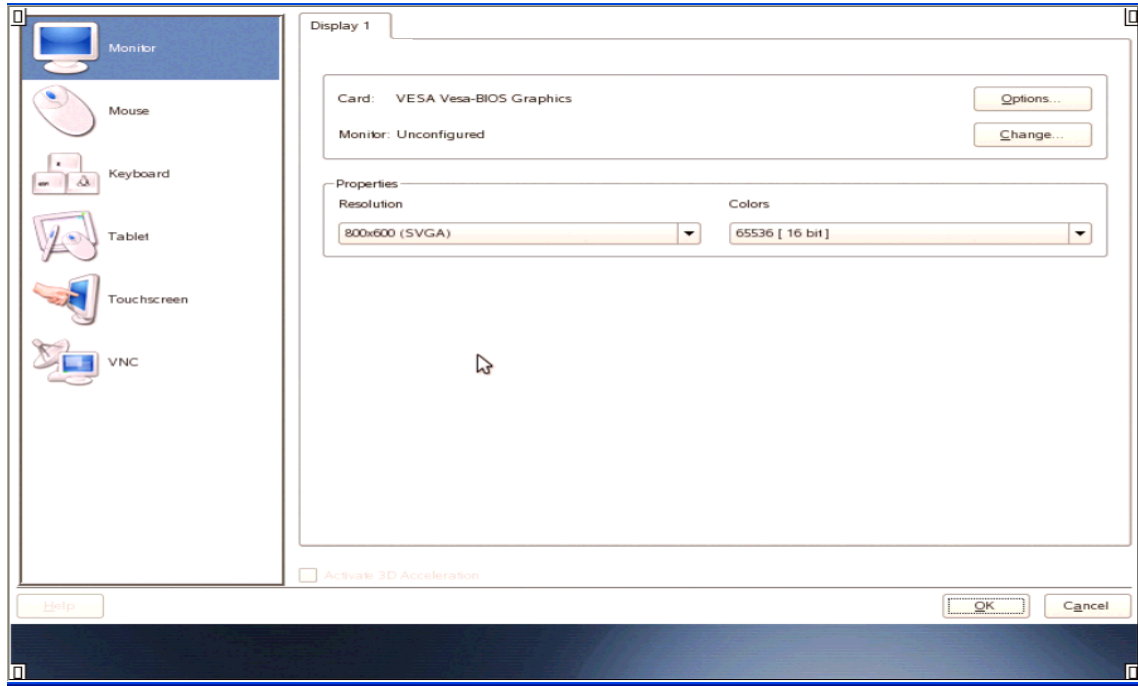
6. Select the **Change Configuration** option.

Figure 38 Automatic Graphics System Setup Window



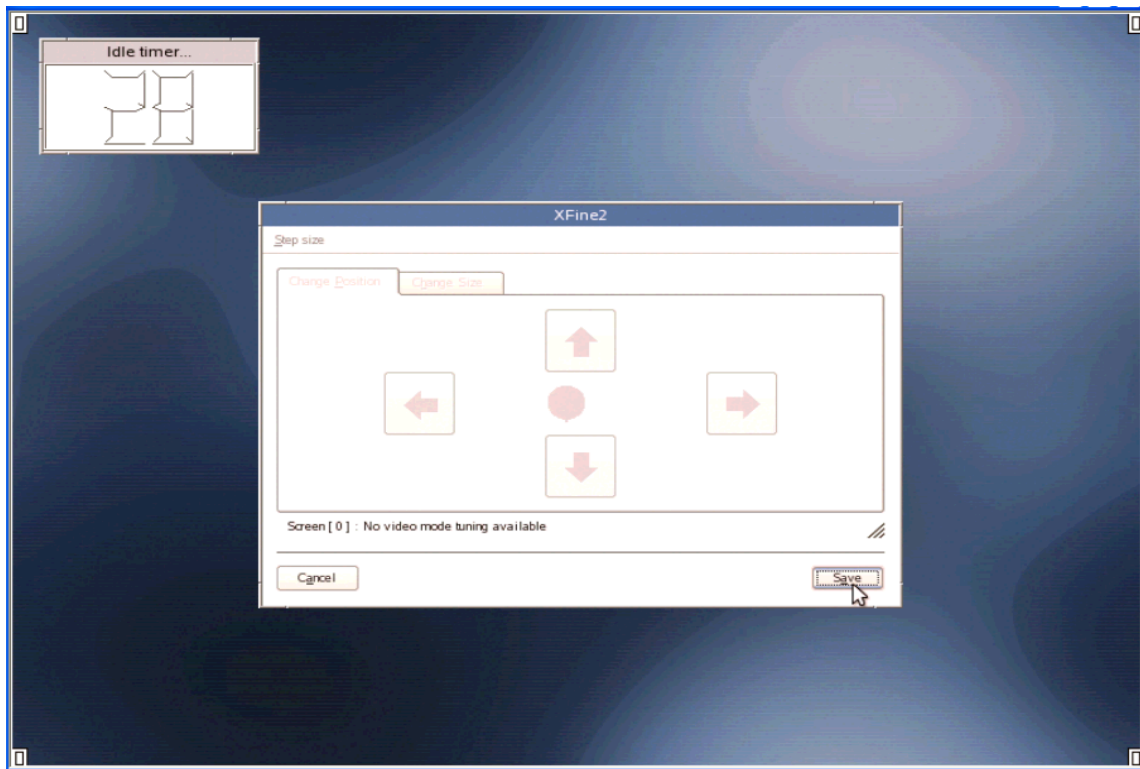
7. Click **OK**.
The following screen appears.

Figure 39 Monitor Display Settings Screen



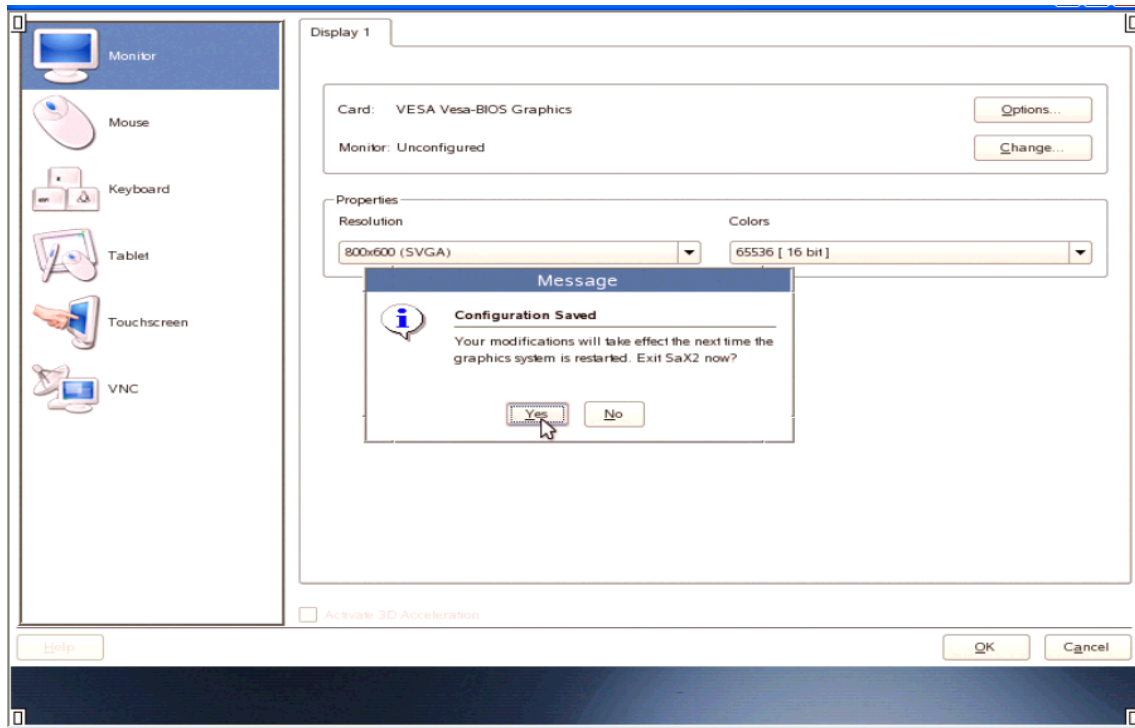
8. Save the configuration.

Figure 40 Save Configuration Screen




The following window appears.

Figure 41 Configuration Saved Confirmation Window



9. Click **Yes** to clear the message.
 10. Click **OK** to continue.
- The following screen appears.

Figure 42 Terminal Screen



2.2.10 Installing the SuSE Linux Enterprise Server 11 and 12 Driver

Use the instructions in this section to install the SLES-11-GA/SP1/SP2/SP3/SP4 and SLES 12 operating systems. The installation method is the same for both 32-bit and 64-bit installations. The EFI mode installation is applicable only to 64-bit operating systems.

1. Prepare a driver disk using a USB drive.
You cannot use a floppy disk because of the size of the MegaSR DUD image.
 - a. Copy the DUD image file to any Linux system.
Ex: `megasr-14.00.0722.2010-1-sles11-sp1-x86_64.img` for SLES11-SP1-64bit
 - b. Create a directory and mount the DUD image to that directory using the following command.
Syntax: `mount -o loop <source> <destination_folder>`
Ex: `mkdir image`
`mount -o loop megasr-14.00.0722.2010-1-sles11-sp1- x86_64.img image`
 - c. Copy the contents of the `image` directory to a USB drive.
Use a formatted USB disk to create the driver disk.
2. Configure the BIOS and OpROM settings.
3. Connect the driver disk and reboot the system.
4. Choose Installation option after the system has restarted.
The following screen appears.

Figure 43 SuSE Installation Screen



5. Press F6 and click **Yes** to load the driver.

NOTE This option appears only if the installation is performed using non EFI mode.

To install the operating system using the EFI mode, perform the following steps.

1. Press Tab when the ELILO boot loader begins loading.
For the SLES 11, SP2 and earlier OS image, the following Window appears.


Figure 44 ELILO Prompt Screen



Type `linux brokenmodules=iscsi` at the ELILO prompt and press Enter.


For the SLES 11, SP3 operating system, the following window appears.

Figure 45 SuSE Linux Enterprise Server 11 Installation Screen



- a. Press E to edit the commands before booting the system.
- b. Append the `linuxefi` line with `brokenmodules=ahci,iscsi` option.
Ex: `linuxefi /boot/x86_64/loader/linux brokenmodules=ahci`
- c. Press Enter to boot the selected operating system.
The following screen appears.

Figure 46 SuSE Linux Enterprise Server Installation Screen



2. Provide `brokenmodules=isci` as a boot argument.

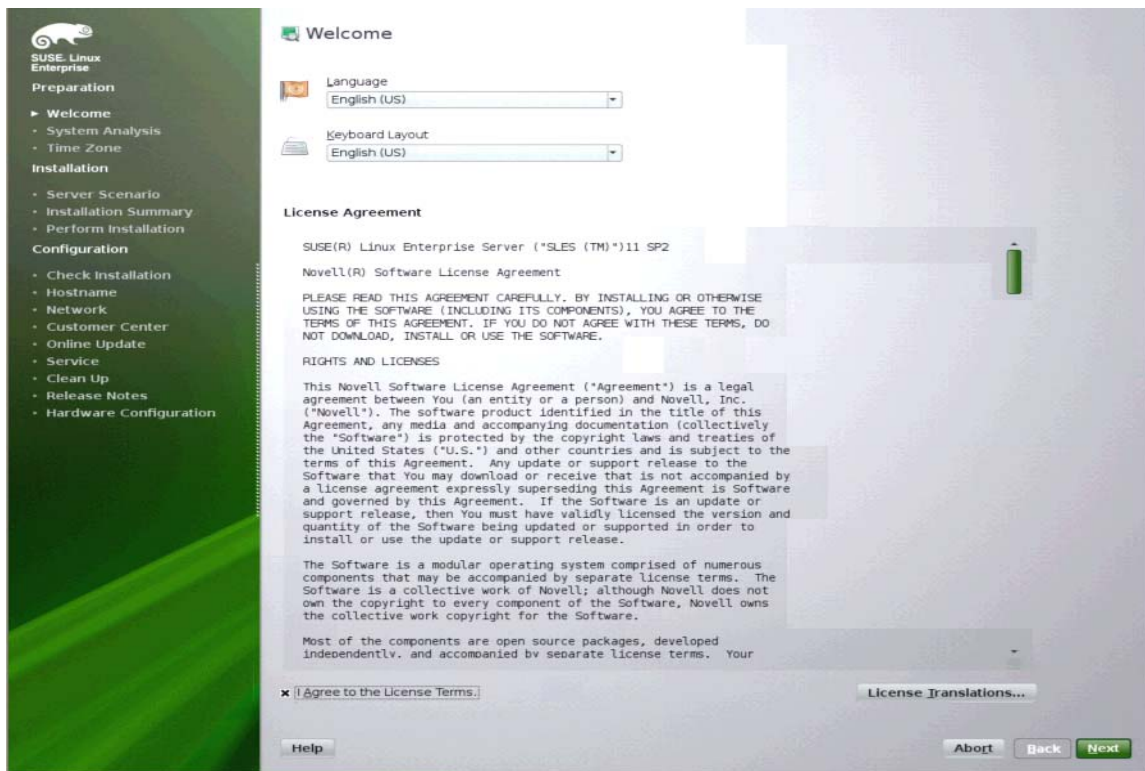
This argument ensures that the SCU controller is claimed by the MegaSR driver.

Avago supports AHCI controllers on select hardware. On that hardware, for the MegaSR driver to claim both SCU and AHCI controllers, use the following command.

```
brokenmodules=isci,ahci
```

The following screen appears after the SLES switches to the installation GUI mode.

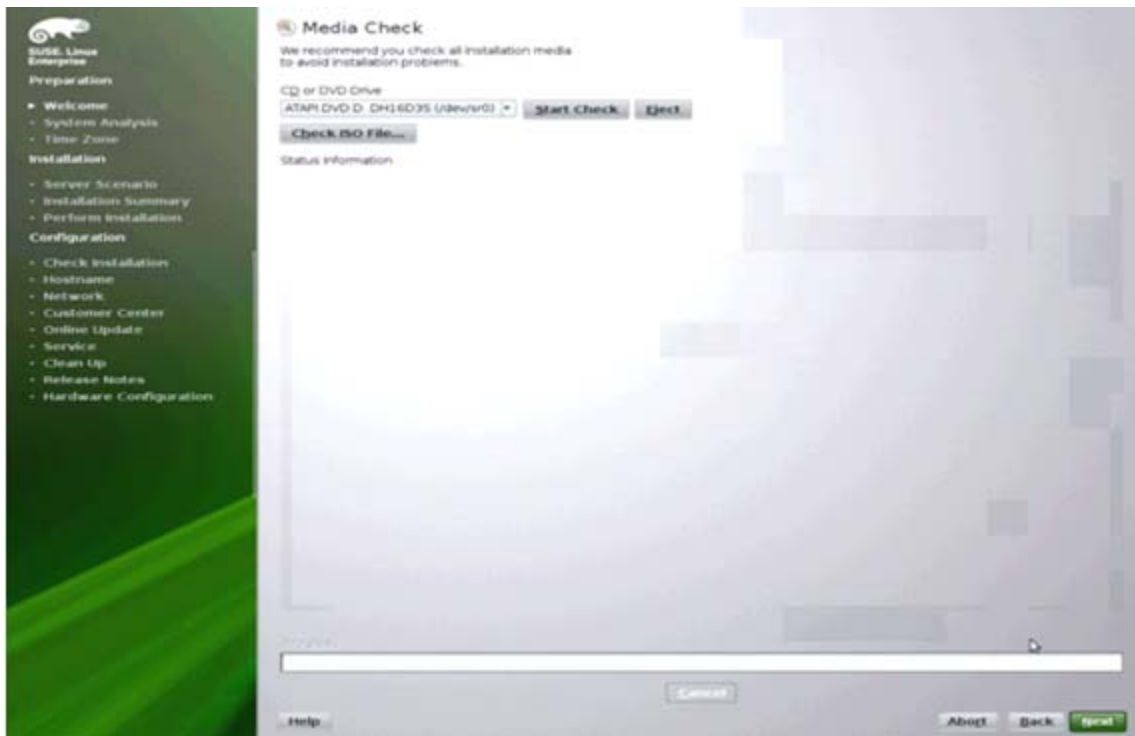
Figure 47 SuSE Linux Enterprise Preparation – Welcome Screen



3. Accept the Licensing agreement and click **Next** to proceed.

The following screen appears.


Figure 48 SuSE Linux Enterprise Preparation – Media Check Screen



This screen shows from which installation media to start the OS installation. In this case, navigate to the appropriate USB-DVD-ROM.

4. Click **Next** to proceed.
The following screen appears.


Figure 49 SuSE Linux Enterprise Preparation – Installation Mode Screen



5. Select the desired installation mode and click **Next** to proceed.

NOTE Select **New Installation** when performing a fresh installation.
The following screen appears.

Figure 50 SuSE Linux Enterprise Installation – Server Base Scenario Screen

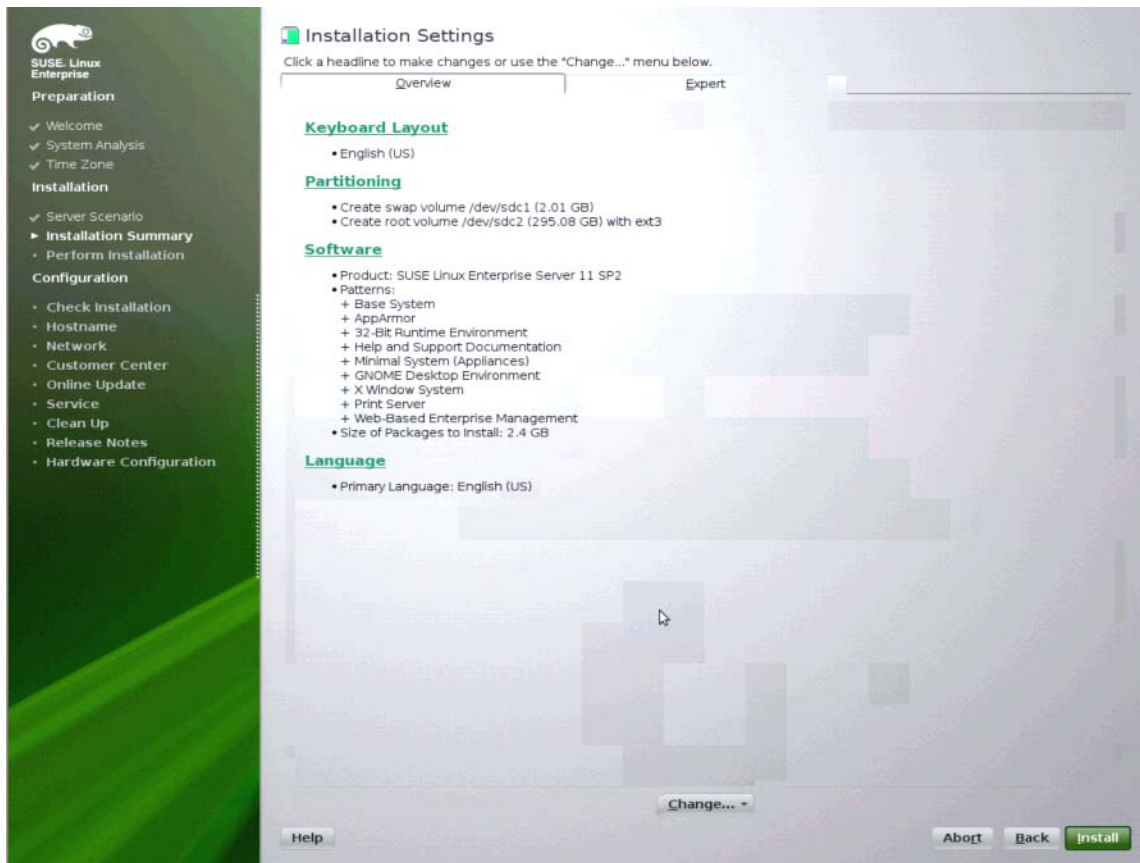


Three options are available during the installation.

- Physical Machine, which is also available for fully virtualized guests
- Virtual Machine, which is for para-virtualized environments like Xen
- Xen Virtualization Host, if virtual machines must be created


6. Select the installation option you want to use for the installation and click **Next** to continue. The following screen appears.

Figure 51 SuSE Linux Enterprise Installation – Settings Summary



7. Select the **Partitioning** option if you need to change the partition parameters.
It is recommended that you leave this option unchanged, if you are not clear on what you are trying to achieve.
For more details on partitioning, refer to the SLES/RHEL installation guide.
8. Select the **Software** option to install additional software packages.
The following screen appears.

Figure 52 Software Selection and System Tasks Screen



If you require additional software packages, check them and click **OK** to continue.

9. Reboot the system from the HDD after completing the installation.

If a Xen kernel is required, follow the instruction provided in the following section.

2.2.10.1 Loading the Xen Kernel RPM on SLES-11

NOTE The Xen kernel does not work if the installation is performed using EFI mode.

NOTE The Xen kernel is not supported in a 32-bit OS. This is a SLES limitation.

1. Install the DKMS package before installing the `megasr_xen_rpm` package.

NOTE For SLES-11, SP1, the latest DKMS package does not work. Use the `DKMS-2.0.5-1.noarch` package. For SLES-11, SP2 use the latest DKMS package.

2. Extract the `megasr_xen_rpm` package from the DUD image.
3. Follow Step 1 in [Section 2.2.10, Installing the SuSE Linux Enterprise Server 11 and 12 Driver](#) to extract the RPM driver.

You can distinguish the Xen kernel DUD image using the `ext` label mentioned with its name.

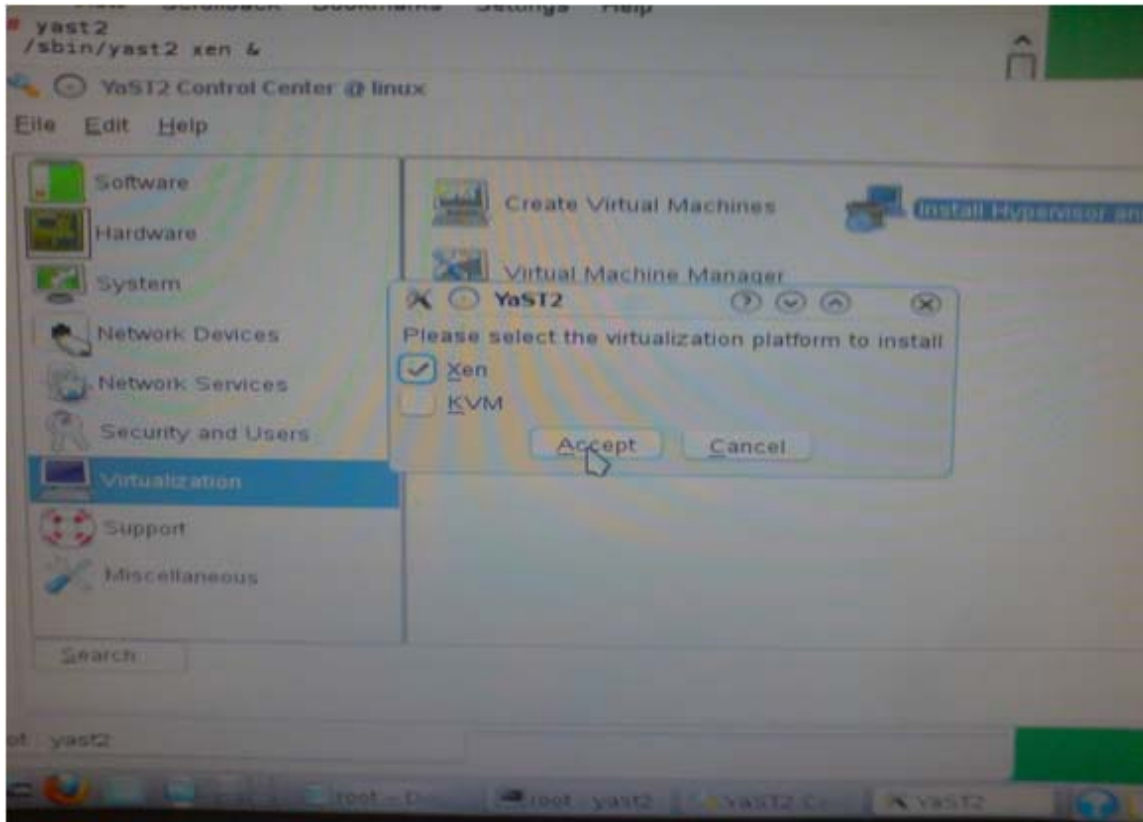
Ex: `megasr-14.00.0722.2010-1-sles11-ext-sp1-x86_64.img`

4. Install the Xen kernel RPM driver.
5. Reboot the system and type `yast2` in the terminal window.

6. Select **Virtualization** and **Install Hypervisor and Tools**.

The following screen appears.

Figure 53 YaST2 Installation Window



7. Choose the Xen kernel on next reboot to boot into the SLES Xen kernel.

2.2.10.2 Updating or Loading Software RAID Driver Kernel RPM on SuSE 12

Perform the following steps to update or load the already installed Software RAID Driver module on SLES 12 Operating Systems.

1. Create a directory and mount the DUD image to that directory using the following command:

```
mkdir /lib/modules/<new-ver>/updates
```

2. Copy the module </lib/modules/<old kver>/updates/megasr.ko> file to the newly created directory </lib/modules/<new kver>/updates directory>

3. Run the following command:

```
depmod <kver>
```

4. Add the <dracut config file> to the module

If you have not created the Dracut Config file yet, you must create one using the following commands:

- a. create /etc/dracut.d.conf/10-megasr.conf
- b. #touch /etc/dracut.d.conf/10-megasr.conf
- c. #vi /etc/dracut.d.conf/10-megasr.conf

- d. `#cd /etc/dracut.conf.d/`
 - e. `/etc/dracut.conf.d# ls`
 - f. `02-early-microcode.conf 10-megasr.conf 99-debug.conf`
 - g. `/etc/dracut.conf.d # cat 10-megasr.conf`
 - h. `add_drivers+="megasr"`
5. Run the `<dracut --force --kver <kernel version>` command.

2.2.11 Installing the SLES Operating System on Grantley Servers

This guideline can be used for the installation of SLES-11-GA/SP1/SP2/SP3/SP4 and SLES 12 OSs on Grantley servers. This Installation procedure explains how the MegaSR driver claims the SATA controller and the AHCI driver claims the sSATA controller at the same time. The installation method is common for both 32-bit and 64-bit installations. The EFI mode of installation is applicable only to 64-bit OS.

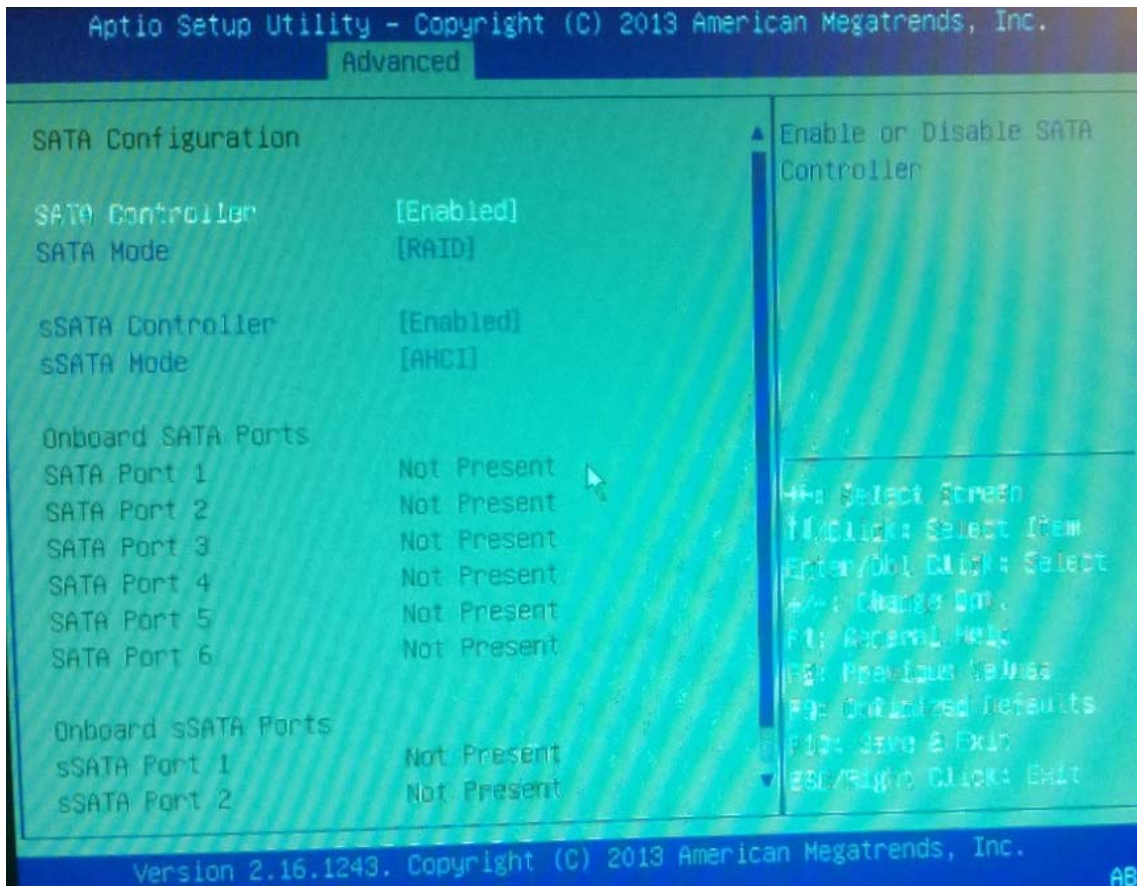
Existing and earlier servers have only one onboard AHCI controller that can be allowed to configure either in RAID mode or in AHCI mode.

The new Grantley Servers have two onboard AHCI controllers (SATA and sSATA) and has the provision to configure each controller to either RAID mode or AHCI mode. At the same time, one controller can be in RAID mode and other in AHCI mode. Both in AHCI mode and in RAID mode this is also possible.

2.2.11.1 Operating System Installation in RAID Mode

1. Prepare a driver disk using USB disk.
You cannot use a floppy disk because of the size of the MegaSR DUD image.
 - a. Copy the DUD image file to any Linux system.
Ex: `megasr-14.00.0722.2010-1-sles11-sp1-x86_64.img` for SLES11-SP1-64bit
 - b. Create a directory and mount the DUD image to that directory using the following command.
Syntax: `mount -o loop <source> <destination folder>`
Ex: `mkdir image`
`mount -o loop megasr-14.00.0722.2010-1-sles11-sp1- x86_64.img image`
 - c. Copy the contents of the sub-directories named 01 and 02 in the image directory to a USB drive.
Use a formatted USB disk to make the driver disk.
 - d. Connect a PD to the SATA controller that is configured in RAID mode.
 - e. Create RAID 0.
2. Configure the BIOS and OpROM settings.

Figure 54 BIOS and OpROM Setting Screen



3. Connect the driver disk and reboot the system.
4. Choose the installation option.
5. Press F6 and click **Yes** to load the driver.

This option appears only if the installation is performed in non EFI mode.

Figure 55 Driver Installation



6. Type `brokenmodules=ahci` as the boot argument.

This argument ensures that the AHCI controller is claimed by the MegaSR driver.

To install the SLES 11, SP3 operating system using EFI mode, perform the following steps.

1. Press Tab when the ELILO boot loader begins loading.

Type `linux brokenmodules=iscsi` at the ELILO prompt and press Enter.

For SLES 11 SP3 the following window appears.


Figure 56 SuSE Linux Enterprise Server 11 Installation Screen



- a. Press E to edit the commands before booting the system.
- b. Append the `linuxefi` line with `brokenmodules=ahci,isci` option.
Eg: `linuxefi /boot/x86_64/loader/linux brokenmodules=ahci`
- c. Press Enter to boot the selected operating system.

The following window appears after the SLES switches to the GUI installation mode.

Figure 57 Grantley Server Welcome Screen



2. Accept the Licensing agreement and click **Next** to continue.

The following screen shows from which installation media to start the OS installation. In this case, a USB-DVD-ROM is used.

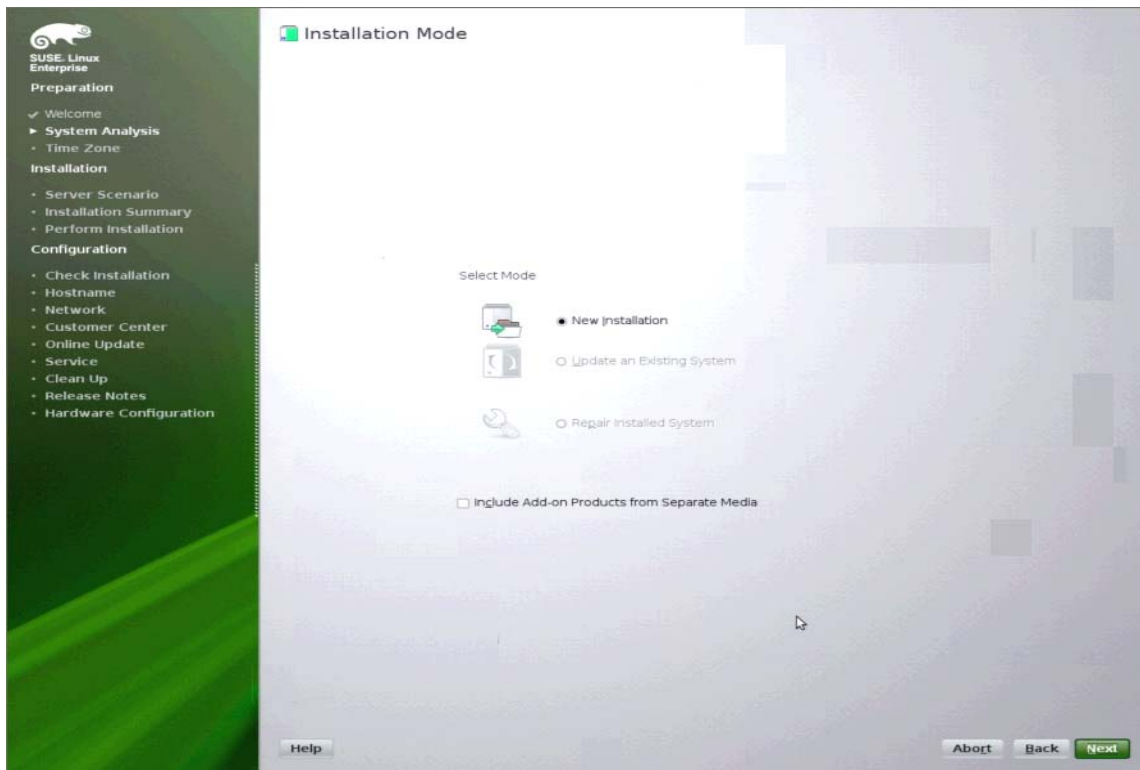
Figure 58 Grantley Media Check Screen



3. Click **Next** to continue.
4. Choose the desired installation mode and click **Next** to continue.

NOTE Select **New Installation** when performing a fresh installation.


Figure 59 Grantley Installation Mode Screen



Three options are available during the installation.

- Physical Machine, which is also available for fully virtualized guests
 - Virtual Machine, which is for para-virtualized environments like Xen
 - Xen Virtualization Host, if virtual machines must be created
5. Select the **Partitioning** option if you need to change the partition parameters.
It is recommended that you leave this option unchanged, if you are not clear on what you are trying to achieve.
For more details on partitioning, refer to the SLES/RHEL installation guide.
 6. Select the **Software** option to install additional software packages.
The following screen appears.

Figure 60 Software Selection and System Tasks Screen



Check any software packages you want to install and click **OK** to continue.

7. Reboot the system from the HDD after completing the installation.

The MegaSR driver now claims both the controllers. To claim the SATA controller by the MegaSR driver and the sSATA controller by the AHCI driver at the same time, perform the following steps.

1. Remove the `blacklist ahci` line from the `/etc/modprobe.d/blacklist` file.
2. Perform the following command to install the AHCI module.

```
modprobe -i ahci
```

3. Add the following statement to load the MegaSR driver in the `/etc/modprobe.conf` file, before loading the AHCI driver.

```
install ahci /sbin/modprobe megasr 2>&1 | : ;  
/sbin/modprobe -ignore-install ahci
```

The preceding statement is used to change the Device Driver load order.

4. Perform the following command to create an `initrd` image.

```
mkinitrd
```

5. Reboot the machine.

After rebooting the machine, the MegaSR and AHCI drivers should load successfully. The MegaSR driver should claim SATA controller (which is configured in RAID mode) and the AHCI driver should claim the sSATA controller (which is in AHCI mode).

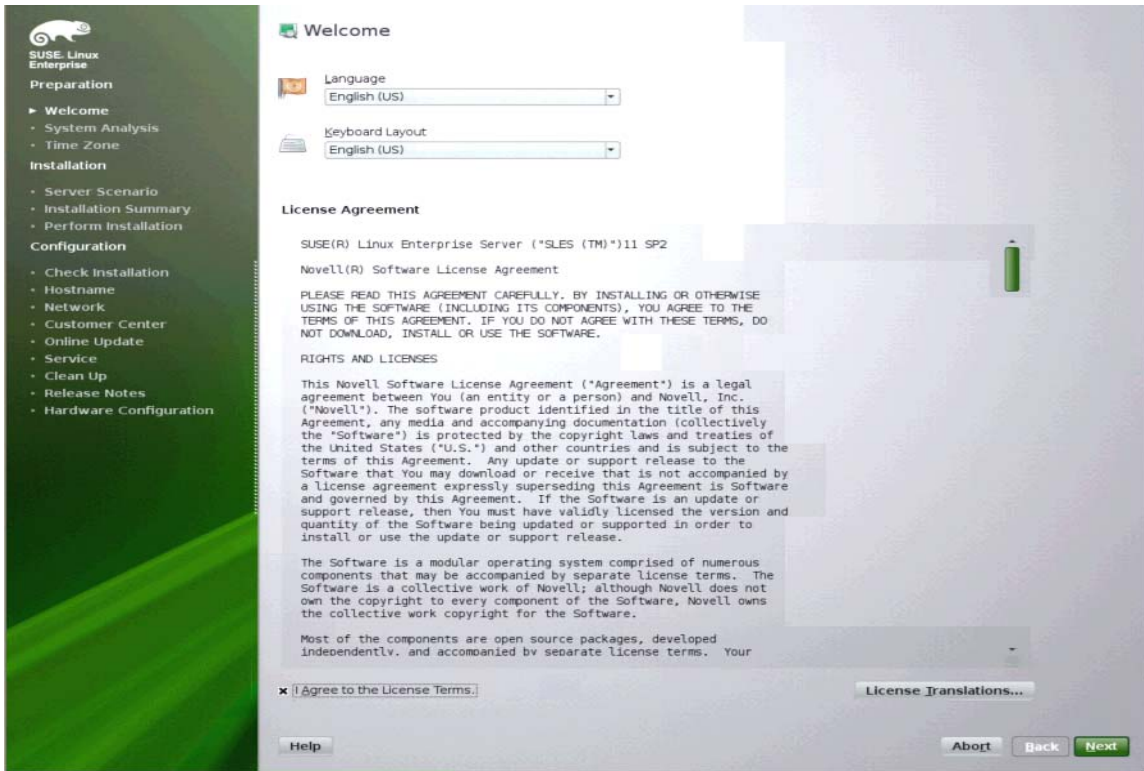
2.2.11.2 Operating System Installation in AHCI Mode

Connect a physical drive to the sSATA controller that is configured in AHCI mode.

1. Configure the BIOS and OpROM settings.
2. Start the installation.

The following window appears after the SLES switches to the installation mode GUI.

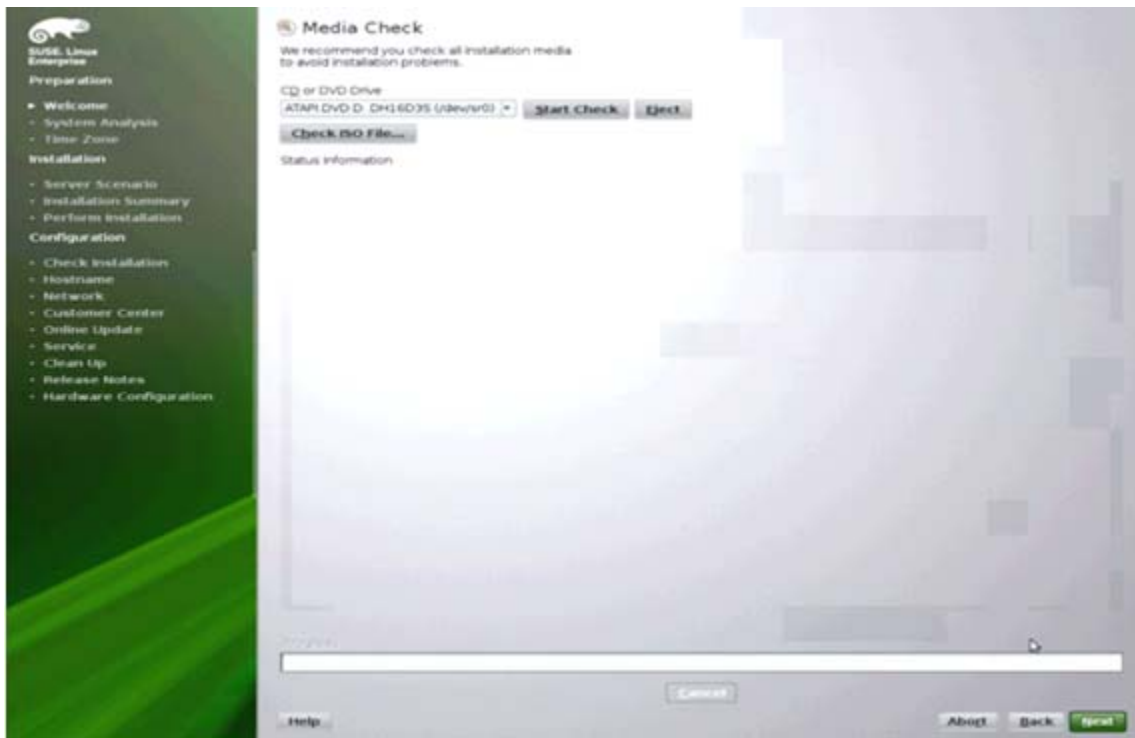
Figure 61 Grantley ACHI Mode Welcome Screen



3. Accept the Licensing agreement and click **Next** to continue.


The following shows from which installation media to start the OS installation.

Figure 62 Grantley ACHI Mode Media Check Screen



4. Select the media from which to load the operating system and click **Next** to continue. The following screen appears.

Figure 63 Grantley ACHI Mode Installation Mode Screen




5. Choose the desired installation mode and click **Next** to continue.

NOTE Select **New Installation** when performing a fresh installation.

Three options are available during the installation.

- Physical Machine, which is also available for fully virtualized guests
 - Virtual Machine, which is for para-virtualized environments like Xen
 - Xen Virtualization Host, if virtual machines must be created
6. Select the **Partitioning** option if you need to change the partition parameters.
It is recommended that you leave this option unchanged, if you are not clear on what you are trying to achieve.
For more details on partitioning, refer to the SLES/RHEL installation guide.
 7. Select the **Software** option to install additional software packages.
 8. If additional software packages are required, check them and click **Next** to continue.

Figure 64 Grantley ACHI Mode Software Selection and System Tasks



9. Reboot the system from the hard drive after finishing the installation.

The AHCI driver now claims both controllers. To claim a SATA controller using the `megasr` driver and the sSATA controller using the AHCI driver at the same time, perform the steps in the following section.

2.2.11.3 Installing the Software RAID Driver (RPM) Package

1. The Software RAID Driver RPM package uses the DKMS package, which can be found at the following link (can be subject to change):

<http://linux.dell.com/dkms/>

2. Download the latest stable RPM package.

The format is as follows:

```
dkms-x.x.xx.x-x.noarch.rpm
```

3. Save the ROM driver to a folder on the system.
4. Open a Terminal window and navigate to the folder.
5. Install the `dkms` RPM driver first using the following command:

```
rpm -ivh dkms-x.x.xx.x-x.noarch.rpm
```

6. Install the `megasr-dkms` RPM driver using the following command:

```
rpm -ivh megasr-x.x.xx.x-x.noarch.rpm
```

7. Edit the `/etc/modprobe.d/unsupported-modules` file after installing the RPM driver.
8. Change the entries as follows:

```
allow_unsupported_modules 1
```

9. Add the following statement to load the MegaSR driver in the `/etc/modprobe.conf` file, before loading the AHCI Driver.

```
install ahci /sbin/modprobe megasr 2>&1 |:;  
/sbin/modprobe -ignore -install ahci
```

Use the previous statement to change the device driver load order.

10. To create the `initrd` image, perform the following command.

```
mkinitrd
```

11. Reboot the machine.

After the machine reboots, the MegaSR and AHCI drivers should load successfully. The MegaSR driver should claim SATA controller (which is configured as RAID mode) and AHCI driver should claim sSATA controller (which is in AHCI mode).

Chapter 3: Software RAID Configuration Utility

Use the configuration utility (CU) to configure disk drive groups and virtual drives, and to perform other configuration tasks in a pre-boot environment.

You can perform the following functions using the Avago Software RAID Configuration Utility:

- Select a configuration method for drive groups and virtual drives
- Create drive groups
- Define virtual drives
- Initialize virtual drives
- Access controllers, virtual drives, and drive groups to view their properties
- Create hot spare drives
- Verify that the redundancy data in virtual drives using RAID level 1, RAID level 5, and RAID level 10 is correct
- Rebuild failed drives
- Reconstruct virtual drives after changing RAID levels or adding a drive to a drive group
- Select a MegaRAID host adapter

NOTE If the configuration utility does not display, go into the BIOS setup and disable Quick boot, Fast boot, Silent boot, Intel® Rapid boot, and Quick POST, then reboot. If you are still unable to access the configuration utility, check for a system BIOS upgrade.

3.1 Performing a Quick Configuration

This section provides high-level instructions to quickly configure drive groups and virtual drives with the Embedded RAID utility. These instructions are intended for users that are familiar with configuration utilities and tools. See [Section 3.4, Configuring Drive Groups and Virtual Drives](#), on page 74, for detailed configuration instructions. To ensure the best performance, select the optimal RAID level for the virtual drive you create. For an explanation of RAID levels, see [Section 1.2, RAID Overview](#), on page 8.

Perform the following steps to configure drive groups and virtual drives using the Embedded RAID utility:

1. Boot the system.
2. The Avago Software RAID CU banner appears for each supported controller in sequence.


NOTE The Avago Software RAID CU can support more than one controller in the system. For example, some systems support an AHCI and SCU controller; others might support two AHCI controllers (such as SATA and sSATA).

3. Press Ctrl+M beneath the appropriate Avago Software RAID CU banner to start the configuration utility for that particular controller.
4. Select **Configure** on the **Management Menu** screen.
5. Select a configuration method from the **Configuration** menu (**Easy Configuration**, **New Configuration**, or **View/Add Configuration**).
6. Create drive groups using the available drives.
7. Designate hot spare disks (optional).
8. Define the virtual drive(s) using the space in the drive groups.
9. Initialize the new virtual drives.

3.2 Management Menu

The **Management Menu** screen appears when you start the Embedded RAID utility.

Figure 65 Software RAID Configuration Utility Management Menu Screen



NOTE The minimum screen resolution for the Embedded RAID graphical user interface (GUI) is 640 × 480.

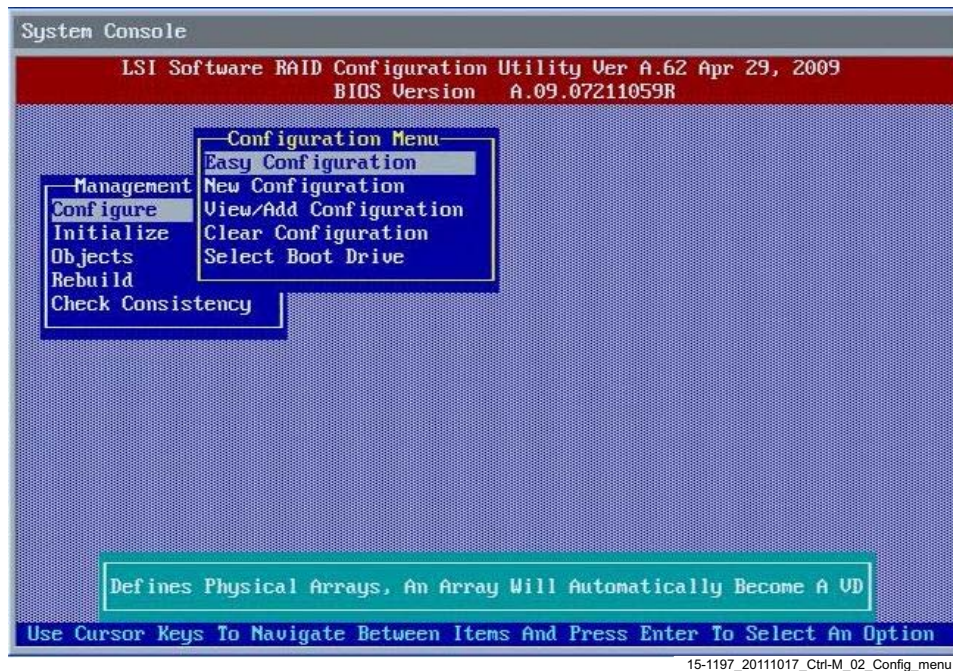
3.3 Configuration Menu

Use the **Configuration Menu** screen to configure drive groups and virtual drives. This section describes the configuration options.

3.3.1 Configuration Menu Options

The **Configuration Menu** screen provides four methods to modify and/or create a virtual drive configuration: **Easy Configuration**, **New Configuration**, **View/Add Configuration**, and **Clear Configuration**, as shown in the following figure.

Figure 66 Software RAID Configuration Menu



The following table describes these methods. The **Configuration Menu** has an **Advance** submenu that lets you set specific options. The available options depend upon the configuration method you use.

Table 1 Software RAID Configuration Utility – Configuration Menu

Option	Description
Easy Configuration	The Easy Configuration option automatically associates every drive group with one virtual drive. Through the Advance Menu , the Easy Configuration option lets you modify the RAID level and stripe size. Section 3.5.2, Using the Easy Configuration Option , on page 75, provides detailed instructions.
New Configuration	The New Configuration option lets you modify the RAID level, stripe size, virtual drive size, and disk spanning (associating virtual drives with multiple drive groups). If you select the New Configuration option, the CU deletes the existing configuration information on the selected controller when saving the new configuration. Section 3.5.3, Using the New Configuration Option , on page 79, provides detailed instructions.
View/Add Configuration	The View/Add Configuration option lets you control the same virtual drive parameters as the New Configuration option without disturbing the existing configuration information. The View/Add Configuration option also lets you enable the Configuration on Disk feature. Section 3.5.4, Using the View/Add Configuration Option , on page 82, provides detailed instructions.
Clear Configuration	The Clear Configuration option erases the current configuration information. Section 3.6, Clearing a Storage Configuration , on page 85, provides detailed instructions.
Specify Boot Drive	The Specify Boot Drive option lets you specify a virtual drive as the boot drive on the controller, if you have created virtual drives. Section 3.7, Configuring a Bootable Virtual Drive , on page 86, provides detailed instructions.

3.4 Configuring Drive Groups and Virtual Drives

The following sections provide detailed instructions for configuring drive groups and virtual drives with the Embedded RAID utility.

NOTE Use drives with the same capacity when you create a storage configuration. If you use drives with different capacities in one drive group, the CU limits each drive to the capacity of the smallest drive.

The number of physical drives in a specific drive group determines the possible RAID levels that you can implement with the drive group.

- RAID level 0 requires from one to eight physical drives.
- RAID level 1 requires two physical drives.
- RAID level 5 requires three to eight physical drives.
- RAID level 10 requires four, six, or eight physical drives.

3.5 Creating a Storage Configuration

This section explains how to use the Embedded RAID utility to configure RAID drive groups and virtual drives to create storage configurations.

3.5.1 Selecting the Configuration Method

The `Easy Configuration` option automatically associates each drive group with one virtual drive. Follow these steps to open the **Configuration Menu** screen, and select a configuration method:

1. On the **Management Menu** screen, highlight **Configure** with your cursor, and press Enter.
The **Configuration Menu** screen appears, as shown in [Figure 66, Software RAID Configuration Menu](#).
2. Select a configuration option.
 - `Easy Configuration`
Automatically associates every drive group with one virtual drive.
 - `New Configuration`
Clears the existing configuration and lets you create a new configuration.
 - `View/Add Configuration`
Retains the existing storage configuration and adds new drives to it (this method does not cause any data loss).
 - `Clear Configuration`
Clears the existing configuration.

CAUTION If you choose to clear the existing configuration or to create a new configuration, the system deletes all the existing data in the configuration. Make a backup of any data that you want to keep before you choose an option.


3. Press Enter.
The configuration screen appears for the configuration option that you selected.
A dialog box warns that you will lose data if you select the `Clear Configuration` option or the `New Configuration` option.
The following sections describe the configuration steps for each configuration method.

3.5.2 Using the Easy Configuration Option

Follow these instructions to create a configuration with the `Easy Configuration` option, with or without redundancy:

1. On the **Management Menu** screen, highlight **Configure** with your cursor, and press Enter.
The **Configuration Menu** screen appears, as shown in [Figure 66, Software RAID Configuration Menu](#).
2. Use your cursor to highlight **Easy Configuration** and press Enter.
The following screen appears.

Figure 67 Easy Configuration – Array Selection Menu



15-1197_20111017_Ctrl-M_04_Easy_config_menu2

3. Press the space bar to select a drive and add it to the drive group.
4. Repeat step 3 to add additional drives to the drive group.
5. Press F10 to continue configuration.
The **Select Configurable Array(s)** dialog box appears, as shown in the following figure.

Figure 68 Drive Group Selection

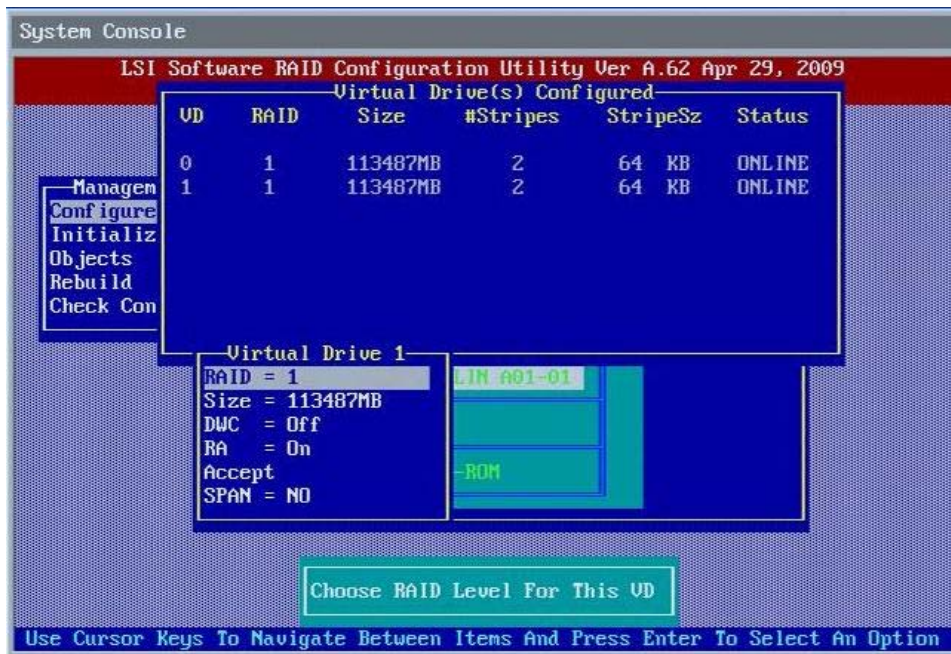


15-1197_20111017_Ctrl-M_05_Easy_config_menu3

6. Press the space bar to select the drive group.
7. Press F10 to continue configuration.

The **Virtual Drive(s) Configured** dialog box appears, as shown in the following screen. This screen shows the virtual drive number, RAID level, drive group size, number of stripes, stripe size, and drive status.

Figure 69 Virtual Drive Parameters Dialog Box



15-1197_20111017_Ctrl-M_07_Easy_config_menu5

8. Change the virtual drive options from the defaults listed on the screen as needed.
Use the cursor keys to navigate between the virtual drive parameters and press Enter to select an option.
The following list briefly describes the virtual drive options:
 - **RAID Level**
The drop-down menu lists the possible RAID levels for the virtual drive. In some cases, only one RAID level is available, depending on the number of drives in the drive group.
 - **Size**
This setting specifies the capacity of the virtual drive.
 - **Disk Write Cache Policy**
When the disk Write Cache policy is **On**, a write transaction is considered to be complete when all the data has been written to the disk cache. When the disk Write Cache policy is **Off**, the write transaction is complete only when the data has been written to the disk.
 - **Read Ahead Policy**
When the disk Read Ahead policy is **On**, extra data is read sequentially ahead of the data that is actually requested, and this extra data is stored in cache memory. If the additional read-ahead data is then requested, it can be read faster from the cache than from the disk directly. This setting speeds up reads for sequential data, but there is little improvement when accessing random data.
 - **Accept**
Select this option to accept the virtual drive parameters.
 - **SPAN**
Choose whether to span drive groups. This setting is available only for RAID 10 drive groups.
9. Highlight **Accept** and press Enter after you select your virtual drive parameters.
The virtual drive configuration appears, as shown in the following screen.

Figure 70 Virtual Drive Configuration Parameters



15-1197_20111017_Ctrl-M_11_Easy_config_menu9

10. Press any key to continue.

11. Click **Yes** at the prompt to save the configuration.

The Avago Software RAID Configuration Utility begins a background initialization of the virtual drives.

NOTE New RAID 5 virtual drives require a minimum number of drives for a background initialization to start. If fewer drives than the minimum exist, the background initialization will not start.

3.5.3 Using the New Configuration Option

If you select the **New Configuration** option, the CU deletes the existing configuration information on the selected controller when it saves the new configuration.

Follow these instructions to create a configuration with the **Easy Configuration** option, with or without redundancy:

1. On the **Management Menu** screen, highlight **Configure** with your cursor, and press Enter. The **Configuration Menu** screen appears, as shown in [Figure 66, Software RAID Configuration Menu](#).
2. Highlight **New Configuration** with your cursor, and press Enter. The following screen appears. This screen displays the drives.

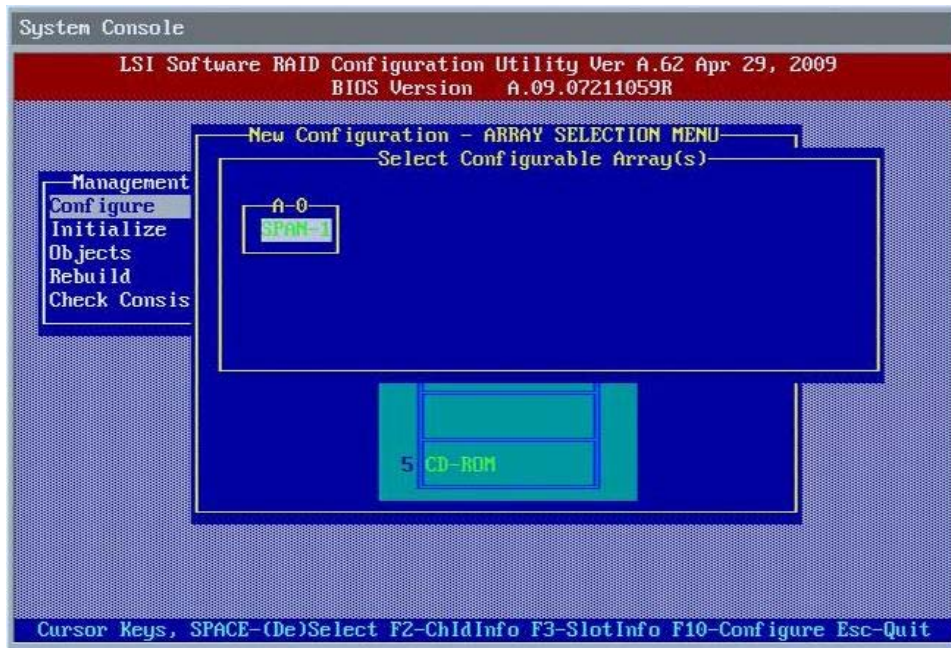
Figure 71 New Configuration – Array Selection Menu



15-1197_20111017_Ctrl-M_12_New_config_menu1

3. Press the space bar to select a drive and add it to the drive group.
4. Repeat step 3 to add additional drives to the drive group.
5. Press F10 to continue the configuration. The **Select Configurable Array(s)** dialog box appears, as shown in the following figure.

Figure 72 Drive Group Selection

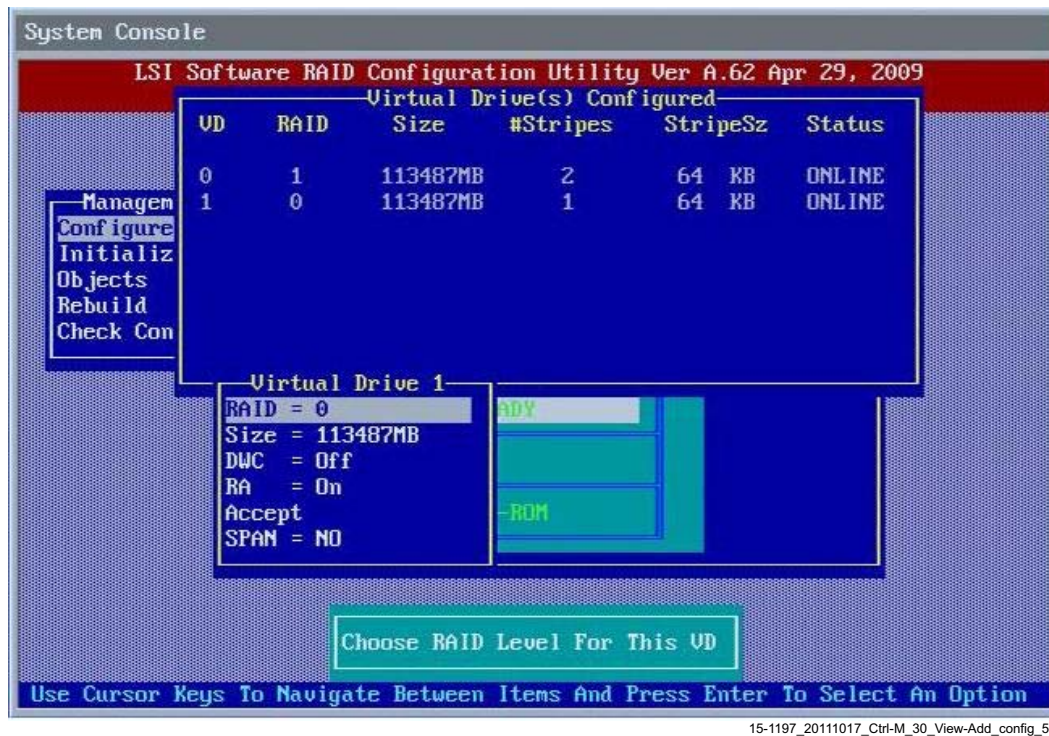


15-1197_20111017_Ctrl-M_14_New_config_menu3

6. Press the space bar to select the drive group.
7. Press F10 to continue the configuration.

The **Virtual Drive(s) Configured** dialog box appears, as shown in the following screen. This screen shows the virtual drive number, RAID level, drive group size, number of stripes, stripe size, and drive status.

Figure 73 Virtual Drive(s) Configured Dialog Box



8. Change the virtual drive options from the defaults listed on the screen as needed. Use the cursor keys to navigate between the virtual drive parameters and press Enter to select an option. The following list provides brief explanations of the virtual drive options:
 - **RAID Level**
 The drop-down menu lists the possible RAID levels for the virtual drive. In some cases, only one RAID level is available, which depends on the number of drives in the drive group.
 - **Size**
 This setting specifies the capacity of the virtual drive.
 - **Disk Write Cache Policy**
 You can disable the `Write Cache` option when you create a virtual drive, but you can enable this option later using the configuration utilities. When the `disk Write Cache` option is **On**, a write transaction is considered to be complete when all the data has been written to the disk cache. When the `disk Write Cache` option is **Off**, the write transaction is complete only when the data has been written to the disk.
 - **Read Ahead Policy**
 When the `disk Read Ahead` option is **On**, extra data is read sequentially ahead of the data that is actually requested, and this extra data is stored in cache memory. If the additional read-ahead data is then requested, it can be read faster from the cache than from the disk directly. This setting speeds up reads for sequential data, but there is little improvement when accessing random data.
 - **Accept**
 Select this option to accept the virtual drive parameters.
 - **SPAN**
 Choose whether to span drive groups. This setting is available only for RAID 10 drive groups.
9. Highlight **Accept** with your cursor, and press Enter after you select your virtual drive parameters. The virtual drive configuration appears, as shown in [Figure 70, Virtual Drive Configuration Parameters](#).

10. Press any key to continue.
11. Click **Yes** at the prompt to save the configuration.
The Avago Software RAID Configuration Utility begins a background initialization of the virtual drives.
New RAID 5 virtual drives require a minimum number of drives for a background initialization to start. If fewer drives exist than the minimum, the background initialization will not start.

3.5.4 Using the View/Add Configuration Option

The View/Add Configuration option lets you control the same virtual drive parameters as the New Configuration option without disturbing the existing configuration information.

Follow these instructions to create a configuration with the Easy Configuration option, with or without redundancy:

1. On the **Management Menu** screen, highlight **Configure** with your cursor, and press Enter.
The **Configuration Menu** screen appears, as shown in the following figure.

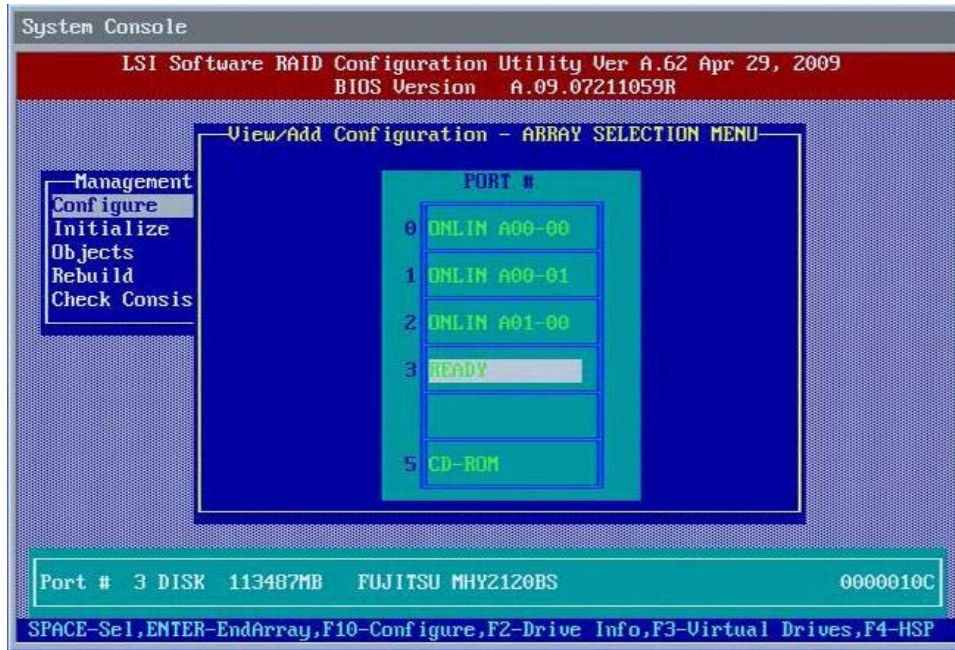
Figure 74 Configuration Menu Screen



15-1197_20111017_Ctrl-M_26_View-Add_config_1

2. Highlight View/Add Configuration with your cursor, and press Enter.
The following screen appears.

Figure 75 View/Add Configuration – Array Selection Menu



15-1197_20111017_Ctrl-M_28_View-Add_config_3

3. Press the space bar to select a drive and add it to the drive group.
4. Repeat step 3 to add additional drives to the drive group.
5. Press F10 to continue the configuration.

The **Select Configurable Array(s)** dialog box appears, as shown in the following figure.

Figure 76 Select Configurable Arrays Dialog Box

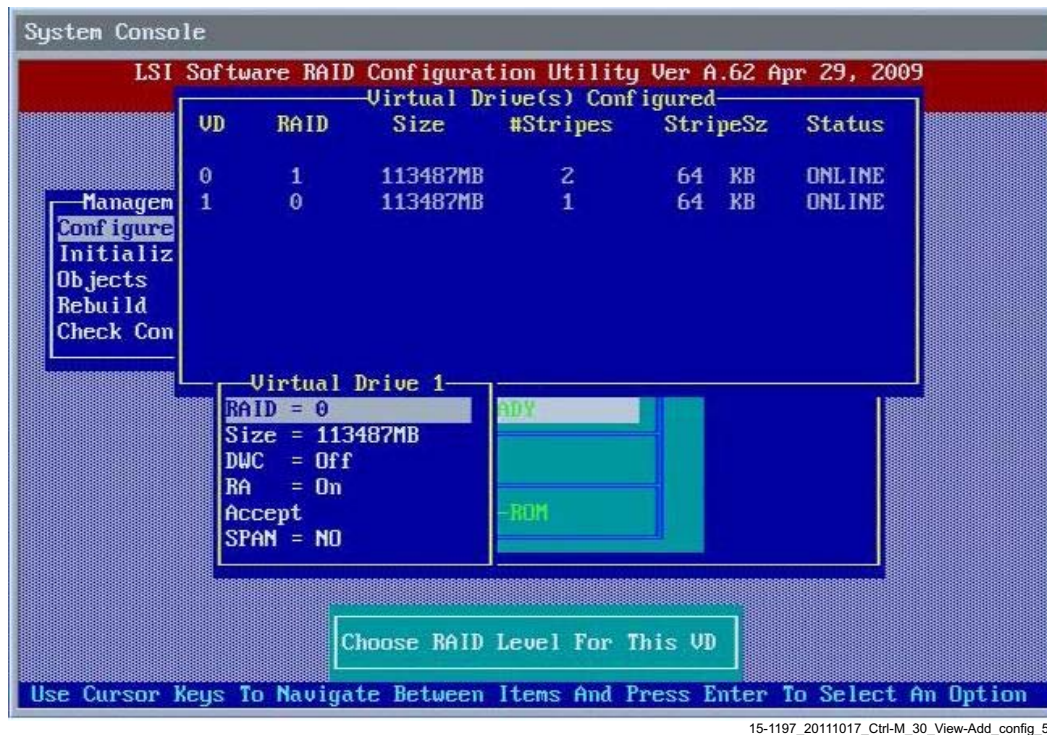


15-1197_20111017_Ctrl-M_29_View-Add_config_4

6. Press the space bar to select the drive group.
7. Press F10 to continue the configuration.

The **Virtual Drive(s) Configured** dialog box appears, as shown in the following screen. This screen shows the virtual drive number, RAID level, drive group size, number of stripes, stripe size, and drive status.

Figure 77 Virtual Drive Parameters Dialog Box



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

Use the cursor keys to navigate between the virtual drive parameters and press Enter to select an option.

The following list contains brief explanations of the virtual drive options:

- **RAID Level**

The drop-down menu lists the possible RAID levels for the virtual drive. In some cases, only one RAID level is available, which depends on the number of drives in the drive group.

- **Size**

This setting specifies the capacity of the virtual drive.

- **Disk Write Cache Policy**

You can disable the disk Write Cache option when you create a virtual drive, but you can enable this option later using the configuration utilities. When the disk Write Cache policy is **On**, a write transaction is considered to be complete when all the data has been written to the disk cache. When the disk Write Cache policy is **Off**, the write transaction is complete only when the data has been written to the disk.

- **Read Ahead Policy**

When the disk Read Ahead policy is **On**, extra data is read sequentially ahead of the data that is actually requested, and this extra data is stored in cache memory. If the additional read-ahead data is then requested, it can be read faster from the cache than from the disk directly. This setting speeds up reads for sequential data, but there is little improvement when accessing random data.

- **Accept**

Select this option to accept the virtual drive parameters.

— **SPAN**

Choose whether to span drive groups. This setting is available only for RAID 10 drive groups.

9. Highlight **Accept** with your cursor, and press Enter.

The virtual drive configuration appears.

10. Press any key to continue.

11. Click **Yes** at the prompt to save the configuration.

The Avago Software RAID Configuration Utility begins a background initialization of the virtual drives.

New RAID 5 virtual drives require a minimum number of drives for a background initialization to start. If fewer drives exist than the minimum required, the background initialization will not start.

3.6 Clearing a Storage Configuration

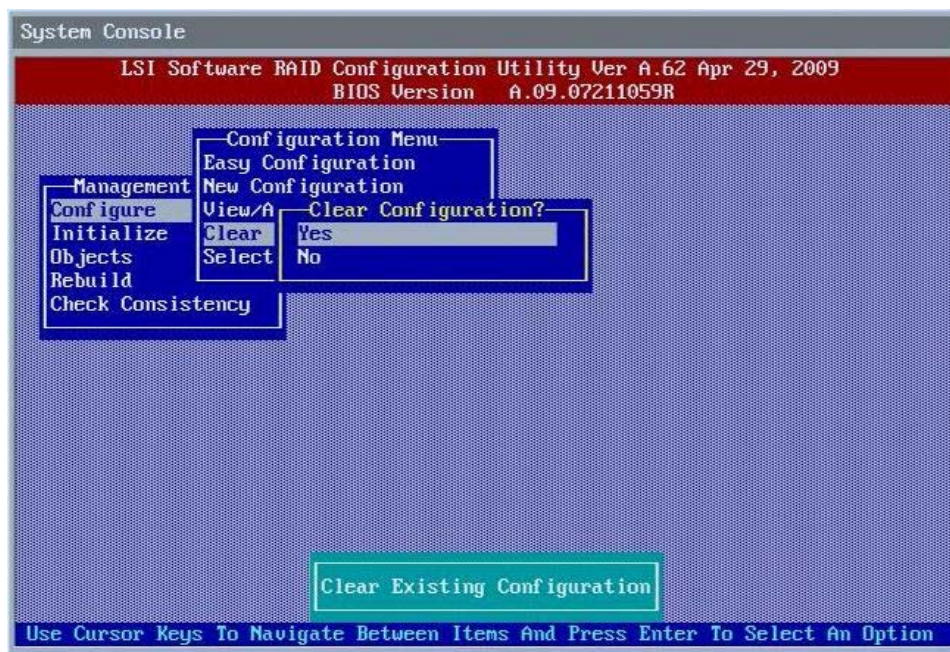
CAUTION Before you clear a storage configuration, be sure to back up all the data you want to keep.

To clear a storage configuration, follow these steps:

1. On the **Management Menu** screen, select **Configure > Clear Configuration**.

The following screen appears.

Figure 78 Clear Configuration Option



15-1197_20111017_Ctr+M_32_Clear_config_2

2. At the prompt, select **Yes** to confirm and press Enter.

The virtual drive is deleted from the configuration.

3.7 Configuring a Bootable Virtual Drive

The default boot virtual drive is LD 0. If you change the boot drive to another virtual drive, the BIOS and the CU preserve this change.

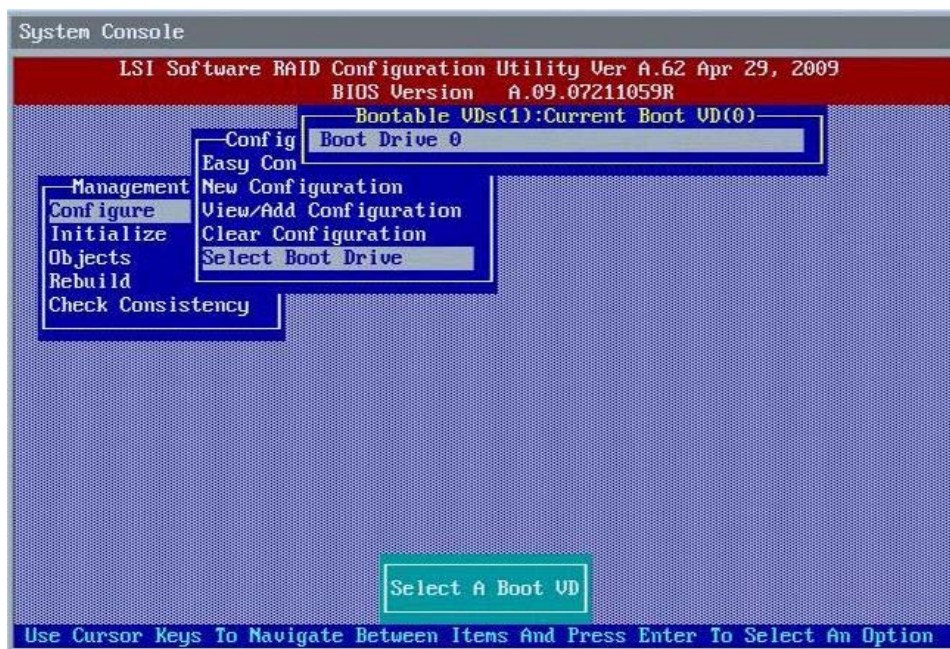
NOTE If you delete the new boot virtual drive, you *must* configure another virtual drive to boot the system. The CU does not automatically select a different boot virtual drive.

Follow these steps to configure a bootable virtual drive:

1. On the **Management Menu** screen, select **Configure > Select Boot Drive**.

The following screen appears.

Figure 79 Select Boot Drive Option



15-1197_20111017_Ctrl-M_34_Select_boot_drive2

2. Select a virtual drive from the list to be the designated boot drive.

The virtual drive is configured to be the boot drive.

3.8 Initializing Virtual Drives

This section explains the two methods used to initialize a virtual drive with the MegaRAID Storage Manager Utility.

If the *Fast Init* property is enabled, fast initialization is used. In fast initialization, the Embedded RAID utility quickly writes zeroes to the first and last 8-MB regions of the new virtual drive.

If the *Fast Init* property is not enabled, the utility performs a complete initialization on the virtual drive. This process can take a long time if the physical disk drives are large.


CAUTION When you initialize a virtual drive, all the existing data on the virtual drive is erased.

3.8.1 First Initialization Method

Follow these steps to initialize a virtual drive using the **Initialize** menu.

1. On the **Management Menu** screen, highlight **Initialize** with the cursor, and press Enter.
The list of virtual drives appears.
2. Use the cursor to highlight the virtual drive you want to initialize (if more than one virtual drive exists).
3. Press Enter.
The following screen appears.

Figure 80 Initializing a Virtual Drive – First Method



15-1197_20111017_Ctrl-M_35_Initialize

4. Press F10.
5. Select **Yes** at the prompt and press Enter to begin the initialization.
A graph shows the progress of the initialization until it is complete.
6. After the initialization is complete, press Esc to return to previous menus.
If you press Esc while initialization is in progress, the following options appear:
 - **Stop**
(Available only if *Auto Resume* is enabled on the controller: **Management Menu > Objects > Adapter > AutoResume**.)
The initialization is stopped, and the CU stores the percentage of the initialization already completed. If *Auto Resume* is enabled, and if *Fast Init* is not enabled, the initialization resumes where it left off when you restart it, instead of starting over from 0 percent.

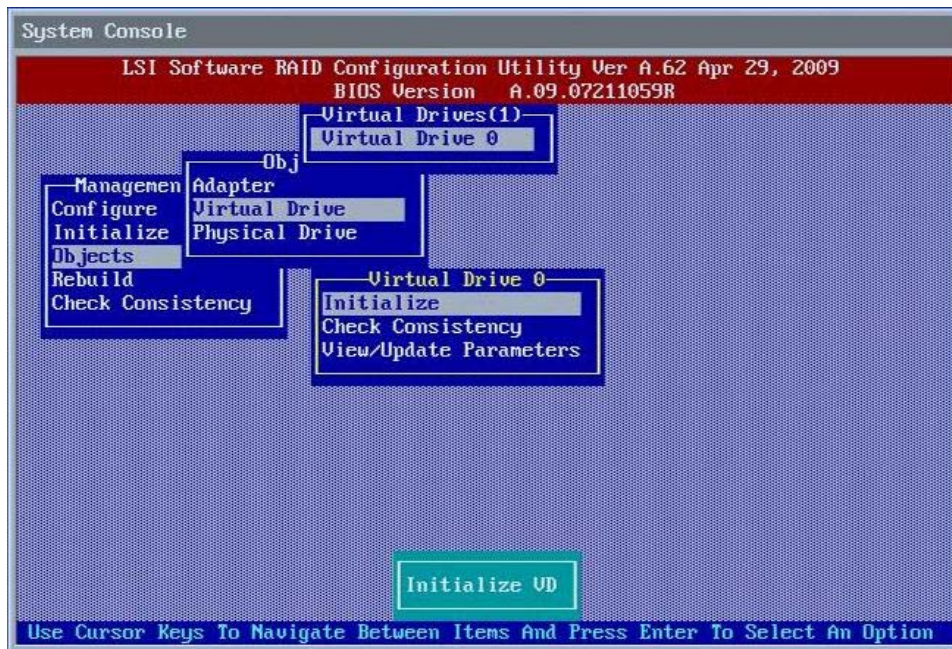
- **Continue**
The initialization continues normally.
- **Abort**
The initialization is completely aborted. If you restart initialization, it begins at 0 percent.

3.8.2 Second Initialization Method

Follow these steps to initialize a virtual drive using the **Objects** menu.

1. On the **Management Menu** screen, select **Objects > Virtual Drive**.
A list of configured virtual drives appears.
2. Use the cursor to highlight the virtual drive you want to initialize (if more than one virtual drive exists), and press Enter.
The following screen appears.

Figure 81 Initializing a Virtual Drive – Second Method



3. Highlight **Initialize** with your cursor on the virtual drive submenu, and press Enter.
4. Select **Yes** at the prompt, and press Enter.
The utility displays a bar graph showing the initialization progress.
5. When initialization completes, press Esc to return to the previous menu.
If you press Esc while initialization is in progress, the **Stop**, **Continue**, and **Abort** options are available, as explained in [Section 3.8.1, First Initialization Method](#), on page 87.

3.9 Rebuilding a Drive

The Embedded RAID utility lets you rebuild a drive of a redundant drive group if the drive group has a failed drive. If the failed drive is still good (that is, if the drive is physically present and its capacity is greater than or equal to the defined capacity of the drive group), it will be rebuilt. If the drive is too small, an error message appears and the CU does not allow the drive to be rebuilt.

NOTE You cannot rebuild a failed drive if the drives' capacity is even 1 byte smaller than the defined capacity of the drive group.

Follow these steps to rebuild a drive:

1. On the **Management Menu** screen, highlight **Rebuild** and press Enter.
2. When the list of drives appears, highlight the failed (*FAIL*) drive that you want to rebuild and press the spacebar to select it.
3. After you select the drive, press F10 to start the *rebuild* process, and then select **Yes** at the confirmation prompt.

The *rebuild* process begins, and a graph shows the progress of the rebuild until it is complete. Although the CU changes the disk drive state to *Rebuild* at this point, the change does not appear on the screen while the *rebuild* operation is in progress.

If the CU detects a media error on the source drive during the *rebuild* operation, it initiates a sector read for that block. If the sector read fails, the CU adds entries to the SBBM table, writes this table to the target drive, and displays an error message.

Additional error messages appear if the SBBM table is 80 percent full or 100 percent full. If the SBBM table is completely full, the *rebuild* operation is aborted, and the drive is marked as *FAIL*.

4. When the *rebuild* operation is complete, the CU displays the message that the *rebuild* operation is successful.
5. Press Esc to display the **Management Menu** screen.

The state of the rebuilt disk drive changes from *FAIL* to *ONLIN*.

If you press Esc while the *rebuild* operation is running, the following options display:

— **Stop**

Available only if the *AutoResume* property is enabled on the adapter: **Management Menu > Objects > Adapter > AutoResume.**

The *rebuild* operation is stopped, and the CU stores the percentage of the *rebuild* operation already completed. If the *AutoResume* property is enabled, the *rebuild* operation resumes where it left off when you restart it, instead of starting over from 0 percent.

— **Continue**

The *rebuild* operation continues normally.

— **Abort**

The *rebuild* operation is completely aborted and the disk drive remains in the *FAIL* state. If you restart the *rebuild* operation, it begins at 0 percent.

3.10 Creating a Global Hot Spare Drive

The Embedded RAID utility lets you create global hot spare drives to protect against data loss. A *hot spare* is an unused drive that you can use to rebuild data from a failed drive and re-establish redundancy, in the case of a disk failure in a redundant RAID drive group (RAID 1, RAID 5, or RAID 10).


NOTE Dedicated hot spare drives are not supported by the Embedded RAID utility.

NOTE When you select a drive to change into a global hot spare, be sure it is the same type of drive as the drives in the drive group that it will protect.

You can create a hot spare when you configure a new storage configuration, as described in the previous sections. To add a hot spare drive to an existing redundant storage configuration, follow these steps:

1. On the **Management** menu, select **Configure > View/Add Configuration**.
2. Select **Physical Drive**.
A list of physical drives appears.
3. Highlight an unconfigured drive or a *Ready* drive with your cursor, and press Enter.
The following screen appears.

Figure 82 View/Add Configuration – Array Selection Menu



15-1197_20111017_Ctrl-M_40_Hotspare_creation

4. In the **HotSpare** dialog box, select **Yes** and press Enter.
5. Select **Yes** from the pop-up menu to create the hot spare drive.
This procedure creates a hot spare for the drive group.

3.11 Checking Data Consistency

The Check Consistency feature verifies the consistency of the data on the physical drives that are part of RAID 1, RAID 5, or RAID 10 virtual drives. The Embedded RAID utility automatically corrects any differences found in the data when a consistency check is run.

Follow these steps to check consistency:

1. On the **Management Menu** screen, highlight **Check Consistency** with your cursor, and press Enter.
A list of configured virtual drives appears.

2. Use the cursor to highlight a virtual drive (if there is more than one virtual drive).
3. Press the spacebar to select the virtual drive.

NOTE

If you select a RAID 0 virtual drive, a message appears stating that a *Check Consistency* operation cannot be performed. To continue, deselect the virtual drive, highlight a redundant virtual drive, and press the spacebar again.

4. Press F10.
5. At the prompt, select **Yes** to start the *Check Consistency* operation, and press Enter.

A graph shows the progress of the *Check Consistency* operation until it is complete.

If the Embedded RAID utility finds any data inconsistencies while comparing the source drive and the target drive, the utility fixes the inconsistency by writing the source data to the target drive. When this happens, a message notifies you that inconsistent data exists on the drives and that the repair has been performed.

If you press Y, the program skips the bad block and continues. If you press N, the program aborts the consistency check. The same message appears if the program finds a hard media error on the target drive.

If you press Esc while the *Check Consistency* operation is running, the following options appear:

— **Stop**

Available only if the *AutoResume* property is enabled on the adapter: **Management Menu > Objects > Adapter > AutoResume**

The *Check Consistency* operation is stopped, and the CU stores the percentage of the task already completed. If the *AutoResume* property is enabled, the *Check Consistency* operation resumes where it left off when you restart it, instead of starting over from 0 percent.

— **Continue**

The *Check Consistency* operation continues normally.

— **Abort**

The *Check Consistency* operation is completely aborted. If you restart the operation, it begins at 0 percent.

3.12 Displaying and Changing Controller Properties

This section explains how you can use the Embedded RAID utility to display and change the properties for the controllers.

3.12.1 Displaying and Changing Controller Properties

Follow these steps to display the properties of a controller.

1. On the **Management Menu** screen, select **Objects > Adapter**.
The list of controllers appear in a dialog box.
2. Highlight a controller with your cursor, and press Enter.
The following screen appears.

Figure 83 Controller Properties Screen



15-1197_20111017_Ctrl-M_42_Objects_Adapter_options

The following table describes the entries/options listed on the controller properties screen. Avago recommends that you leave these options at their default settings to achieve the best performance, unless you have a specific reason to change them.

Table 2 Controller Properties and Values

Property	Options	Default
Rebuild Rate	0 to 100 (percentage of system resources) Use this option to select the rebuild rate for drives connected to the selected controller. The rebuild rate is the percentage of system resources dedicated to rebuilding a failed drive. The higher the number, the more system resources devoted to a rebuild.	30
Chk Const Rate	0 to 100 (percentage of system resources) Use this option to select the amount of system resources dedicated to consistency checks of virtual drives connected to the selected controller.	30
FGI Rate (Foreground Initialization Rate)	0 to 100 (percentage of system resources) Use this option to select the amount of system resources dedicated to foreground initialization of virtual drives connected to the selected controller.	30
BGI Rate (Background Initialization Rate)	0 to 100 (percentage of system resources) Use this option to select the amount of system resources dedicated to background initialization of virtual drives connected to the selected controller.	30
Disk WC (Disk Write Cache)	Off, On You can disable the Disk Write Cache option when you create a virtual drive, but you can enable this option later using the configuration utilities. When the Disk Write Cache option is On , a write transaction is considered to be complete when all the data has been written to the disk cache. When Disk Write Cache is Off , the write transaction is complete only when the data has been written to the disk.	Off (Write Through Enabled)

Table 2 Controller Properties and Values (Continued)

Property	Options	Default
Read Ahead (RA)	On, Off When the disk Read Ahead option is On , extra data is read sequentially ahead of the data that is actually requested, and this extra data is stored in cache memory. If the additional read-ahead data is then requested, it can be read faster from the cache than from the disk directly. This setting speeds up read operations for sequential data, but there is little improvement when accessing random data.	On
BIOS State	Enable, Disable Use this option to enable or disable the BIOS for the selected controller. 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.	Enable
Cont On Error	No, Yes Enable this option if you want the boot process to continue when the controller BIOS encounters an error during boot-up.	Yes
Fast Init	Enable, Disable A fast initialization quickly writes 0s to the first and last 10-MB regions of the new virtual drive and then completes the initialization in the background. It is seldom necessary to use this option, because the virtual drive was already initialized when you created it.	Enable
Auto Rebuild	On, Off The Auto-Rebuild option allows a failed drive to be replaced and the data automatically rebuilt by <i>hot-swapping</i> the drive in the same drive bay. The RAID drive group continues to handle requests while the rebuild occurs.	On
Auto Resume	Enable, Disable When enabled, you can stop a consistency check, rebuild, or initialization, and resume it later where it left off, instead of aborting it and starting over.	Enable
Disk Coercion	None, 128MB, 1GB The Disk Coercion option is a tool to force drives of varying capacities to be the same capacity so they can be used in a drive group. The coercion mode options are None, 128MB-way, and 1GB-way. The number you choose depends on how much the drives from various vendors vary in their actual size.	1GB
Factory Default	Returns the settings to the factory default settings.	N/A

3. To change the value of a controller property, use the cursor to highlight the property, and press Enter.
4. Select or type a different value for the property, and press Enter.
When you are finished, press Esc until you return to the **Management Menu** screen.

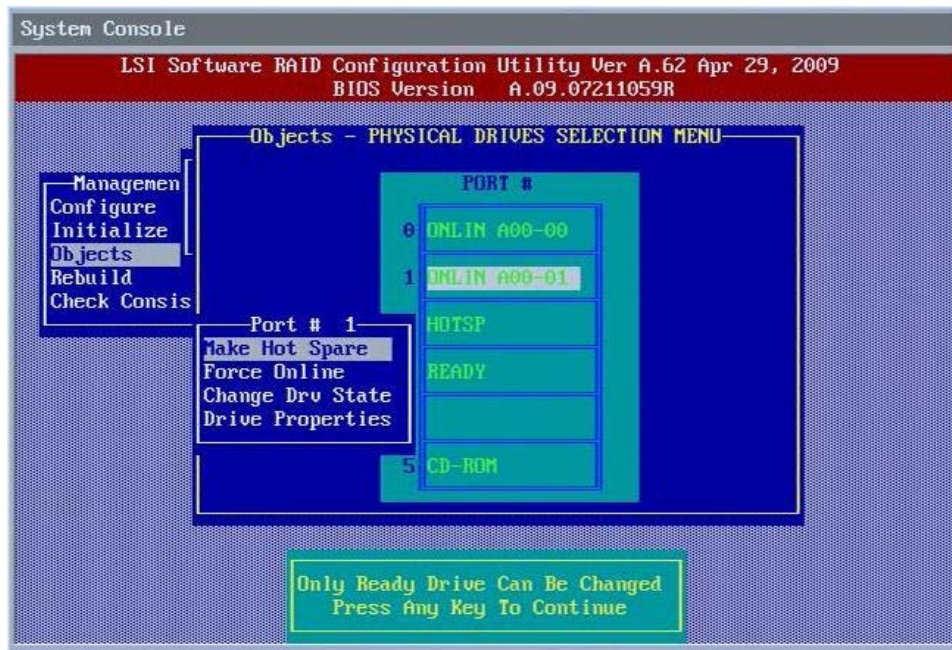
3.12.2 Displaying and Changing Drive Properties

The Embedded RAID utility displays properties, policies, and operations for virtual drives.

Follow these steps to display the virtual drive information.

1. On the **Management Menu** screen, select **Objects > Adapter**.
The list of controllers appear in a dialog box.
2. Highlight a controller with your cursor, and press Enter.
The **Physical Drive Selection Menu** screen appears.
3. Highlight a drive with your cursor, and press Enter.
The following screen appears. This screen lists the actions you can perform.

Figure 84 Virtual Drive Parameters Screen



15-1197_20111017_Ctrl-M_49_Objects_VD_options6

4. Select the action you want to perform and press Enter.

The options are:

- **Make Hot Spare**

Only a drive in *Ready* state can be changed to a hot spare.

- **Force Online**

Only a drive in *Failed* state can be forced online.

- **Change Drv State**

Change an online (*Onlin*) drive to *Failed* state or a hot spare to *Ready* state.

- **Drive Properties**

Display the device type, capacity, manufacturer name and model, and revision number.

3.12.3 Viewing or Changing Virtual Drive Properties

You can disable the write cache option when you create a virtual drive, but you can enable this option later using the configuration utilities. When the disk Write Cache option is **On**, a write transaction is considered to be complete when all the data has been written to the disk cache. When the disk Write Cache option is **Off**, the write transaction is complete only when the data has been written to the disk.

When the disk Read Ahead option is **On**, extra data is read sequentially ahead of the data that is actually requested, and this extra data is stored in a cache. If the additional read-ahead data is then requested, it can be read faster from the cache than from the disk directly.

NOTE When the disk Write Cache option is **On**, a danger that data could be lost exists if the power fails before the cached data is written to disk.

Follow these steps to view the disk Write Cache or Read Ahead settings, or enable the disk Write Cache setting:

1. On the **Management Menu** screen, select **Objects > Virtual Drive**.
The list of virtual drives appears.
2. Highlight a virtual drive with your cursor, and press Enter.
The list of virtual drive parameters appears.
3. Highlight **Disk WC** with your cursor, and press Enter.
The **On/Off** dialog box opens.
4. Change the setting to **On** to enable the disk Write Cache setting.
The disk Write Cache option is enabled.

3.13 Forcing Drives Online or Offline

To force a drive online or offline, follow these steps:

1. On the **Management Menu**, select **Objects > Physical Drive**.
2. Highlight a physical drive that is a member of an drive group, and press Enter.
3. Select **Force Offline** or **Force Online** from the menu.
 - If the drive was online, its status changes to *FAIL*.
 - If the drive was offline, its status changes to *ONLINE*.

3.14 Sector Size Not Supported

The Avago Software RAID CU does not support 4K sector size drives. If 4K drives are plugged into the system, the software RAID CU detects the 4K sector size drive and displays an *Unsupported* message to the user.

Chapter 4: Human Interface Infrastructure Configuration Utility

The Human Interface Infrastructure (HII) Configuration Utility (CU) is used to configure controllers, drive groups, and virtual drives, and to perform other configuration tasks in a pre-boot environment.

The HII CU supports the hot plug feature. You can add and remove devices to a computer while the computer is running and the operating software automatically recognizes the change.

This chapter describes how to configure controllers, drive groups, and virtual drives with the HII CU. To ensure the best performance, select the optimal RAID level for the virtual drive you create. For an explanation of RAID levels, see [Section 1.2, RAID Overview](#), on page 8.

4.1 Managing Controllers

This section explains how you can use the HII CU to view and change the properties for controllers, and clear a configuration.

4.1.1 Viewing Controller Properties

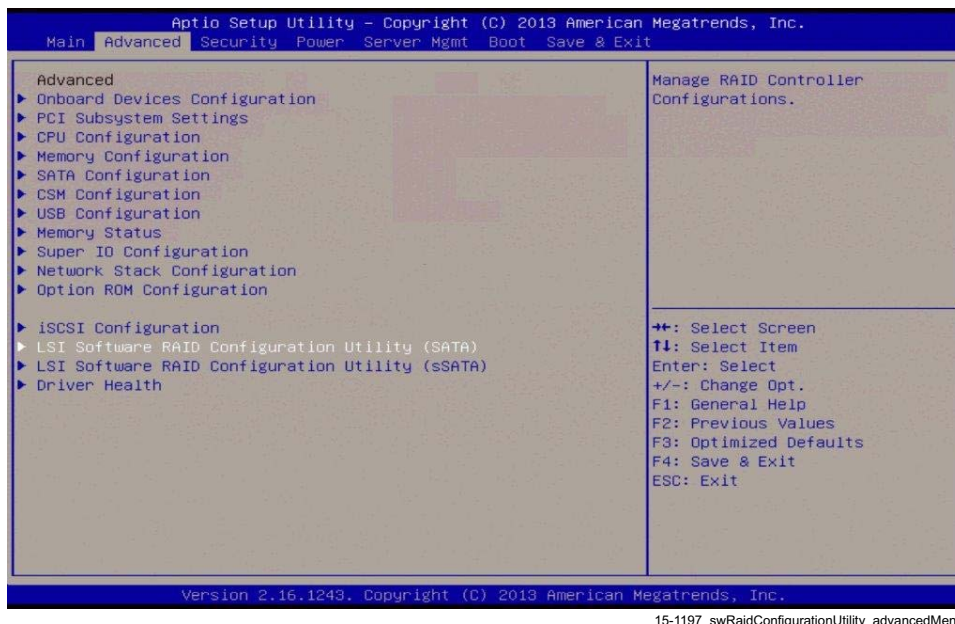
The HII CU displays information for SATA and secondary SATA (sSATA) controllers.

Perform the following steps to view the controller properties.

1. Choose the **Avago Software RAID Configuration Utility** for SATA or sSATA controllers from the **Advanced** tab in the **Setup Utility** and press Enter.


NOTE If the **Avago Software RAID Configuration Utility** is not visible under the **Advanced tab**, refer to the *OEM User Guide*.

Figure 85 Select the Avago Software RAID Configuration Utility



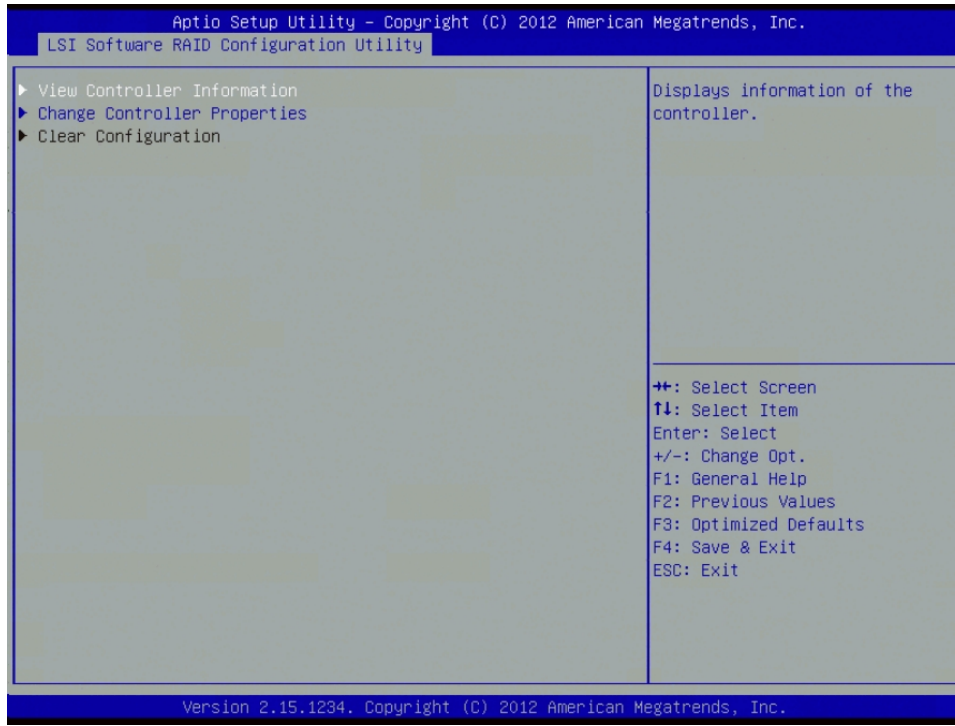
2. Highlight **Controller Management** on the **Configuration Options** screen and press Enter.

Figure 86 RAID Configuration Options Screen



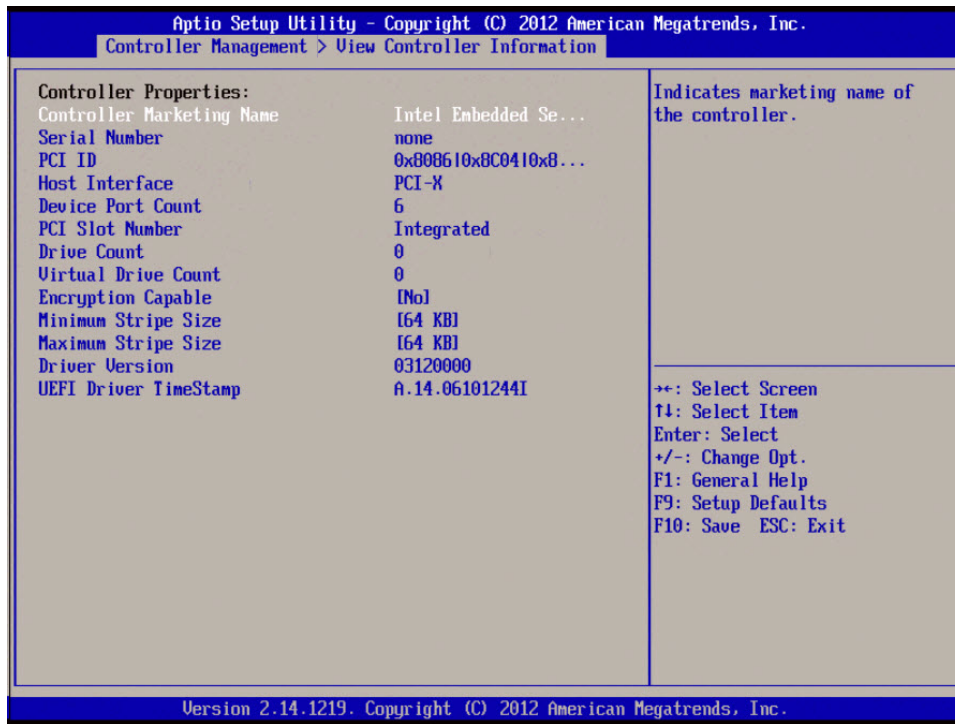
The **Controller Management** screen appears, as shown in the following figure.

Figure 87 Controller Management Screen



3. Highlight **View Controller Information** and press Enter.
The following screen appears.

Figure 88 Controller Management – View Controller Information Screen



The information on this screen is read-only. This screen presents basic information, such as the number of virtual drives that are defined on this controller and the number of drives connected to the controller.

The following table defines the controller properties.

Table 3 Controller Properties

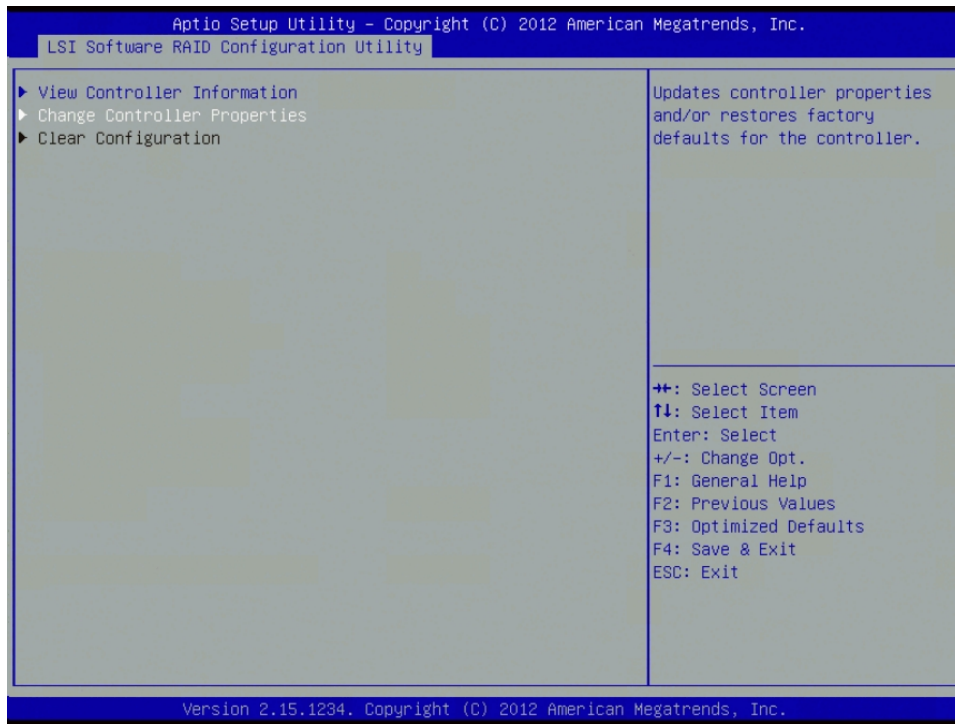
Property	Description
Controller Marketing Name	Embedded MegaRAID
Serial Number	The manufacturer-assigned serial number.
PCI ID	The ID number for the Peripheral Component Interconnect (PCI) local bus.
Host Interface	The type of interface used by the computer host system, such as PCI Express® (PCIe®).
Device Port Count	The maximum number of ports supported by the software RAID controller in which devices (such as CD-ROM and disks) can be connected.
PCI Slot Number	The number of the PCI slot in which the selected controller is installed.
Drive Count	The number of drives connected to the selected controller.
Virtual Drive Count	The number of virtual drives configured on the controller currently.
Encryption Capable	Indicates whether the controller offers the ability to encrypt data on the drives. This solution provides data protection in the event of theft or loss of physical drives.
Minimum Stripe Size	The minimum length of the data segments that the controller writes across multiple drives, not including the parity drives. The default minimum stripe size is 64 KB.
Maximum Stripe Size	The maximum length of the data segments that the controller writes across multiple drives, not including the parity drives. The default maximum stripe size is 64 KB.
Driver Version	The driver version of the EFI driver.
UEFI Driver Timestamp	The UEFI driver compilation time stamp.

4.1.2 Changing Controller Properties

You can use the HII CU to change the properties for a controller. Perform the following steps to change information for a controller.

1. Highlight **Controller Management** on the **Configuration Options** screen (Figure 86) and press Enter.
The following screen appears.

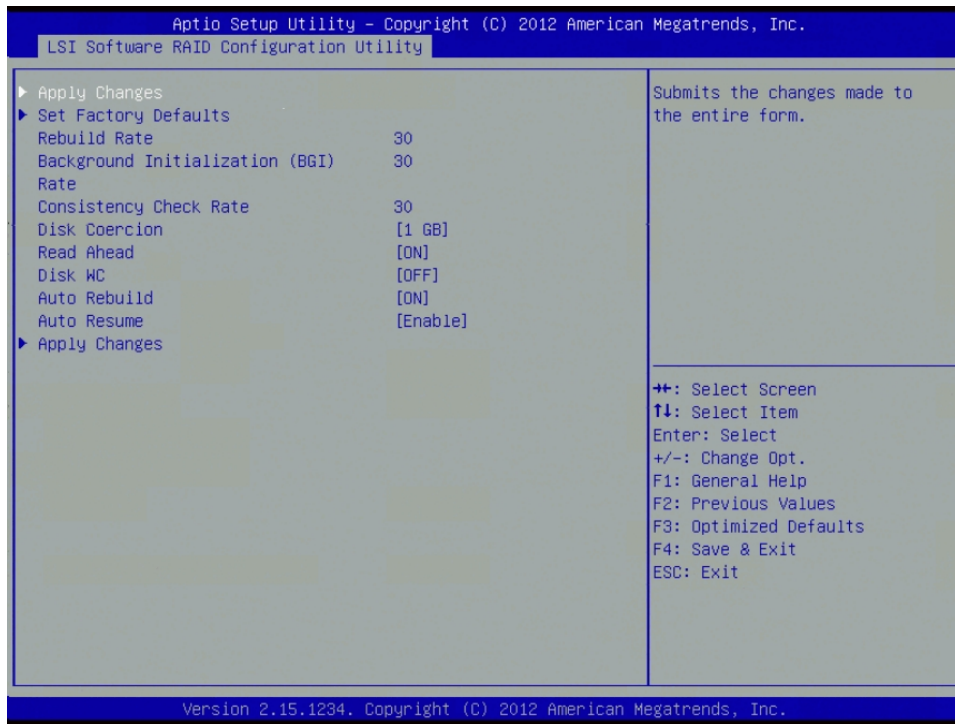
Figure 89 Controller Management Screen



15-1197_swRaidConfigUtility_changeControllerProperties

2. Highlight **Change Controller Properties** and press Enter.
The following screen appears.

Figure 90 Controller Management – Change Controller Properties Screen



3. Change the controller properties as desired.
 The following table defines these controller properties.

Table 4 Changing Controller Properties

Property	Use	Description
Set Factory Defaults	Use the arrow keys to move the cursor to this property and press Enter. On the Confirm screen, select Yes to confirm your selection.	Resets factory default values for all of the controller properties.
Set Boot Devices		Selects the virtual drive to use as the boot device.
Rebuild Rate	Use the arrow keys to move the cursor to this property. Press the plus key (+) to increase the rate or press the minus key (-) to decrease the 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. The default value is 30 percent.
Background Initialization (BGI) Rate	Use the arrow keys to move the cursor to this property. Press the plus key (+) to increase the rate or press the minus key (-) to decrease the rate.	Background initialization is a check for media errors on the drives when you create a virtual drive. It is an automatic operation that starts five minutes after you create the virtual drive. This check ensures that striped data segments are the same on all of the drives in the drive group. The default value is 30 percent.
Consistency Check Rate	Use the arrow keys to move the cursor to this property. Press the plus key (+) to increase the rate or press the minus key (-) to decrease the rate.	A consistency check is an operation used to verify that all stripes in a virtual drive with a redundant RAID level are consistent and that it automatically fixes any errors. The consistency check rate is the rate at which consistency check operations are run on a computer system. The default value is 30 percent.

Table 4 Changing Controller Properties (Continued)

Property	Use	Description
Disk Coercion	Use the arrow keys to move the cursor to this property. Press the plus key (+) to set the value to 128 MB , 1 GB , or None .	Drive coercion is a tool to force drives of varying capacities to the same capacity so they can be used in a drive group. The coercion mode options are None , 128MB-way , and 1GB-way . The number you choose depends on how much the drives from various vendors vary in their actual size.
Read Ahead	Use the arrow keys to Enable or Disable .	When the disk Read Ahead option is On , extra data is read sequentially ahead of the data that is actually requested, and this extra data is stored in cache memory. If the additional read-ahead data is then requested, it can be read faster from the cache than from the disk directly. This setting speeds up read operations for sequential data, but there is little improvement when accessing random data.
Disk WC	Use the arrow keys to Enable or Disable .	You can disable the disk Write Cache option when you create a virtual drive, but you can enable this option later using the configuration utilities. <ul style="list-style-type: none"> ■ When the disk Write Cache option is On, a write transaction is considered complete when all the data has been written to the disk cache. ■ When the disk Write Cache option is Off, the write transaction is complete only when the data has been written to the disk.
Auto Rebuild	Use the arrow keys to Enable or Disable .	Auto-rebuild allows a failed drive to be replaced and the data automatically rebuilt by <i>hot-swapping</i> the drive in the same drive bay. The RAID drive group continues to handle requests while the rebuild occurs.
Auto Resume	Use the arrow keys to Enable or Disable .	When enabled, you can stop a <i>consistency check</i> , <i>rebuild</i> , or <i>initialization</i> operation, and resume it later where it left off, instead of aborting it and starting over.

4. Highlight **Apply Changes** and press Enter to register your changes.
 Your changes are registered and the controller properties are changed.

4.1.3 Clearing Configurations

Perform the following steps to clear all existing configurations on the controller.

1. Highlight **Controller Management** on the **Configuration Options** screen (Figure 86) and press Enter.
2. Highlight **Clear Configuration** on the **Controller Management** screen (Figure 87) and press Enter.
3. Highlight **Yes** to confirm your selection on the **Confirm** screen.
 This action clears the existing configurations.

4.2 Managing Virtual Drives

You can use the **Virtual Drive Management** screen to create virtual drive configurations, view and change the virtual drive properties, delete virtual drives, initialize drives, and perform other tasks related to the virtual drives.

NOTE You cannot create a configuration across SATA and sSATA controllers.

4.2.1 Configuring Virtual Drives

This section provides detailed instructions to configure drive groups and virtual drives with the HII CU.

NOTE Use drives with the same capacity when you create a storage configuration. If you use drives with different capacities in the same drive group, the CU limits each drive to the capacity of the smallest drive.

The number of physical drives in a specific array determines the possible RAID levels that you can implement with the array.

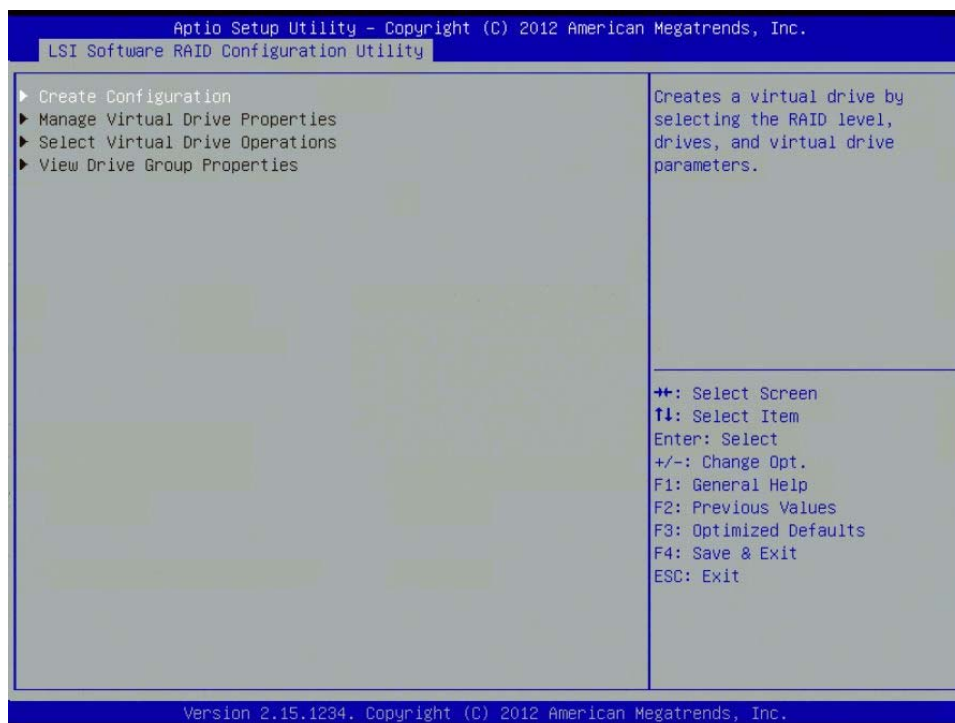
- RAID 0 requires from one to eight physical drives.
- RAID 1 requires two physical drives.
- RAID 5 required three to eight physical drives.
- RAID 10 requires four, six, or eight physical drives.

NOTE The stripe size is read-only. You cannot change the settings for this property.

Follow these steps to access the **Virtual Drive Management** screen and create a virtual drive configuration.

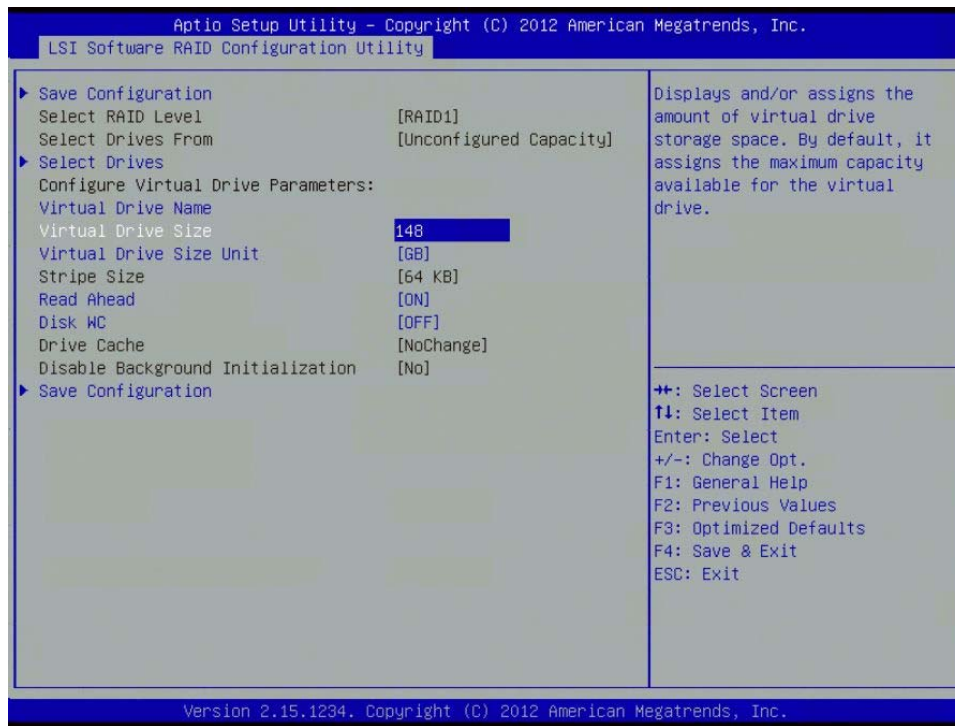
1. Highlight **Virtual Drive Management** on the **Configuration Options** screen (Figure 86) and press Enter. The following screen appears.

Figure 91 Virtual Drive Management Screen



2. Highlight **Create Configuration** and press Enter. The following screen appears.

Figure 92 Virtual Drive Management – Create Configuration Screen



3. Use the arrow keys to select any highlighted fields (one at a time) for which you want to change the setting, and press Enter.
4. Select the setting for each property that you want to change from the default.
 You can change the settings for the following fields on this screen:

Table 5 Virtual Drive Management Property Settings

Property	Description
Select RAID Level	The possible RAID levels for the virtual drive. See Section 1.2, RAID Overview , on page 8, for more information about the RAID levels.
Select Drives From	The sources that you can use to select drives for the virtual drive. The options are Unconfigured Capacity and Free Capacity .
Select Drives	Select this button and a screen appears that lists Unconfigured Good drives or free capacity , which depend on the value you selected in the Select Drive From field.
Virtual Drive Name	Enter the name of the virtual drive.
Virtual Drive Size	Enter the capacity of the virtual drive. Normally, this value is the full capacity of the drive. You can specify a smaller capacity if you want to create other virtual drives on the same drive group.
Virtual Drive Size Unit	Enter the unit of capacity you want to use for the virtual drive. The options are MB, GB, and TB.
Stripe Size	A stripe consists of the data segments that the RAID controller writes across multiple drives, not including parity drives. The default is 64 KB . This setting is read-only.

Table 5 Virtual Drive Management Property Settings (Continued)

Property	Description
Read Ahead	When the disk <code>Read Ahead</code> option is On , extra data is read sequentially ahead of the data that is actually requested, and this extra data is stored in cache memory. If the additional read-ahead data is then requested, it can be read faster from the cache than from the disk directly. This setting speeds up reads for sequential data, but there is little improvement when accessing random data.
Disk WC	You can disable the disk <code>Write Cache</code> option when you create a virtual drive, but you can enable this option later using the configuration utilities. <ul style="list-style-type: none">■ When the disk <code>Write Cache</code> option is On, a write transaction is considered to be complete when all the data has been written to the disk cache.■ When the disk <code>Write Cache</code> option is Off, the write transaction is complete only when the data has been written to the disk.
Disable Background Initialization	Use this option to select the amount of system resources dedicated to background initialization of virtual drives connected to the selected controller.

4.2.2 Managing Virtual Drive Properties

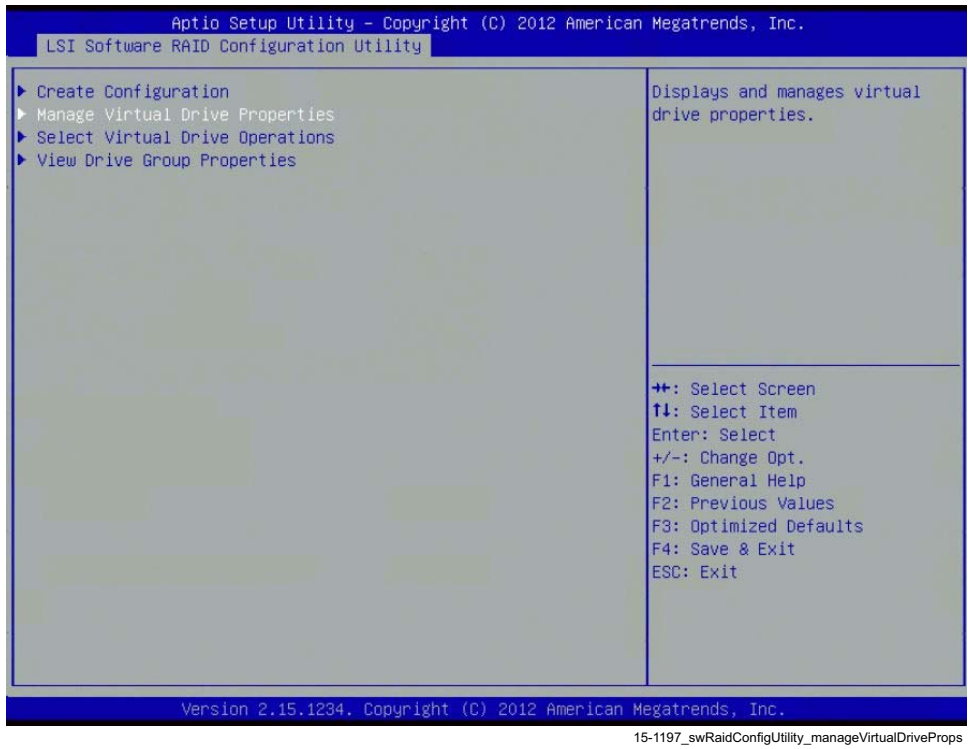
After you create a virtual drive, you can use the **Virtual Drive Management** screen to change the name of the virtual drive.

NOTE You can change only the virtual drive name for the virtual drive that is selected. The other virtual drive properties are read-only.

Perform the following steps to access the virtual drive properties screen and change the name of the virtual drive.

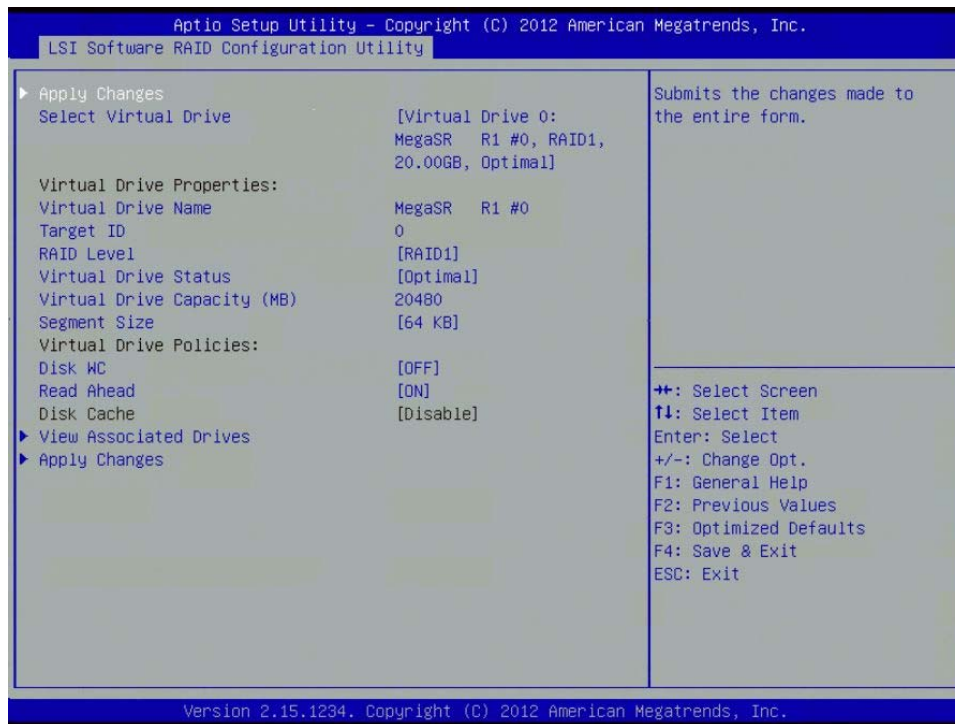
1. Highlight **Virtual Drive Management** on the **Configuration Options** screen ([Figure 86](#)) and press Enter.
The following screen appears.

Figure 93 Virtual Drive Management Screen



2. Highlight **Manage Virtual Drive Properties** and press Enter.
The following screen appears.

Figure 94 Virtual Drive Management – Manage Virtual Drive Properties Screen



15-1197_swRaidConfigUtility_applyChanges_manageVirtualDriveProps

3. Highlight the **Virtual Drive Name** field and press Enter.
4. Enter a different name for the virtual drive and press Enter.
5. Highlight **Apply Changes** and press Enter to make the selected change.
The name of the virtual drive is changed.

4.2.3 Selecting Virtual Drive Operations

You can use the **Virtual Drive Management** screen to select and launch operations related to the virtual drive properties.

Perform the following steps to access the virtual drive operations screen and carry out operations for the virtual drive properties.

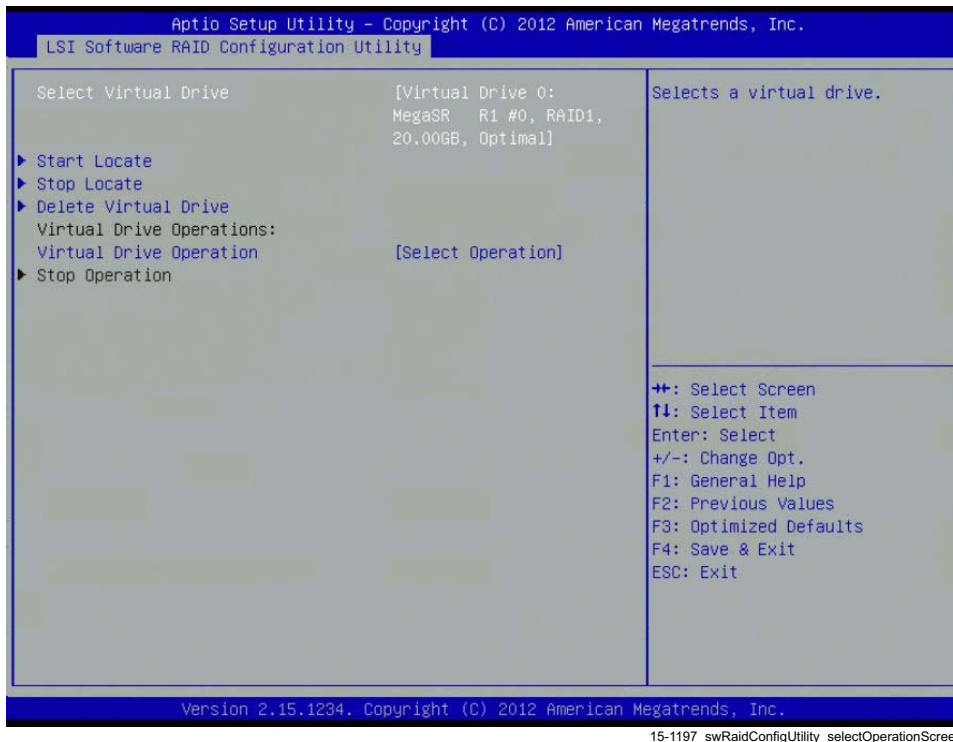
1. Highlight **Select Virtual Drive Operations** on the **Configuration Options** screen (Figure 86) and press Enter.
The following screen appears.

Figure 95 Virtual Drive Management Screen



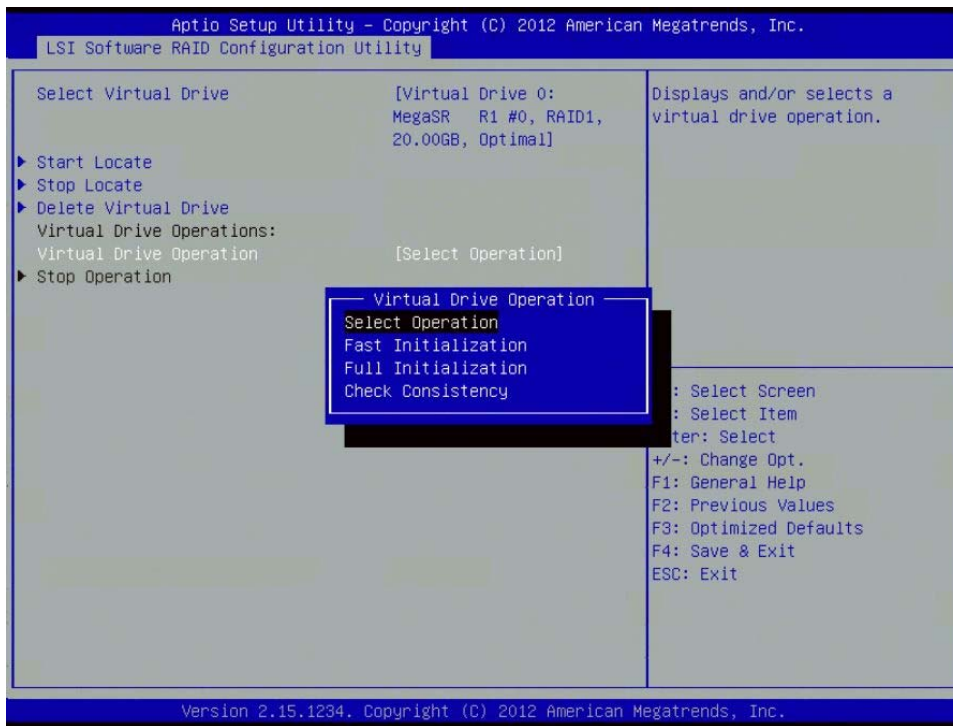
The following screen appears.

Figure 96 Select Virtual Drive Operation



- Highlight **Select Operation** and press Enter.
 The following screen appears.

Figure 97 Select Virtual Drive Operations Screen



- Highlight each operation you want to carry out and press Enter.
 The following two operations appear.

Table 6 Select Virtual Drive Operation Properties

Property	Description
Select Virtual Drive	Select the virtual drive you want to carry out operations on. The screen displays the virtual drive number, virtual drive name, RAID level, virtual drive capacity, and virtual drive status.
Virtual Drive Operation	Select the Fast Initialization or Full Initialization option to initialize this virtual drive. A fast initialization quickly writes 0s to the first and last 10-MB regions of the new virtual drive and then completes the initialization in the background. A full (slow) initialization is not complete until the entire virtual drive has been initialized with 0s. It is seldom necessary to use this option, because the virtual drive was initialized when you created it. NOTE If you have redundant virtual drives, you have the option to perform the Make Data Consistent operation.

- Highlight **Start Operation** and press Enter.

The confirmation screen appears.

NOTE The `Start Operation` option and the `Stop Operation` option toggle is based on the current status of the operation. For example, after you start an operation, the **Start Operation** field is no longer highlighted, and the **Stop Operation** field is highlighted.

5. Highlight **Yes** and press Enter to confirm that you want to carry out the operations.

NOTE After you choose an operation and select `Start Operation`, the progress bar appears, showing zero percent. To see the change in percentages, return to the previous page and then access this page again.

4.2.4 Viewing Drive Group Properties

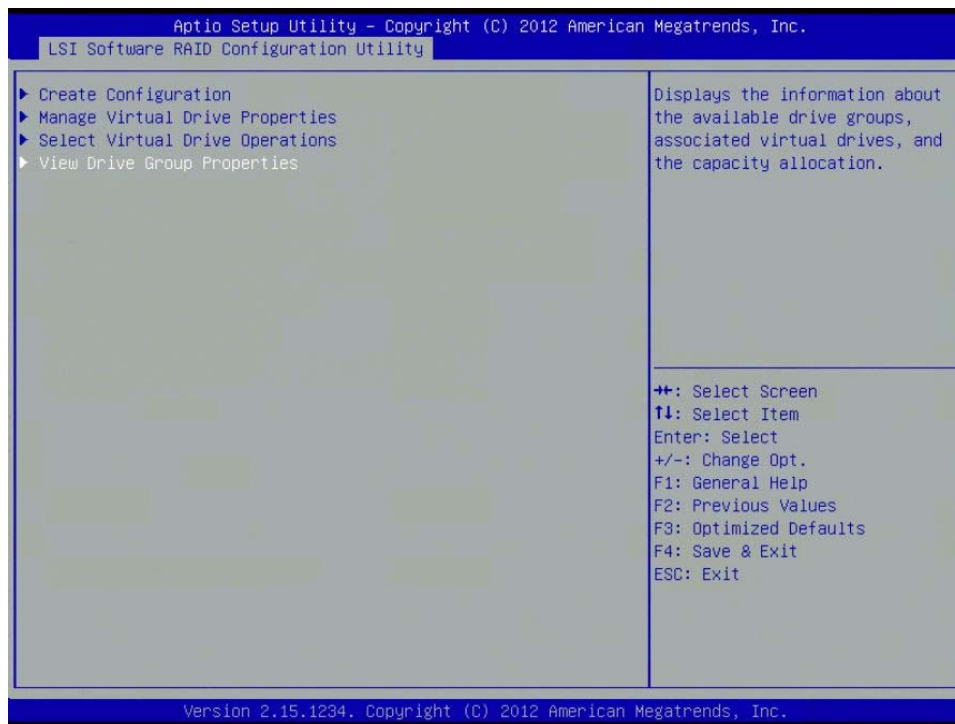
You can use the **Virtual Drive Management** screen to access the **View Drive Group Properties** screen and view the properties. The drive group contains the drives that are used to create the virtual drive.

Perform the following steps to access the **View Drive Group Properties** screen and view the drive group properties.

1. Highlight **Drive Management** on the **Configuration Options** screen (Figure 86) and press Enter.

The following screen appears.

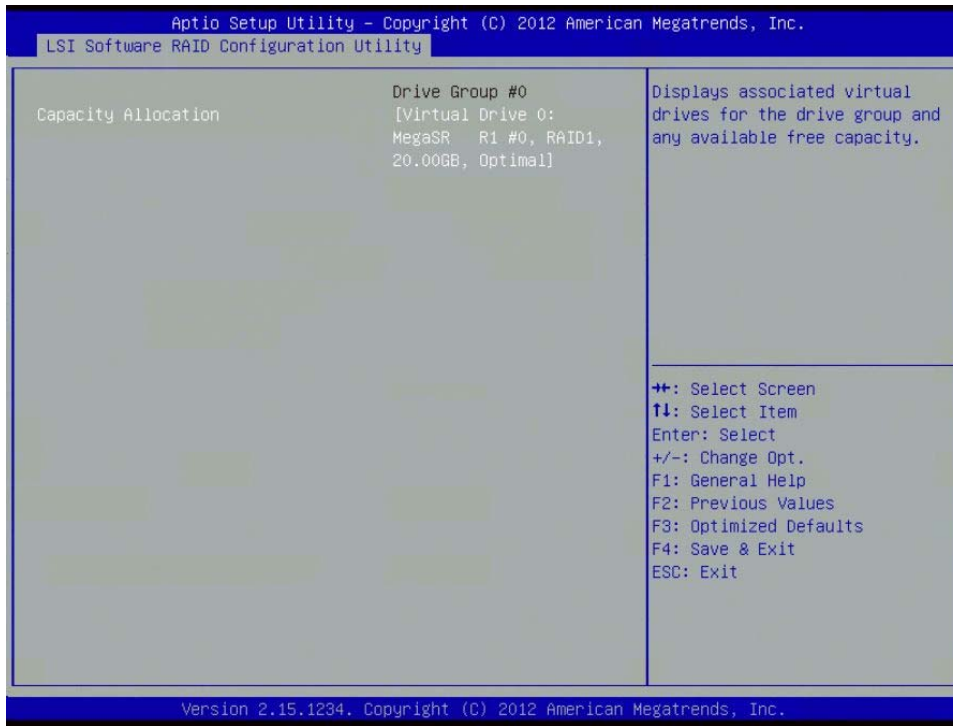
Figure 98 Virtual Drive Management Screen



2. Select **View Drive Properties** and press Enter.

The following screen appears.

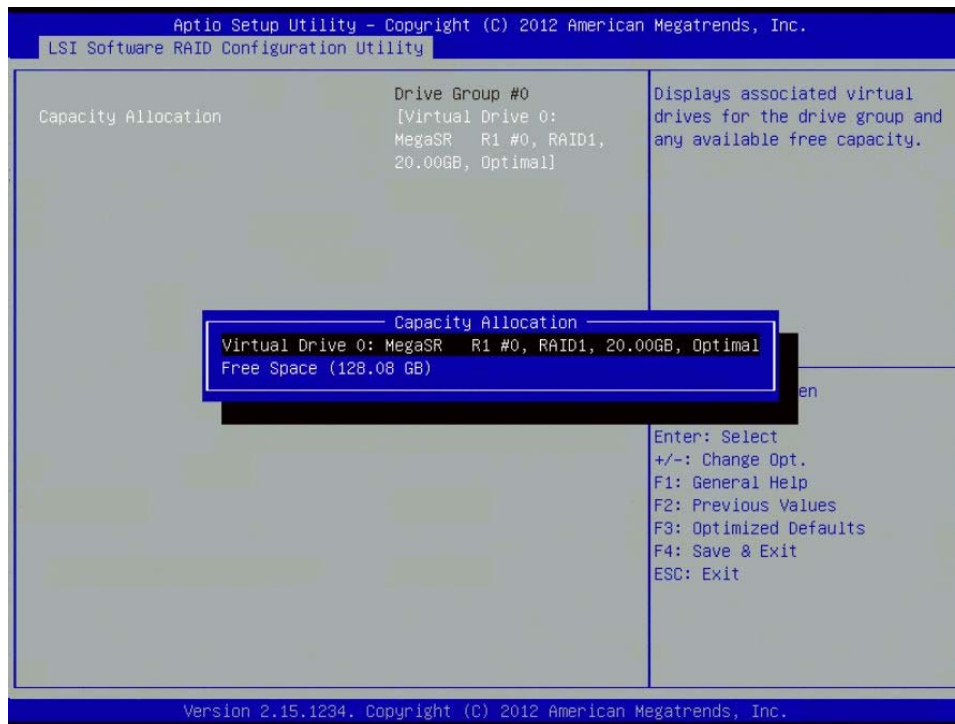
Figure 99 View Drive Group Properties Screen



The screen displays the drive group number, drive group name, RAID level, virtual drive capacity, and virtual drive status.

3. Highlight the drive group and press Enter.
The following screen appears.

Figure 100 Drive Group Properties Screen and Free Space



4.3 Managing Drives

You can use the **Drive Management** screen to view and change the physical drive properties of the drives in a drive group.

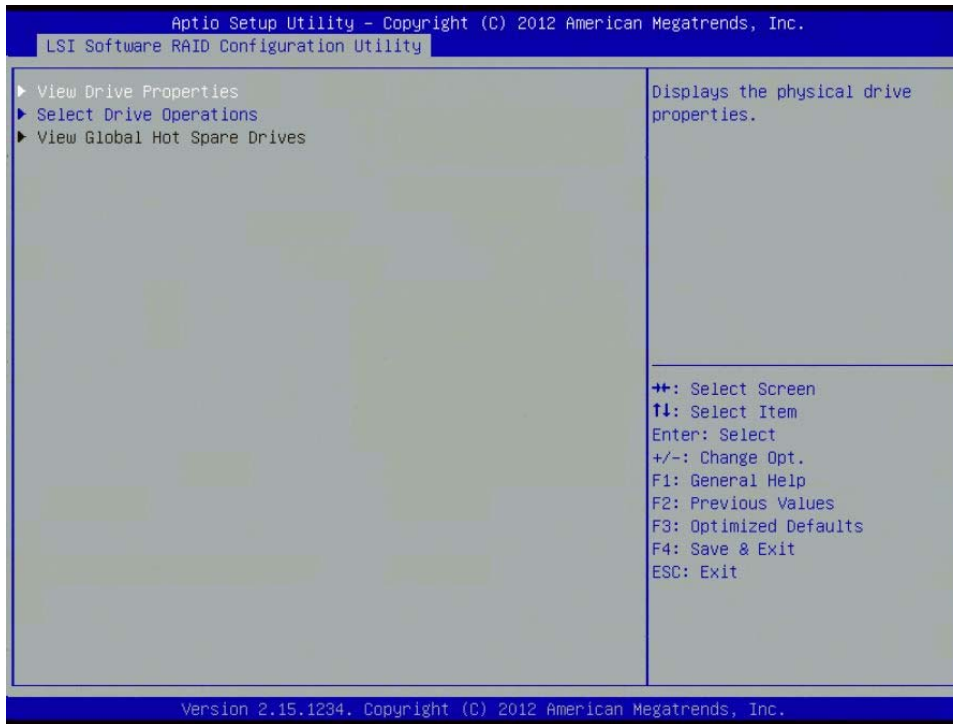
4.3.1 Viewing Drive Properties

Perform the following steps to access the **View Drive Properties** screen and view the values for the properties.

1. Highlight **Drive Management** on the **Configuration Options** screen (Figure 86) and press Enter.

The following screen appears.

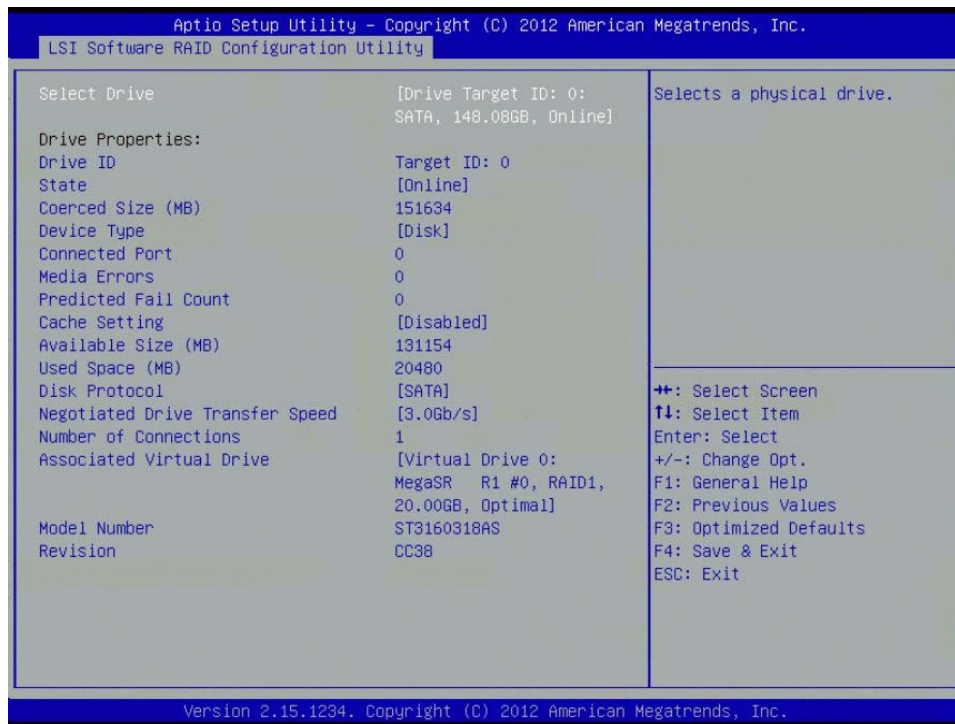
Figure 101 Drive Management Screen



2. Highlight **View Drive Properties** and press Enter.

The following screen appears. This screen contains additional options that you can scroll down to view. The following figure shows the first options. You can scroll down to display the additional options.

Figure 102 Drive Management – View Drive Properties Screen



You can highlight the **Associated Virtual Drive** field on the second **View Drive Properties** screen, and press Enter to view the properties. However, you cannot change the values in this field.

The following table shows the drive properties and their values.

Table 7 Drive Properties

Property	Value
Select Drive	The drive number, drive slot, drive type, drive capacity, and drive status of the selected drive.
Drive ID	The ID and the drive slot of the selected drive.
State	The state of the selected drive, such as <i>Online</i> , <i>Offline</i> , or <i>Unconfigured Good</i> .
Coerced Size (MB)	The capacity to which the selected drive has been coerced to make it compatible with other drives that are nominally the same capacity.
Revision	The firmware revision of the drive.
Device Type	Indicates the type of device selected, such as a drive, tape, or CD/DVD-ROM.
Connected Port	Indicates the port to which the selected drive is connected.
Media Errors	Indicates the number of media errors on the drive. Media errors are physical defects on the drive.
Predicted Fail Count	Indicates the predicted number of drive failures.
Available Size (MB)	Indicates the free capacity of the selected drive.
Used Space (MB)	Indicates the configured space of the selected drive in MB.
Disk Protocol	Indicates the type of drive selected, such as SATA.

Table 7 Drive Properties (Continued)

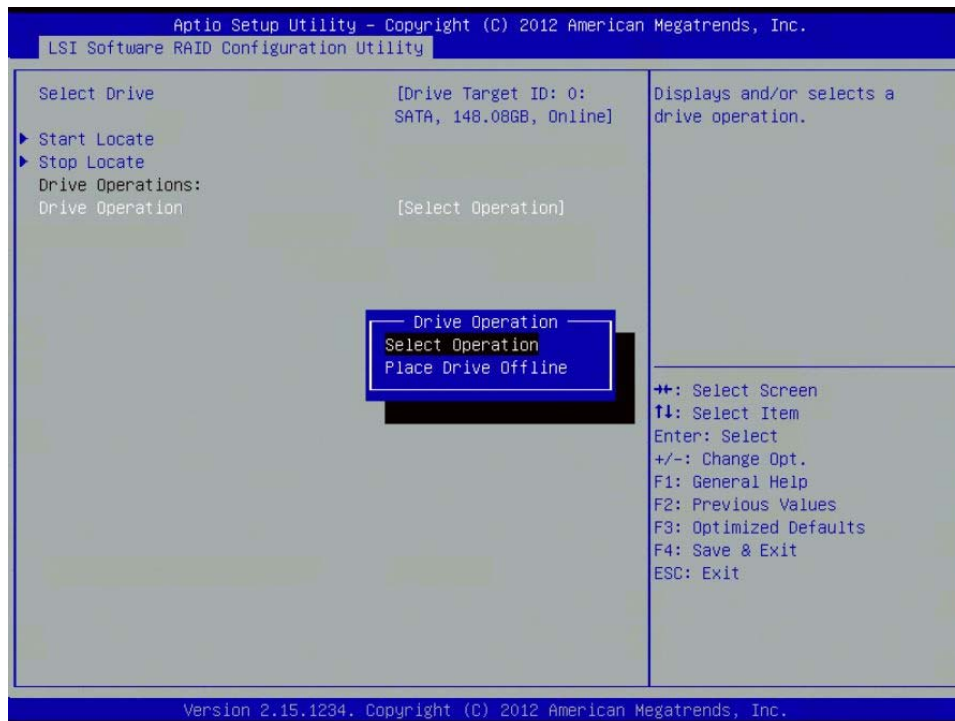
Property	Value
Negotiated Drive Transfer Speed	Indicates the negotiated link speed for the data transfer to or from the selected drive.
Number of Connections	Indicates the number of devices connected.
Associated Virtual Drive	Indicates the virtual drive number, virtual drive name, RAID level, virtual drive capacity, and virtual drive status.

4.3.2 Selecting Drive Operations

Perform the following steps to access the drive operations screen and carry out the operations.

1. Highlight **Drive Management** on the **Configuration Options** screen (Figure 86) and press **Enter**.
 The **Drive Management > Select** screen appears, as shown in Figure 101, *Drive Management Screen*.
2. Highlight **Select Drive Operations** and press Enter.
 The following screen appears.

Figure 103 Drive Management – Select Drive Operations



15-1197_swRaidConfigUtility_driveMgt_selectDriveOperations

3. Highlight an operation and press Enter.

The following table defines the drive operations:

Table 8 Drive Operation Settings

Property	Value
Select Drive	Select the drive on which you want to carry out operations. The screen displays the drive number, slot number, drive type, drive capacity, and drive status.
Drive Operation	Make a drive offline or online. The options are <code>Place Drive Offline</code> and <code>Place Drive Online</code> . NOTE If a good drive is part of a redundant drive group with a hot spare, and you force that drive offline, the data on the drive rebuilds to the hot spare drive. The drive you forced offline goes into the <i>Unconfigured Good</i> state.

4. Highlight **Start Operation** and press Enter.
The confirmation screen appears.
5. Highlight **Yes** and press Enter.
The controller performs the selected operations.

4.4 UDK2010 Support

The HII utility works with both EDK I (EDK 1.06) and EDK II (UDK2010). The EDK is the Intel implementation of the EFI Specification and EDK II (UDK2010) is a cross-platform firmware development environment for the UEFI.

4.5 4K Sector Size Support

The HII utility can be used for both native 512B sector size and 4KB sector sized drives.

Chapter 5: MegaRAID Storage Manager Utility

The MegaRAID Storage Manager software is a configuration and monitoring utility used with the Embedded MegaRAID software. This chapter provides a brief overview of the MegaRAID Storage Manager and explains how to install it on the supported operating systems.

5.1 Overview

The MegaRAID Storage Manager utility can be installed either on the local system that contains the MegaRAID controllers you want to monitor, or on a remote system such as a workstation or a laptop with network connectivity to the system you want to monitor. Because of this, the system hardware and software requirements to install the MegaRAID Storage Manager utility might differ from the system requirements for the MegaRAID controller(s) you are using.

The MegaRAID Storage Manager utility lets you configure, monitor, and maintain storage configurations created under the Embedded RAID utility. The MegaRAID Storage Manager utility graphical user interface (GUI) makes it easy for you to create and manage storage configurations.

NOTE The MegaRAID Storage Manager utility can manage a wide range of MegaRAID controllers. Some MegaRAID Storage Manager utility features are not applicable for Embedded RAID.

5.1.1 Creating Storage Configurations

The MegaRAID Storage Manager utility lets you easily configure the controllers, disk drives, and virtual disks on your workstation or server. The Simple Configuration mode and the Advanced Configuration mode greatly simplify the process of creating drive groups and virtual drives. You can use these configuration modes to create the best possible configuration with the available hardware.

You can create the following types of configurations:

- **Simple configuration** mode
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** mode
Lets you choose additional settings and customize virtual drive creation. This option provides greater flexibility when creating virtual drives for your specific requirements.

See [Section 6.1, Creating a New Storage Configuration](#), on page 123, for the procedures used to create storage configurations.

5.1.2 Monitoring Storage Devices

The MegaRAID Storage Manager utility displays the status of virtual disks, physical disks, and other storage devices 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 disk failures and other events that require immediate attention.

5.1.3 Maintaining Storage Configurations

You can use the MegaRAID Storage Manager utility to perform system maintenance tasks such as running consistency checks on arrays that support redundancy.

5.2 Hardware and Software Requirements

The MegaRAID Storage Manager utility has the following hardware requirements:

- PC-compatible computer with an IA-32 (32-bit Intel Architecture) processor or an EM64T (64-bit) processor and at least 128 MB of system memory

NOTE Use at least 1 GB of system memory.

- Physical drive with at least 50 MB of available free space

The MegaRAID Storage Manager utility supports these operating systems:

- Microsoft Windows Server 2003, 2008, 2008R2, 2012
- Microsoft Windows Workstation 7
- Microsoft Windows Vista Workstation
- Red Hat Enterprise Linux (RHEL) 5 and 6
- SuSE Linux Enterprise Server (SLES) 10 and 11, with latest updates and service packs

NOTE The MegaRAID Storage Manager utility supports the RHEL 6 operating system (OS) with a few prerequisites. By nature, the RHEL 6 OS default installation deploys only limited inbox libraries though it is shipped in the CD image for the full installation. This design from the RHEL 6 OS adds an overhead for the Java® application programming users, and a few required libraries are not deployed during the default RHEL 6 operating system installation. Because of this, users need to port the application in to the native 64-bit version.

Refer to your server documentation and to the operating system documentation for more information on hardware requirements and operating system requirements.

5.3 Installation

This section explains how to install (or reinstall) the MegaRAID Storage Manager utility on your workstation or server.

5.3.1 Installing the MegaRAID Storage Manager Utility on Microsoft Windows


Follow these steps to install the MegaRAID Storage Manager utility on a system running a Microsoft Windows operating system:

1. Unzip the ZIP file that contains the MegaRAID Storage Manager utility in Windows Explorer.
2. Double-click the `setup.exe` file to start the installation program.
3. When the **Welcome** screen appears, click **Next**.

If MegaRAID Storage Manager utility is already installed on this system, the **Program Maintenance** screen appears. Read the screen text and select **Modify**, **Repair**, or **Remove**.

4. When the next screen appears, read and accept the user license, and click **Next**.
The following screen appears.


Figure 104 Customer Information Screen



15-1197_msmUtility_customerInfoScreen

5. 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 the MegaRAID Storage Manager utility to view or change storage configurations.
 - If you select **Only for current user**, the MegaRAID Storage Manager utility shortcuts and associated icons are available only to the user with this user name.
6. Click **Next** to continue.
7. On the next screen, accept the default destination folder, or click **Change** to select a different destination folder.
8. Click **Next** to continue.
The following screen appears.

Figure 105 Setup Type Screen



15-1197_msmUtility_setupTypeScreen

9. Select one of the setup options.

The screen text explains the options.

- Normally, you would select **Complete** if you are installing the MegaRAID Storage Manager utility on a server. This option installs the complete MegaRAID Storage Manager utility.
- Select **Client** if you are installing the MegaRAID Storage Manager utility on a PC that will be used to view and configure servers over a network.
 - The *Master GUI*, *Monitor Configurator*, *Help* files for both the MegaRAID Storage Manager utility and *Monitor Configurator*, and *DebugLog* features are available.
 - The *Popup*, *SNMP*, *Monitor*, *Framework*, *Storelib* and *Storelib-IR*, *StorelibJNI*, and *StorelibIRJNI* features are not available.
- Select **Server** to install only those components required for remote server management.
 - The *Popup*, *SNMP*, *Monitor*, *Framework*, *Storelib* and *Storelib-IR*, *StorelibJNI*, and *StorelibIRJNI* features are available.
 - The *Master GUI*, *Monitor Configurator*, *Help* files for both the MegaRAID Storage Manager utility and *Monitor Configurator*, and *DebugLog* features are not available.
- Select **StandAlone** if you will use MegaRAID Storage Manager utility to create and manage storage configurations on a standalone workstation. All the components are available except for the Network Capability Plugin module.
- Select **Custom** to specify individual features to install. This option lets you select or omit program components.

10. Click **Next** to proceed.

11. Click **Install** to install the program.

12. Click **Finish** when the final **Configuration Wizard** screen appears.

If you select the **Client** installation for a PC 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 utility), you cannot connect to a remote server unless you first edit the `startupui.bat` file. Specifically, you must add the IP address of the remote server to the end of the `startupui.bat` file.

For example, to connect to a remote framework on server 192.168.0.10, add the IP address to the end of the `startupui.bat` file as shown in this example:

```
start JRE\bin\javaw -classpath .;GUI.jar GUI.VivaldiStartupDialog  
ajsgyqkj=71244 192.168.0.10
```

Be sure to include a space in front of the IP address, as shown in the example.

5.3.2 Installing the MegaRAID Storage Manager Utility for Linux

Follow these steps to install the MegaRAID Storage Manager utility on a system running Red Hat Linux or SuSE Linux:

1. Copy the `MSM_linux_installer...tar.gz` file to a temporary folder.
2. Untar the `MSM_linux_installer...tar.gz` file using the following command:

```
tar -zxvf MSM_linux_installer...tar.gz
```

A new disk directory is created.

3. Go to the new disk directory.
4. Find and read the `readme.txt` file in the disk directory.
5. Enter the following command to start the installation:

```
./install.sh
```

If you select the **Client** installation for a PC 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 utility), you cannot connect to a remote server unless you first edit the `startupui.sh` file. Specifically, you must add the IP address of the remote server to the end of the `startupui.sh` file.

For example, to connect to a remote framework on server 192.168.0.10, add the IP address to `startupui.sh` as shown in this example:

```
start JRE\bin\javaw -classpath .;GUI.jar GUI.VivaldiStartupDialog  
ajsgyqkj=71244 192.168.0.10
```

Be sure to include a space in front of the IP address, as shown in the example.

5.3.2.1 Linux Installation Messages

One or more of the following messages can appear while you are installing the MegaRAID Storage Manager utility on a Linux system:

- More than one copy of the MegaRAID Storage Manager utility has been installed.
This message indicates that the user has installed more than one copy of the MegaRAID Storage Manager utility.

NOTE This action can be accomplished 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 the `rpm` files manually before installing the MegaRAID Storage Manager utility with the procedure listed previously.

- The version is already installed.
This message indicates that the version of the MegaRAID Storage Manager utility 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 utility 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.

Chapter 6: Configuration

This chapter explains how to use the MegaRAID Storage Manager utility to create and modify storage configurations on Avago RAID controllers.

The Avago RAID controllers support RAID 0, RAID 1, RAID 5, and RAID 10 storage configurations. The Configuration Wizard lets you create new storage configurations and modify the configurations. To learn more about RAID and RAID levels, see [Section 1.2, RAID Overview](#), on page 8.

NOTE You cannot create or modify a storage configuration unless you are logged on to a server with administrator privileges.

6.1 Creating a New Storage Configuration

You can use the MegaRAID Storage Manager utility to create new storage configurations on systems with Avago RAID 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.

This section describes the virtual drive parameters and explains how to create simple storage configurations and advanced storage configurations.

6.1.1 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. 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* (Default)
The new configuration is not initialized and the existing data on the drives is not overwritten.
 - *Fast Initialization*
This option lets you start writing data to the virtual drive immediately.
 - *Full Initialization*
A complete initialization is performed on the new configuration.
You cannot write data to the new virtual drive until the initialization is complete. This process can take a long time if the drives are large.

NOTE New RAID 5 virtual drives require at least five drives for a background initialization to start.

- **Stripe size**
The setting is 64 KB only.
- **Read policy**

Specify the read policy for this virtual drive:

- *No read ahead*: (Default)

Disables the read ahead capability.

Read ahead capability allows the controller to read sequentially ahead of the requested data and to store the additional data in cache memory, anticipating that the data will be needed soon. Read ahead capability speeds up read operations for sequential data, but little improvement exists when accessing random data.

■ **Write policy**

Specify the write policy for this virtual drive:

- *Write Through*

In this mode, the controller sends a data transfer completion signal to the host when the drive subsystem has received all the data in a transaction.

This option eliminates the risk of losing cached data in case of a power failure.

- *Always Write Back*

In this mode, the controller sends a data transfer completion signal to the host when the disk cache has received all the data in a transaction.

■ **I/O policy**

The I/O policy applies to read operations on a specific virtual drive. It does not affect the read ahead cache.

- *Direct IO* (Default)

In this mode, read operations 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, and Direct IO makes sure 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* (Default)

Allow read/write access.

- *Read Only*

Allow read-only access.

- *Blocked*

Do not allow access.

■ **Disk cache policy** – Select a cache setting for this drive:

- *Enabled*

Enable the disk cache.

- *Disabled*

Disable the disk cache.

- *Unchanged* (Default)

Leave the current disk cache policy unchanged.

6.1.2 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 RAID 10 spanned drives using the simple configuration procedure. To create RAID 10 spanned drives, use the advanced configuration procedure described in [Section 6.1.3, Creating a Virtual Drive Using Advanced Configuration](#), on page 127.

Follow these steps to create a new storage configuration in simple configuration mode.

1. Click the **Dashboard** tab on the main menu screen.
The following screen appears.

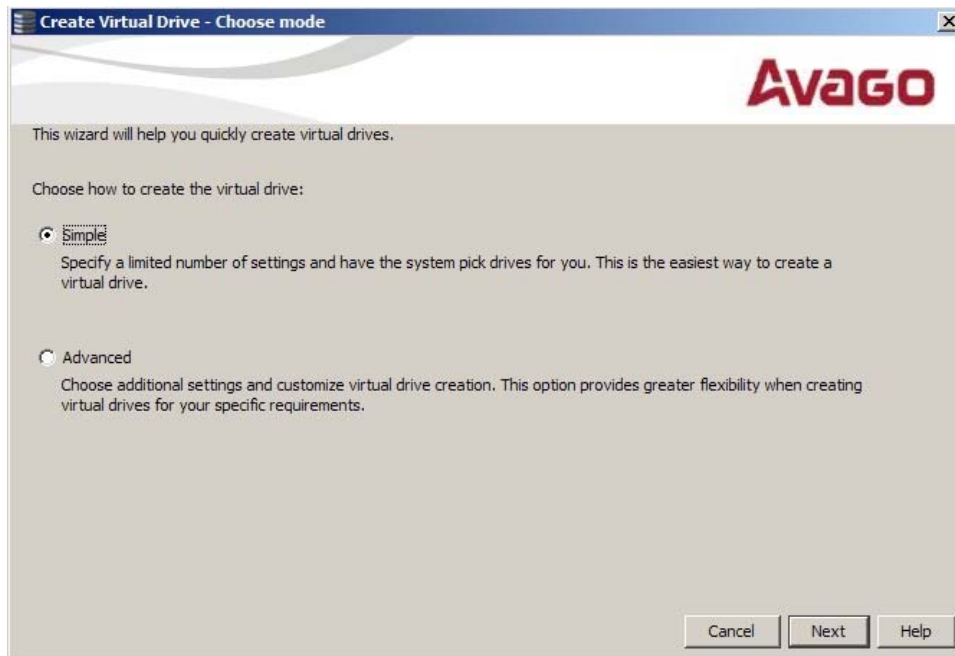
Figure 106 Virtual Drive Creation Menu



15-1197_msmUtility_virtualDriveCreationMenu

2. Click **Create Virtual Drive**.
The following screen appears.

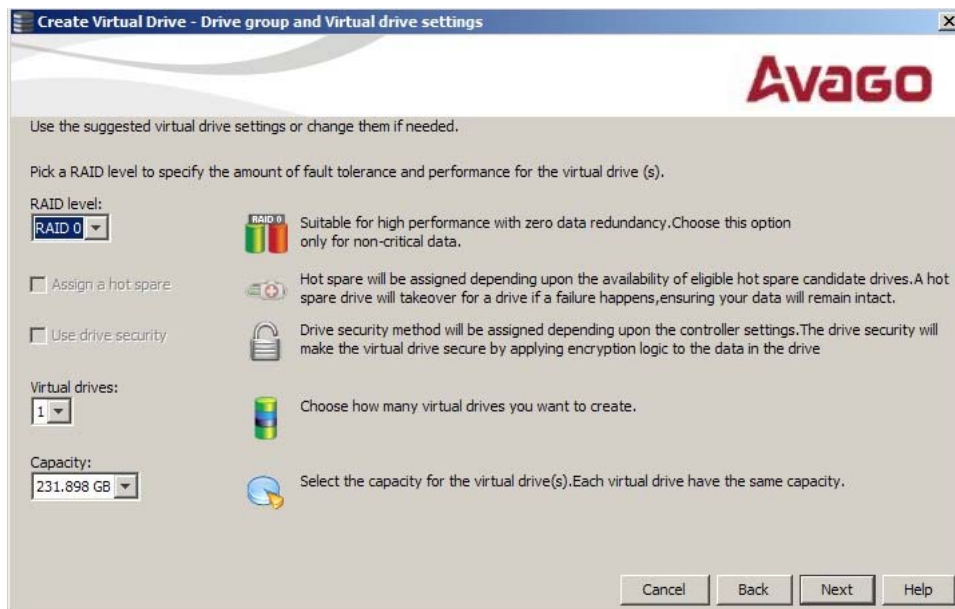
Figure 107 Virtual Drive Simple Configuration Mode



15-1197_msmUtility_virtualDriveSimpleconfigMode

3. Choose **Simple** and click **Next**.
The following screen appears.

Figure 108 Create Virtual Drive Screen



15-1197_msmUtility_createSimpleVirtualDriveScreen

4. Select the RAID level desired for the virtual drive.
When you use simple configuration, the RAID controller supports RAID levels 1 and 5. In addition, it supports independent non-redundant drives (configured as RAID 0). The screen text gives a brief description of the RAID

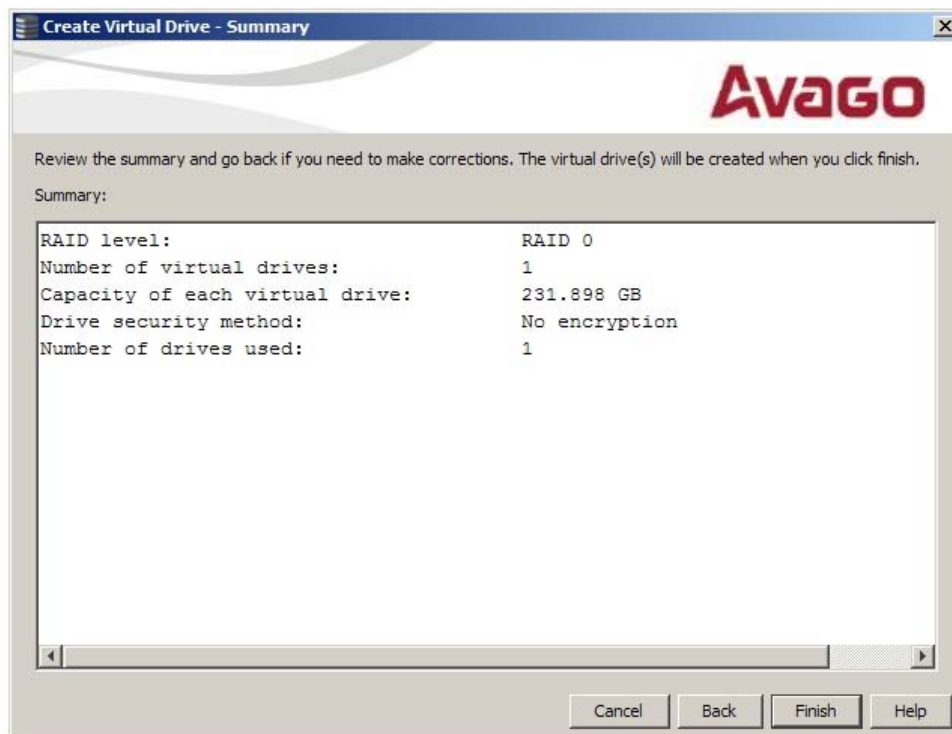
level that you select. The RAID levels that you can choose depend on the number of drives available. To learn more about RAID levels, see [Section 1.2, RAID Overview](#), on page 8.

5. Use the drop-down menu in the **Virtual drives** field to choose how many virtual drives you want to create.
6. Select the capacity of the virtual drive(s).
Each virtual drive has the same capacity.

7. Click **Next**.

The following screen appears. This window shows the selections you made for a simple configuration.

Figure 109 Create Virtual Drive – Summary Window



15-1197_msmUtility_createVirtualDrive_summaryScreen

8. 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.

6.1.3 Creating a Virtual Drive Using Advanced Configuration

The **Advanced** configuration procedure provides an easy way to create a new storage configuration. The **Advanced** configuration gives you greater flexibility than the **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.

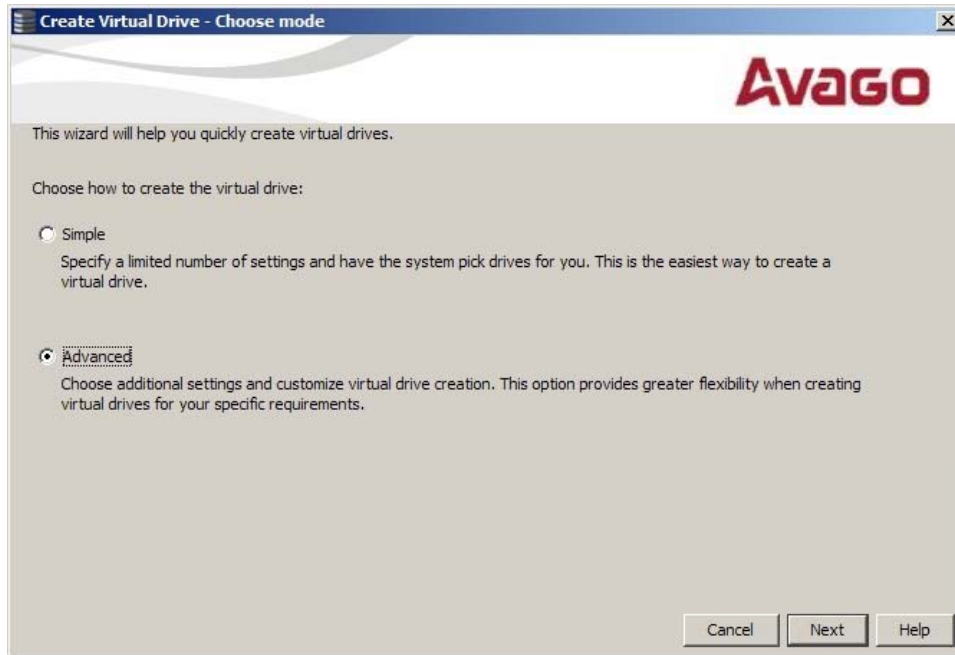
Follow these steps to create a new storage configuration in the **Advanced** configuration mode.

1. Click the **Dashboard** tab on the MegaRAID Storage Manager utility main menu screen.

The **Create Virtual Drive** option appears in the **Actions** section of the **Dashboard**, as shown in [Figure 106, Virtual Drive Creation Menu](#).

2. Click **Create Virtual Drive** in the **Actions** section of the screen.
The following screen appears.

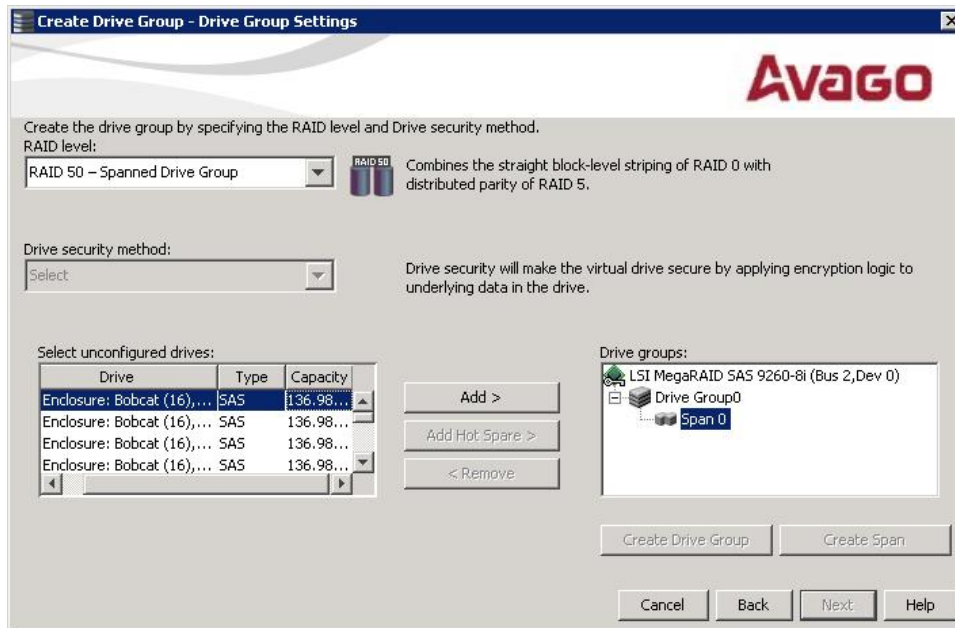
Figure 110 Virtual Drive Advanced Configuration Mode



15-1197_msmUtility_virtualDriveAdvancedConfigMode

3. Click **Advanced**, then click **Next**.
The following screen appears.

Figure 111 Create Drive Group Settings Dialog Box



15-1197_msmUtility_virtualDriveAdvancedConfigMode_createDriveGroup_DriveGroupSettings

4. Perform the following actions in the **Create Drive Group Settings** dialog box:

- a. Select the RAID level desired for the drive group from the drop-down menu.

The drive group appears in the **Drive groups** field.

The RAID controller supports RAID levels 1, 5, and 10. In addition, it supports independent drives configured as RAID 0. The screen text gives a brief description of the RAID level you select. Which RAID levels you can choose depend on the number of drives available.

- b. Select unconfigured drives from the list of drives and click **Add>** to add them to the drive group.

The selected drives appear under **Drive Group {number}**.

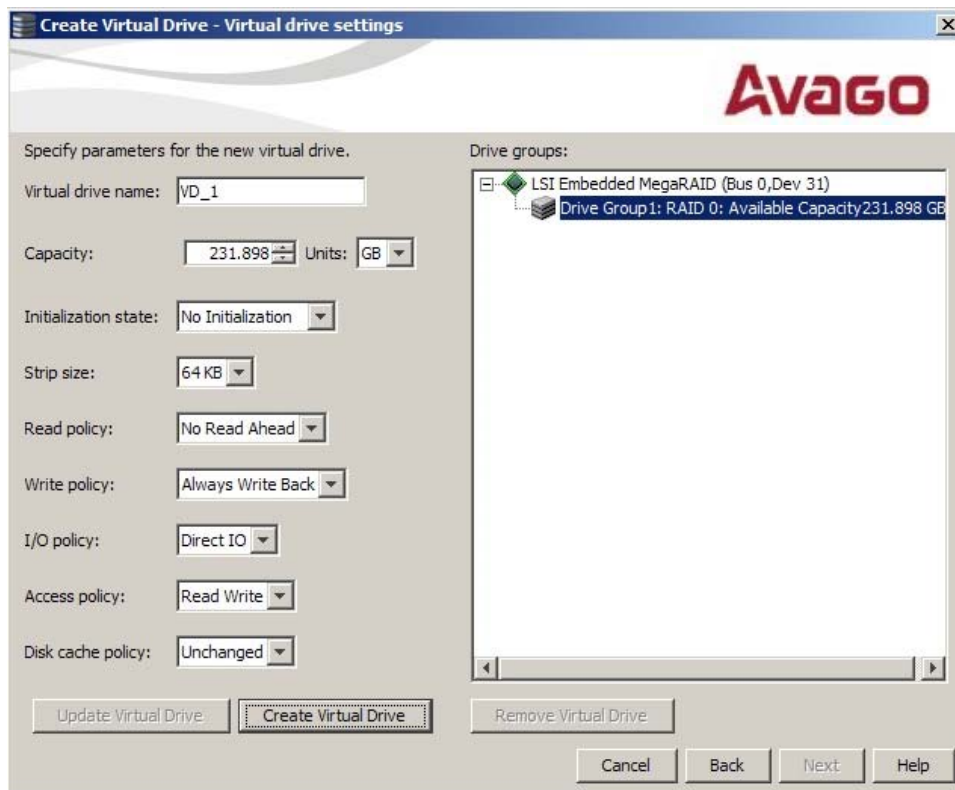
- c. If you are creating a RAID 10 drive group, perform the following steps. If not, go to step d.

NOTE On a RAID 10 array, you can create only one virtual drive, and that virtual drive must occupy the entire space of the RAID 10 array.

- Click **Create Span** to create a second span in the drive group.
 - Select unconfigured drives from the list of drives and click **Add>** to add them to the second drive group.
 - The selected drives appear under the second span, **Span {number}**, below the second drive group, **Drive Group {number}**.
 - Click **Create Drive Group** to make a drive group with the spans.
 - Click **Next** to complete the steps for a RAID 10 configuration.
- d. Click **Create Drive Group** to make a drive group.
- e. Click **Next** to complete this step.

The following screen appears. The drive group and the default virtual drive settings appear. The options to update or remove the virtual drive are grayed out until you create the drive.

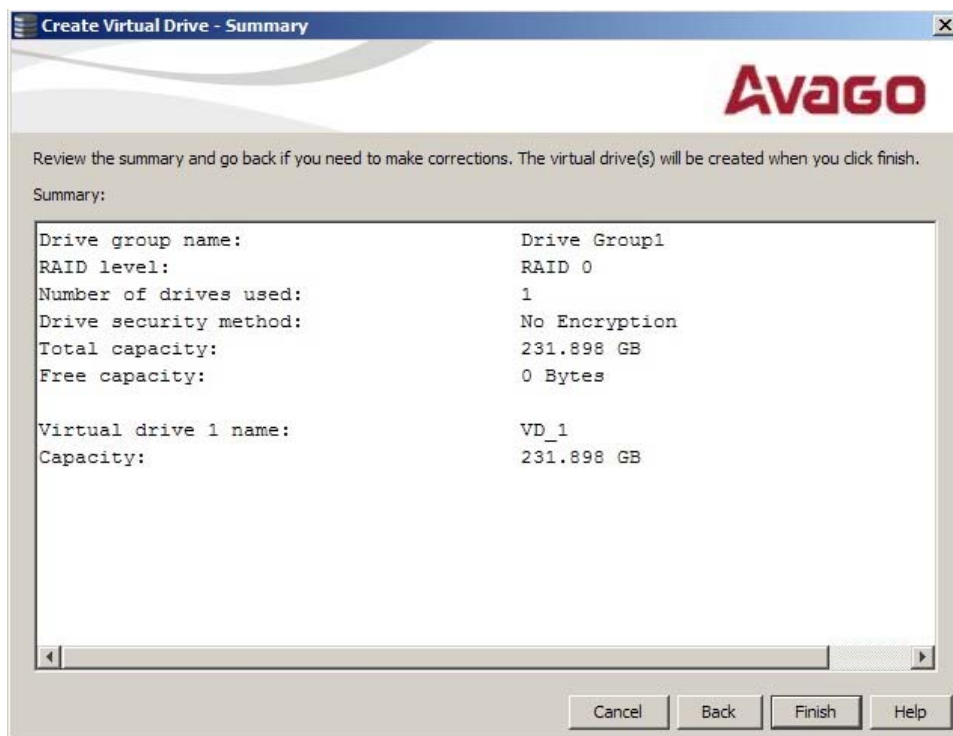
Figure 112 Virtual Drive Settings Dialog Box



15-1197_msmUtility_createVirtualDrive_virtualDriveSettings

5. Select the virtual drive settings to fit your environment.
See [Section 6.1.1, Selecting Virtual Drive Settings](#), on page 123, for more information about the virtual drive settings.
6. Click **Create Virtual Drive**.
The new virtual drive appears under the drive group. The `Update Virtual Drive` and `Remove Virtual Drive` options are available. The `Update Virtual Drive` option allows you to change the virtual drive settings, and the `Remove Virtual Drive` option let you delete the virtual drive.
7. Click **Next**.
The following screen appears. This window shows the selections you made for the advanced configuration.

Figure 113 Virtual Drive Summary Window



15-1197_msmUtility_createVirtualDrive_summary

8. 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.
After the configuration is completed, a dialog box notifies you that the virtual drives have been successfully created. If more drive capacity exists, the dialog box asks whether you want to create more virtual drives. If no more drive capacity exists, you are prompted to close the configuration session.
9. Select **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** screen, as shown in [Figure 108](#). If you select **No**, the utility asks whether you want to close the wizard.
10. If you selected **No** in the previous step, select **Yes** or **No** to indicate whether you want to close the wizard.
If you select **Yes**, the configuration procedure closes. If you select **No**, the dialog box closes and you remain on the same page.

6.2 Changing Adjustable Task Rates

Follow these steps if you need to change the adjustable rates for rebuilds, and other system tasks that run in the background.

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 will 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: _____

1. Select a controller icon in the **Physical** tab or the **Logical** tab in the left panel of the MegaRAID Storage Manager main menu screen.
2. Select **Go To > Controller > Set Adjustable Task Rates** from the menu bar, as shown in the following figure.

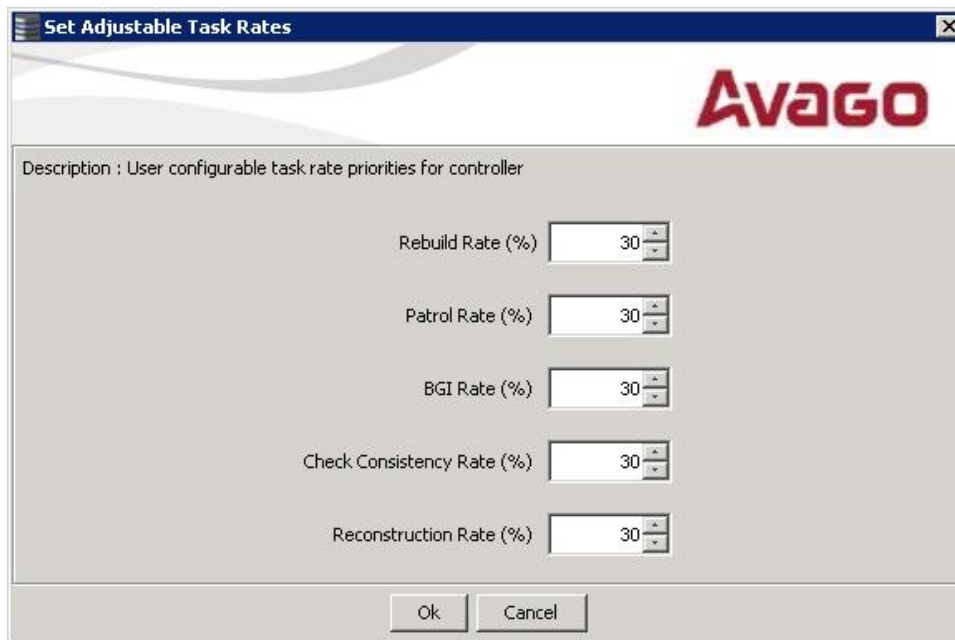
Figure 114 Set Adjustable Task Rates Menu



15-1197_msmUtility_goToController_setAdjustableTaskRates

The following screen appears.

Figure 115 Set Adjustable Task Rates Dialog Box



3. Enter changes, as needed, to the following task rates:

- Rebuild rate
- Patrol read
- BGI rate (for fast initialization)
- Check consistency rate (for consistency checks).
- Reconstruction rate

Each task rate can be set from 0 percent to 100 percent. The higher the number, the faster the activity runs in the background, possibly impacting other system tasks.

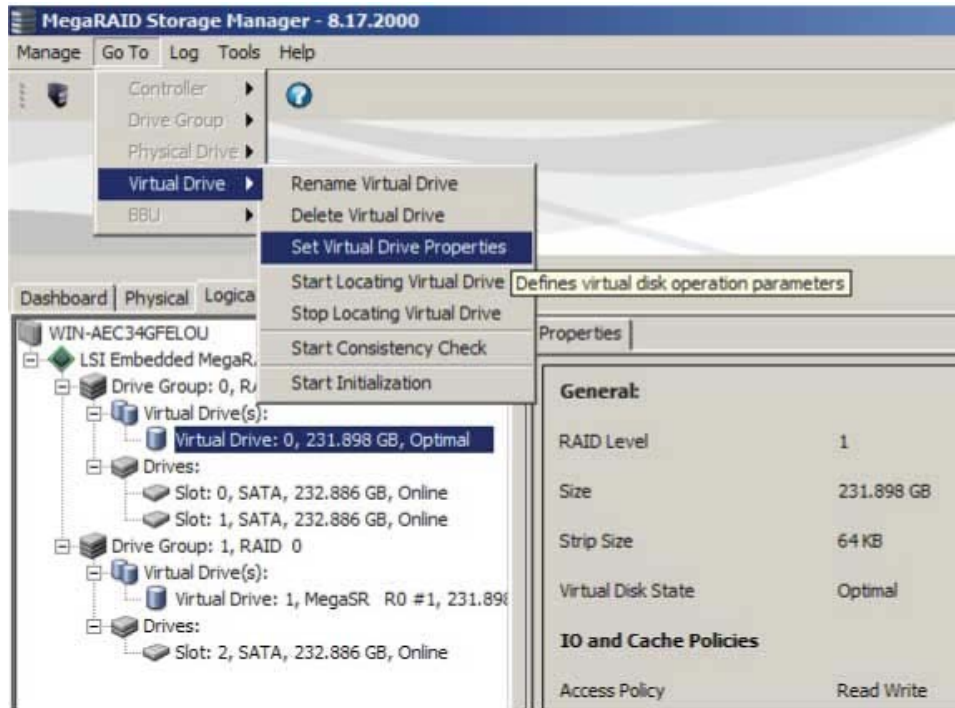
4. Click **OK** to accept the new task rates.
5. When the warning message appears, click **OK** to confirm that you want to change the task rates.

6.3 Changing Virtual Drive Properties

You can change the read policy, the write policy, and the other virtual drive properties at any time after a virtual drive is created. 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 main menu screen.
2. Select **Go To > Virtual Drive > Set Virtual Drive Properties** from the menu bar, as shown in the following figure.

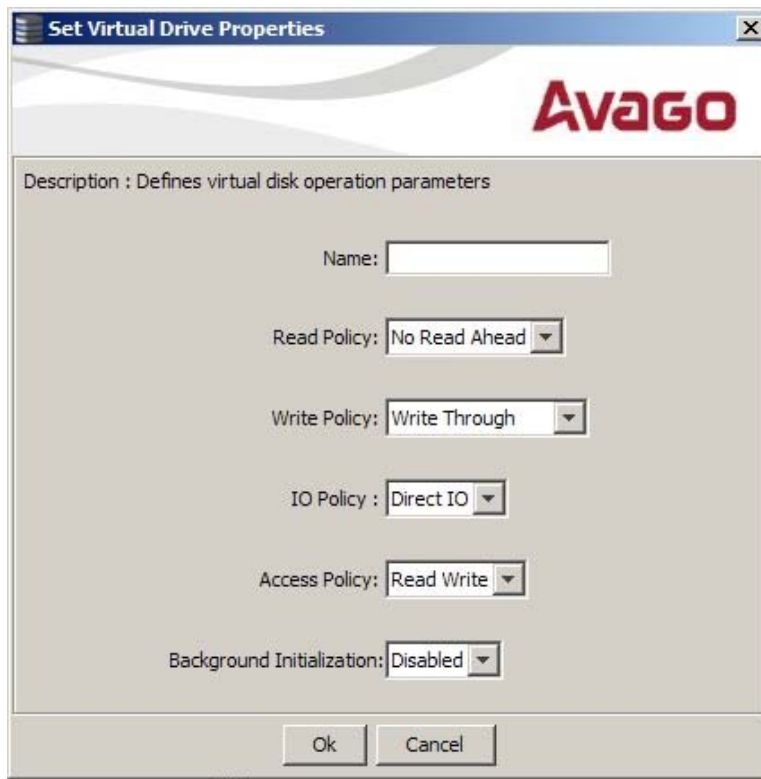
Figure 116 Set Virtual Drive Properties Menu



15-1197_msmUtility_goToVirtualDrive_setVirtualDriveProps

The following screen appears.

Figure 117 Set Virtual Drive Properties Screen



15-1197_msmUtility_setVirtualDrivePropsDialog

3. Change the virtual drive properties as needed.
For information about these properties, see [Section 6.1.1, Selecting Virtual Drive Settings](#), on page 123.
4. Click **Ok** to accept the changes.
The virtual drive settings are updated.

6.4 Deleting a Virtual Drive

CAUTION Be sure to back up the data that is 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 that is 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 whether you want to delete a virtual disk to avoid deleting the virtual disk by mistake.

Chapter 7: System Event and Storage Device Monitors

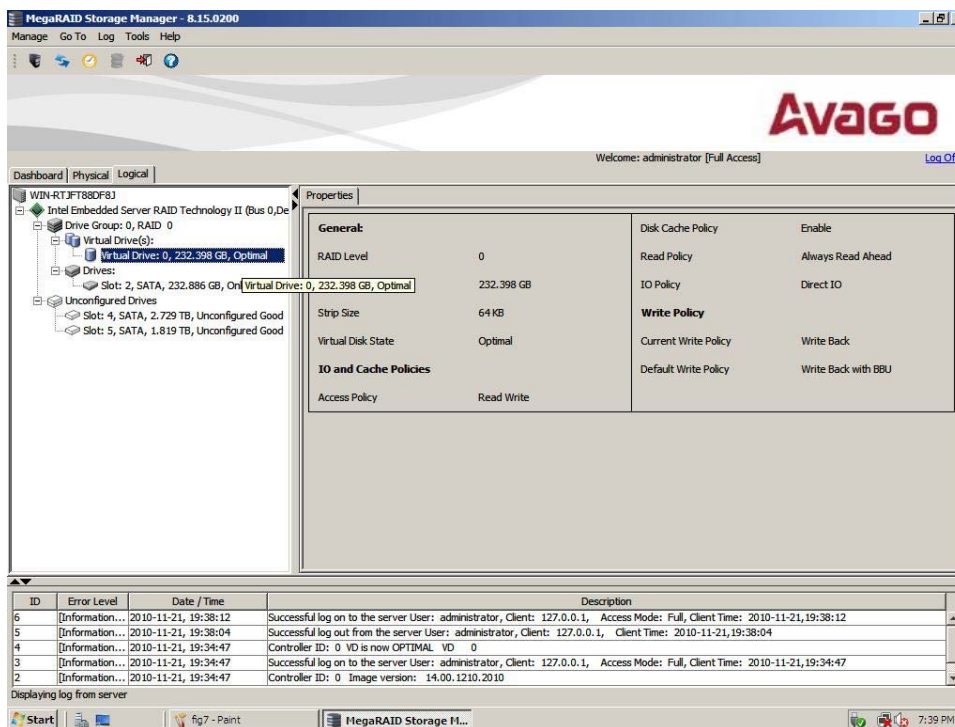
This chapter explains how to use MegaRAID Storage Manager utility to monitor the status of drives, virtual drives, and other storage devices.

7.1 Monitoring System Events

The MegaRAID Storage Manager utility 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 at the bottom of the **MegaRAID Storage Manager** main menu screen, as shown in the following figure.

You can use the MegaRAID Storage Manager utility to alert you about events. Settings exist for the delivery of alerts, the severity level of events, exceptions, and email settings.

Figure 118 Event Information Window



15-1197_msmUtility_eventInformationWindow

Each message that appears in the event log has a severity level that indicates the importance of the event (as shown in the following table), a date and timestamp, and a brief description. You can click an event to display the same information in a window. (For a list of all events, see [Appendix A, Events and Messages](#), on page 166).

Table 9 Event Severity Levels

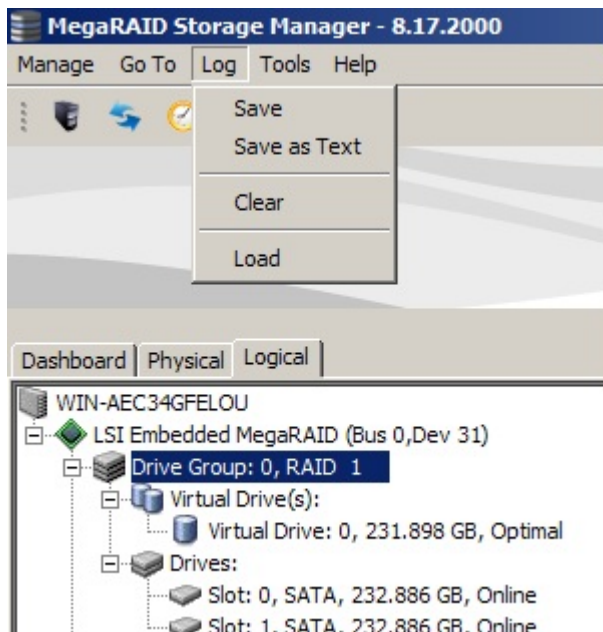
Severity Level	Meaning
Information	Informational message. No user action is necessary.
Warning	Some component might be close to a failure point.
Critical	A component has failed, but the system has not lost data.
Fatal	A component has failed, and data loss has occurred or will occur.

The **Log** menu has four options:

- **Save**
Saves the current log to a *.log file.
- **Save as Text**
Saves the current log in .txt format.
- **Clear**
Clears the current log information. You have the option to save the log first.
- **Load**
Enables you to load a local *.log file.

The following figure shows the log menu.

Figure 119 Log Menu



15-1197_msmUtility_log_Menu

7.2 Configuring Alert Notifications

The *Alert Notification Configuration* feature allows you to control and configure the alerts that the MegaRAID Storage Manager utility sends when various system events occur.

To access this screen, select **Tools > Configure Alerts** on the main menu screen.

The following screen appears.

Figure 120 Alert Notification Configuration Menu



3_03132-00

The **Alerts Notification Configuration** 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 121 Alerts Notification Configuration Screen



3_03133-00

You can select the **Alert Settings** tab to perform the following actions:

- Select the methods for the delivery of alerts.
- Change the severity level of events.
- 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 send to the monitor.

ATTENTION When you load a saved backup file, all unsaved changes made in the current session are lost.

Select the **Mail Server** tab to perform the following actions:

- Enter or edit the sender email address.
- Enter the name of the SMTP server.
- Require authentication of the email server.
- 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 send to the monitor.

ATTENTION When you load a saved backup file, all unsaved changes made in the current session will be 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 a *.xml backup file of the entire alert configuration.
- Load all the values from a previously saved backup into the dialog to edit or send to the monitor.

ATTENTION When you load a saved backup file, all unsaved changes made in the current session will be lost.

7.2.1 Setting Alert Delivery Methods

You can select the methods used to send alert deliveries, which include 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 **Alerts Notification Configuration** screen, click the **Alert Settings** tab.
2. Under the **Alerts Delivery Methods** heading, select one of the severity levels.
3. Click **Edit**.

The following screen appears.

Figure 122 Alert Notification Delivery Methods Dialog Box



3_03134-00

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.

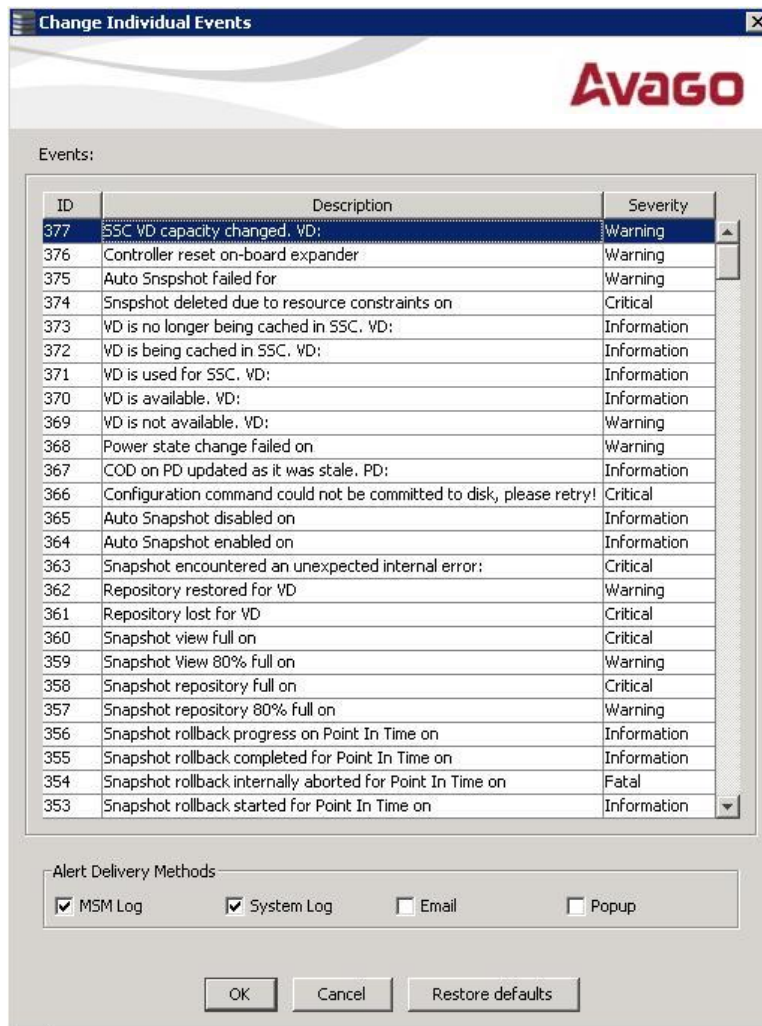
7.2.2 Changing Alert Delivery Methods for Individual Events

You can change the alert delivery options for an event without changing the severity level.

1. On the **Alerts Notification Configuration** screen, click the **Alert Settings** tab.
The **Alerts Setting** portion of the screen appears, as shown in [Figure 121](#).
2. Click **Change Individual Events**.

The following screen appears. The dialog box shows the events by their ID number, description, and severity level.

Figure 123 Change Individual Events Dialog Box



15-1197_monitoring_changeIndividualEventsDialogBox

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. Press **Esc** to return to the **Alerts Notification Configuration** screen.
6. Click **OK**.
This action saves all the changes made to the event.

7.2.3 Changing the Severity Level for Individual Events

See [Table 9](#) for details about the severity levels.

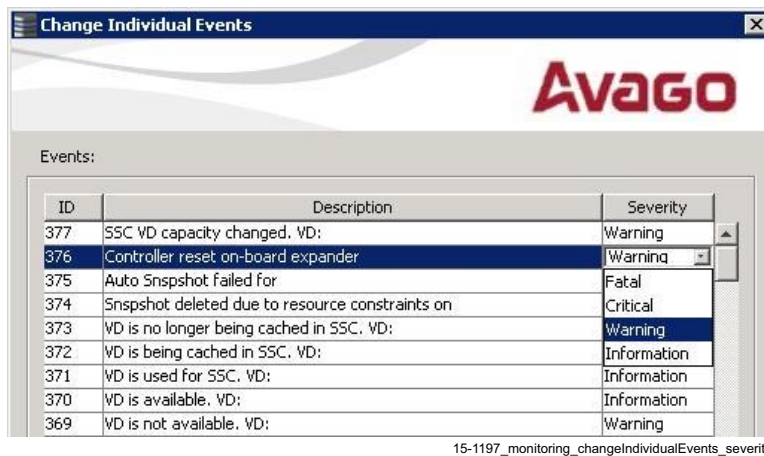
To change the event severity level for a specific event, perform the following steps.

1. On the **Alerts Notification Configuration** screen, click the **Alert Settings** 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 123](#). 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 following screen appears.

Figure 124 Change Individual Events Severity Level Menu

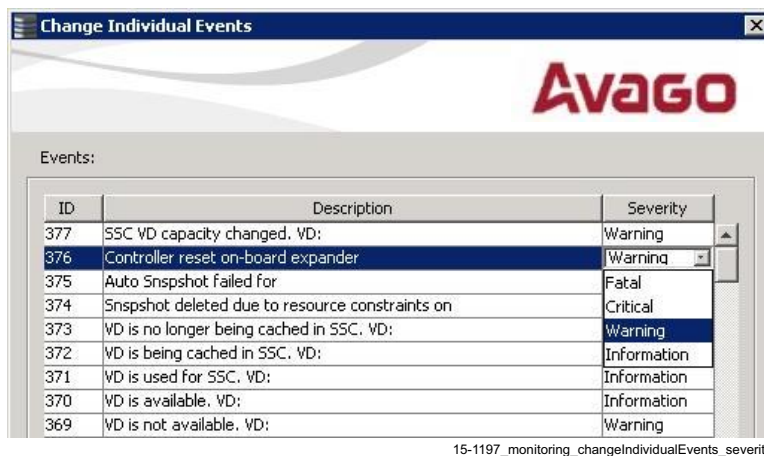


5. Select a different severity level for the event from the menu.
6. Press Esc to return to the **Alerts Notification Configuration** screen.
7. Click **OK** to save all the changes made to the events.

7.2.4 Multiple Events Displayed in a Single Pop-Up Window

You can view multiple events in a single pop-up window, as shown in the following figure.

Figure 125 Pop-Up for Multiple Events



7.2.5 Entering or Editing the Sender Email Address and SMTP Server

You can use the **Alerts Notification Configuration** screen to enter or edit the sender email address and the SMTP server.

1. On the **Alerts Notification Configuration** screen, click the **Mail Server** tab.
The following screen appears.

Figure 126 Mail Server Options



15-1197_monitoring_configureAlertsMailServer_dialogBox

2. Enter a new sender email address in the **Sender email address** field or edit the existing sender email address.
3. Click **OK**.

7.2.6 Authenticating a Server

You can use the **Alerts Notification Configuration** 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.

Perform the following steps to enter or edit the address:

1. On the **Alerts Notification Configuration** screen, click the **Mail Server** tab.
The **Mail Server** options appears, as shown in [Figure 126](#). The authentication check box is selected by default.
2. Enter a user name in the **User name** field.

3. Enter the password in the **Password** field.
4. Click **OK**.

7.2.7 Saving Backup Configurations

You can save an `.xml` backup file of the entire alert configuration. This backup file includes all the settings on the three tabs.

1. On the **Alerts Notification Configuration** screen, click the **Alert Settings** tab, **Mail Server** tab, or **Email** tab.
2. Click **Save Backup**.
The drive directory appears.
3. Enter a file name 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 **Alert Notification Configuration** screen closes.

7.2.8 Loading Backup Configurations

You can load all the values from a previously saved backup into the dialog (all tabs) to edit or send to the monitor.

CAUTION 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 **Alerts Notification Configuration** screen, choose the **Alert Settings** 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.
3. Click **Yes**.
The drive directory appears, from which you can select a backup configuration to load.
4. Select the backup configuration file (it should be in `.xml` format).
5. Click **Open**.
The drive directory disappears.
6. Click **OK**.
The backup configuration is loaded, and the **Alerts Notification Configuration** screen is closed.

7.2.9 Adding Email Addresses of Recipients of Alert Notifications

The **Email** tab portion of the **Alerts Notification Configuration** screen shows the email addresses of recipients of the alert notifications. The MegaRAID Storage Manager utility 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 **Email** tab on the **Event Notification Configuration** screen.
The following screen appears.

Figure 127 Email Settings



15-1197_monitoring_configureAlerts_addingEmailAddresses

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.

7.2.10 Testing Email Addresses of Recipients of Alert Notifications

Use the **Email** tab portion of the **Alerts Notification Configuration** screen to send test messages to the email addresses that you added for the recipients of alert notifications.

1. Click the **Email** tab on the **Event Notification Configuration** screen.
The **Email** section of the screen appears, as shown in the previous figure.
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.

If the MegaRAID Storage Manager utility cannot send an email message to the email address, an error message appears.


7.2.11 Removing Email Addresses of Recipients of Alert Notifications

Use the **Email** tab portion of the **Alerts Notification Configuration** screen to remove email addresses of the recipients of alert notifications.

1. Click the **Email** tab on the **Event Notification Configuration** screen.
 The **Email** section of the screen appears, as shown in [Figure 127](#).
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.


7.3 Monitoring Controllers

NOTE The Embedded MegaRAID Software drivers act as virtual *controllers*. Because these are not actual hardware storage controllers installed in the computer system, some of the controller properties shown in the following screen do not apply to them.

When the MegaRAID Storage Manager utility is running, you can see the status of all controllers in the left panel of the **MegaRAID Storage Manager** main menu screen. 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 9.2.1, Dashboard View/Physical View/Logical View](#), on page 160, 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** main menu screen. The following screen appears.


Figure 128 Controller Properties




Most of the information on this screen is self-explanatory. Note that the *Rebuild Rate*, *Patrol Read Rate*, *Reconstruction Rate*, *Consistency Check Rate*, and *BGI Rate* are all user selectable. For more information about these rates, see [Section 6.2, Changing Adjustable Task Rates](#), on page 131.

The controller properties are defined in [Appendix A, Events and Messages](#), on page 166.

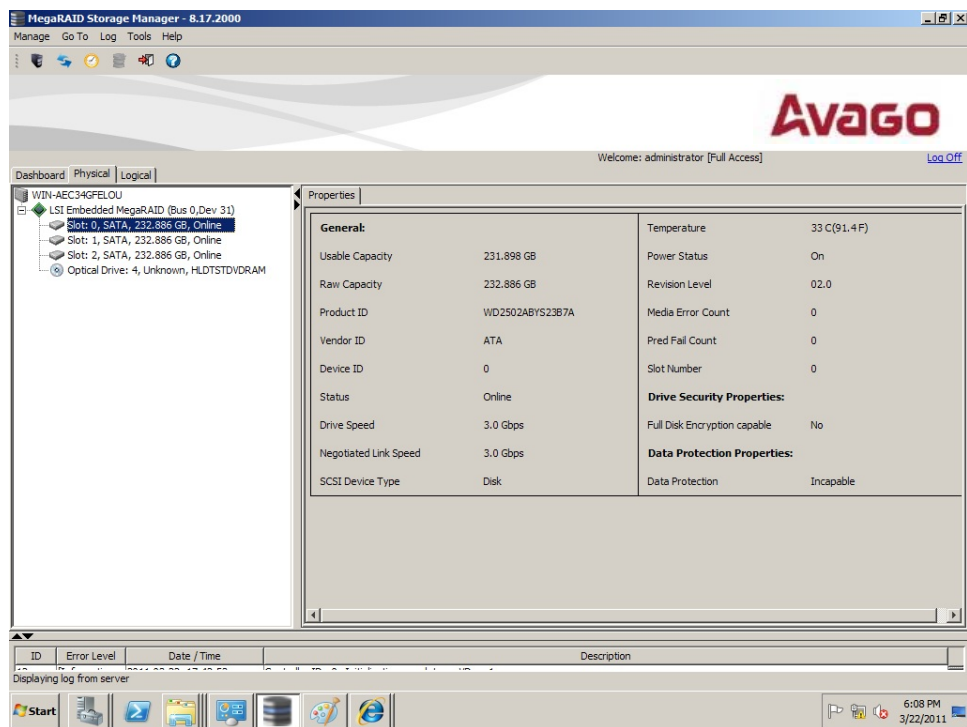
7.4 Monitoring Drives

When the MegaRAID Storage Manager utility is running, you can see the status of all drives in the left panel of the **MegaRAID Storage Manager** main menu screen. 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 9.2.1, Dashboard View/Physical View/Logical View](#), on page 160, or a complete list of device icons.)

To display complete drive information, click a drive icon in the left panel of the **MegaRAID Storage Manager** main menu screen. The following screen appears.

Figure 129 Drive Information



15-1197_monitoring_driveInformationProperties

The information on this panel is self-explanatory. No user-selectable properties exist for physical devices. Icons for other storage devices such as CD-ROM drives and digital audio tape (DAT) drives can also appear in the left panel.

If the drives are in a drive enclosure, you can identify which drive is represented by each drive LED on the enclosure. Follow these steps to locate the drive:

1. Click the drive icon in the left panel.
2. Click **Go To > Physical Drive > Start Locating Drive**.

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.

3. To stop the drive LED on the enclosure from blinking, select **Go To > Physical Drive > Stop Locating Drive**.

To display a graphical view of a drive, click a drive icon in the left panel of the **MegaRAID Storage Manager** main menu screen, and click the **Graphical View** tab. In the 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
- Reserved space is red

When you select a virtual drive from the drop-down menu, the drive space used by that virtual drive appears in green.

7.5 Running a Patrol Read Operation

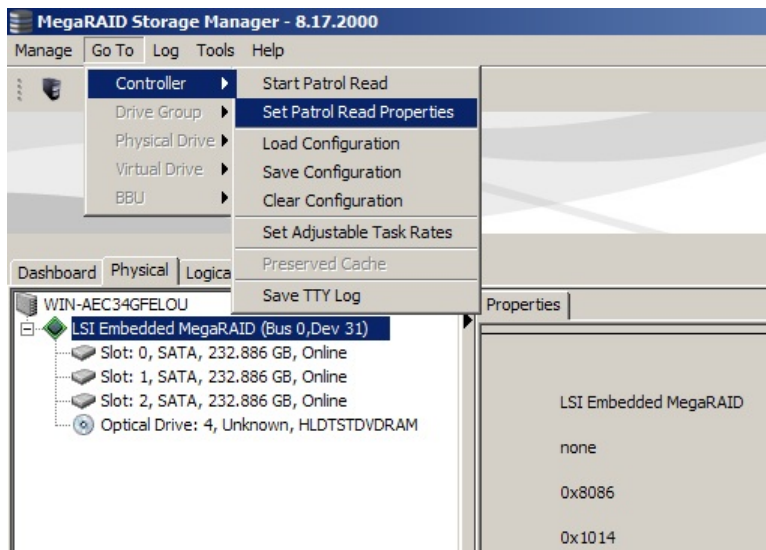
A *patrol read* operation periodically verifies all sectors of drives connected to a controller, which includes the system reserved area in the RAID configured drives. A *patrol read* operation can be used for all RAID levels and for all hot spare drives. This operation is initiated only when the controller 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* operation without changing the properties.

To set the `patrol read` properties and then start a *patrol read* operation, 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 > Set Patrol Read Properties**.

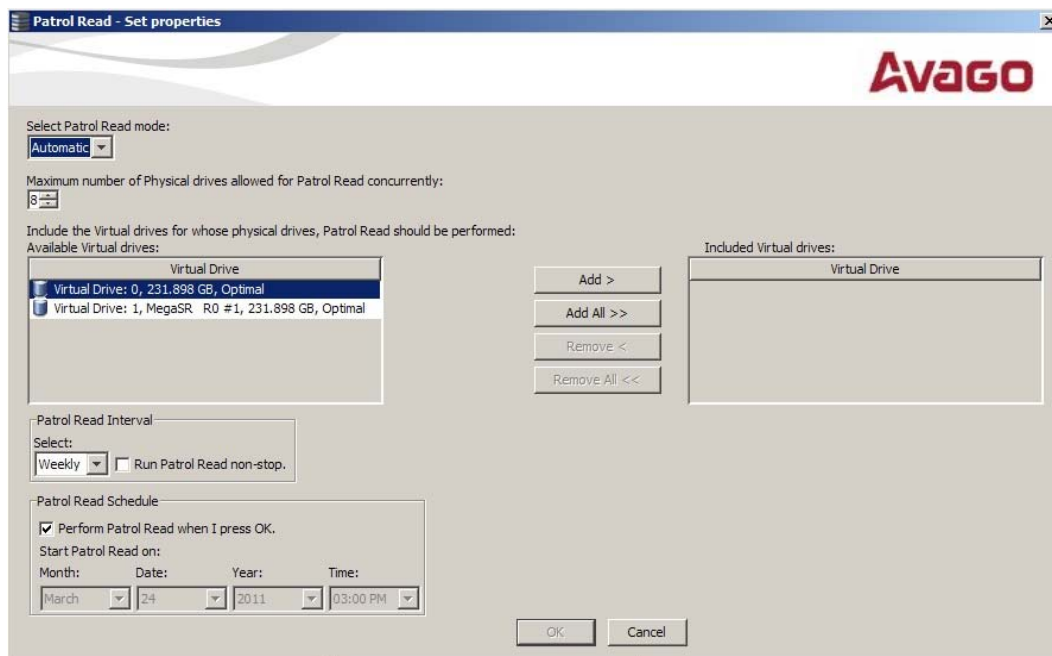
Figure 130 Start Patrol Read Menu



15-1197_monitoring_patrolRead_propertiesMenu

The following screen appears.

Figure 131 Patrol Read Configuration



15-1197_monitoring_patrolRead_propertiesWindow

3. Select a mode for a *patrol read* operation. The options are:

- **Automatic**

The *patrol read* operation runs automatically at the time interval you specify on this screen.

- **Manual**

The *patrol read* operation runs only when you manually start it by selecting **Start Patrol Read** from the controller **Options** panel.

— Disabled

The *patrol read* operation does not run.

- Specify the maximum number of drives to include in the *patrol read* operation. Eight drives is the maximum supported by Embedded MegaRAID Software.
- 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 the virtual drives.
- (Optional) Change the frequency at which the *patrol read* operations will run. The default frequency is weekly (168 hours), which is suitable for most configurations. The other options are hourly, daily, and monthly.

NOTE

Avago recommends that you 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: _____

- (Optional) Set the *patrol read* operation to run at a specific time. The default is for the *patrol read* operation to start when you click **OK** on this screen. To change the default so that the *patrol read* operation starts at a specific time, follow these steps (otherwise, skip this step and proceed to the next step):
 - Uncheck the box **Perform Patrol Read when I click OK**.
 - Select the month, year, day, and time to start patrol read.
- Click **OK** to enable your *patrol read* selections.

NOTE

The *patrol read* operation does not report on its progress while it is running. The *patrol read* status is reported in the event log only.

To start a *patrol read* operation without changing the `patrol read` properties, follow these steps:


- Click a controller icon in the left panel of the **MegaRAID Storage Manager** main menu screen.
- Select **Go To > Controller > Start Patrol Read** in the menu bar.
- When prompted, click **Yes** to confirm that you want to start a *patrol read* operation.


7.5.1 Patrol Read Task Rates

You can change the *patrol read task rate*. The task rate determines the amount of system resources that are dedicated to a *patrol read* operation when it is running. Avago 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 6.2, Changing Adjustable Task Rates](#), on page 131.

7.6 Monitoring Virtual Drives

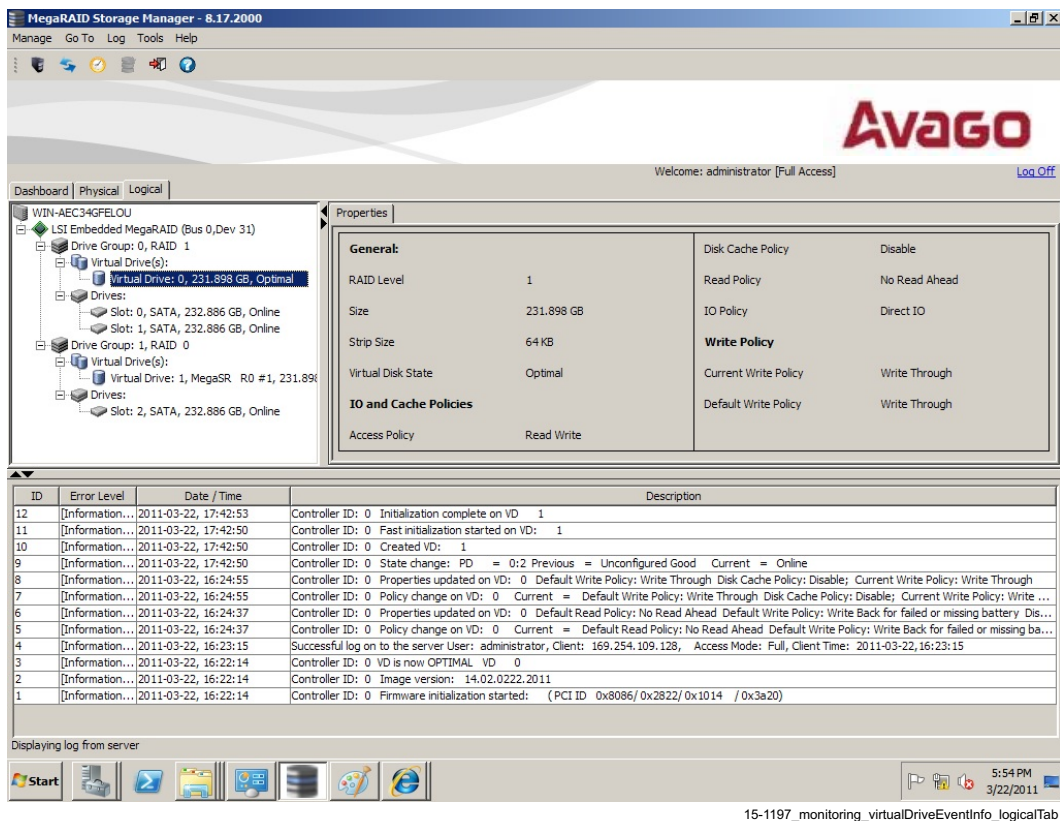
When the MegaRAID Storage Manager utility is running, you can see the status of all virtual drives. If a virtual drive is operating normally, the icon looks like this: . If the virtual drive is running in *Degraded* mode (for example, if a drive

has failed), a small yellow circle appears to the right of the icon: . A red circle indicates that the virtual drive has failed and data has been lost.

When the **Logical** tab is selected, the left panel of the **MegaRAID Storage Manager** main menu screen 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. The following figure shows the **Properties** panel for a virtual drive.

Figure 132 Virtual Drive Properties



The RAID level, stripe size, and access policy of the virtual drive are set when the virtual drive is configured.

NOTE You can change the read policy, write policy, and other virtual drive properties. See [Section 6.3, Changing Virtual Drive Properties](#), on page 132, for the procedure you can use to change these properties.


If the drives in the virtual drive are in an enclosure, you can identify them by making their LEDs blink. To identify the drives, 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.

The LEDs on the drives in the virtual drive start blinking (except for the LEDs for hot spare drives).

3. To stop the LEDs from blinking, click **Go To > Virtual Drive > Stop Locating Virtual Drive**.

7.7 Monitoring Enclosures

When the MegaRAID Storage Manager utility 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.

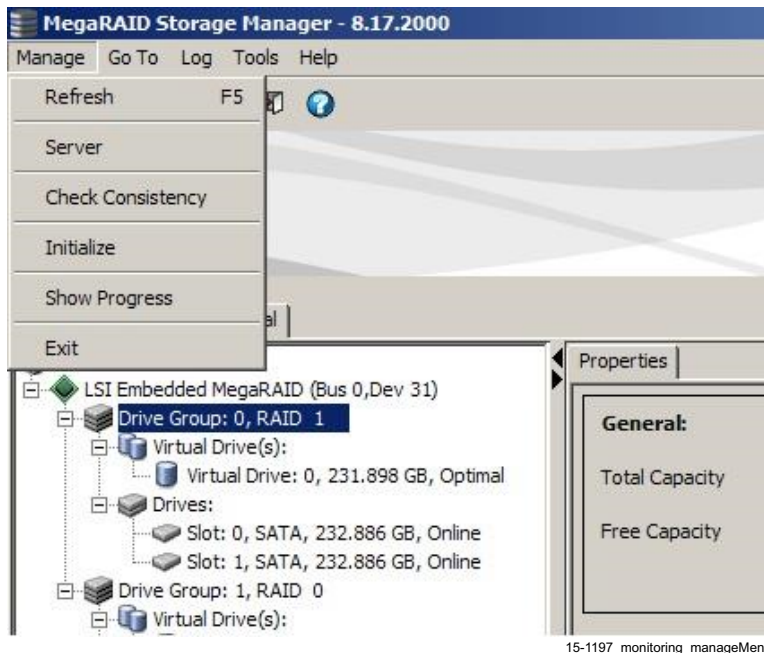
7.8 Monitoring Rebuilds and Other Processes

The MegaRAID Storage Manager utility lets you monitor the progress of rebuilds and other lengthy operations in the **Group Show Progress** window.

Follow these steps to monitor the progress of these operations.

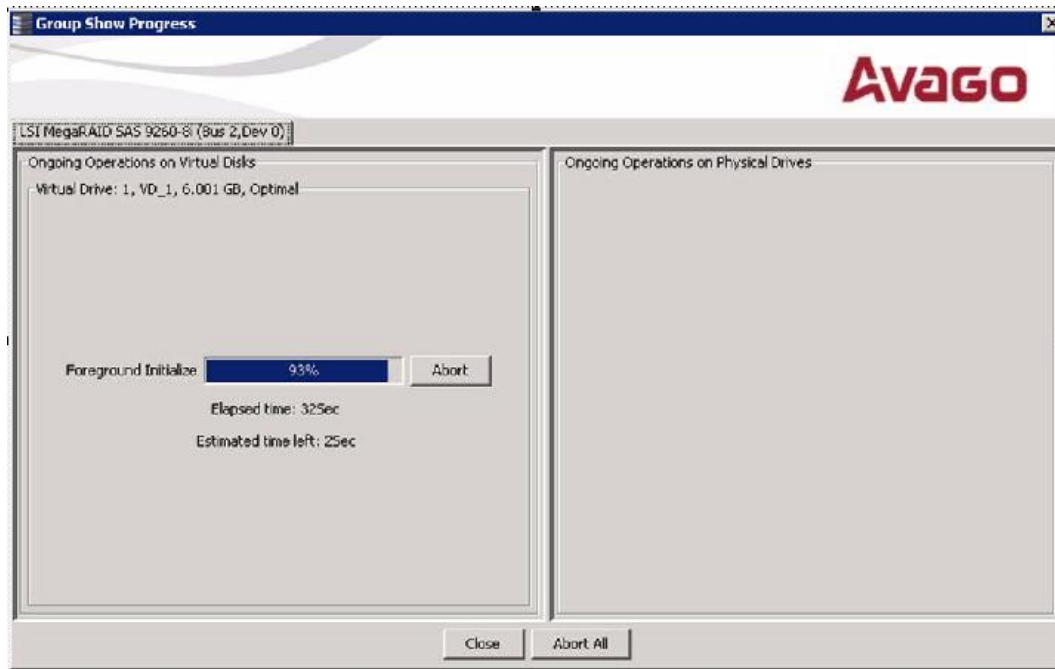
1. Select **Manage > Show Progress** on the menu bar, as shown in the following figure.

Figure 133 Group Show Progress Menu



The following screen appears.

Figure 134 Group Show Progress Window



15-1197_monitoring_groupShowProgressWindow

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 operations appear in this window:

- Background or foreground initialization of a virtual drive (see [Section 8.1, Initializing a Virtual Drive](#), on page 153)
 - Rebuild (see [Section 8.4, Rebuilding a Drive](#), on page 156)
 - Check consistency (see [Section 8.2, Running a Consistency Check Operation](#), on page 154)
2. (Optional) Click **Abort All** to abort all ongoing processes.
 3. Click **Close** to close the window.

Chapter 8: Storage Configuration Management and Maintenance

This chapter explains how to use Embedded RAID to maintain and manage storage configurations.

8.1 Initializing a Virtual Drive

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** main menu screen, and click the icon of the virtual drive that you want to initialize.
2. Select **Go To > Virtual Drive > Start Initialization**.
The initialize dialog box appears.
3. Select the virtual drive(s) to initialize.

CAUTION Initialization erases all data on the virtual drive. Make sure to back up any data you want to keep before you initialize. Make sure the operating system is not installed on the virtual drive you are initializing.

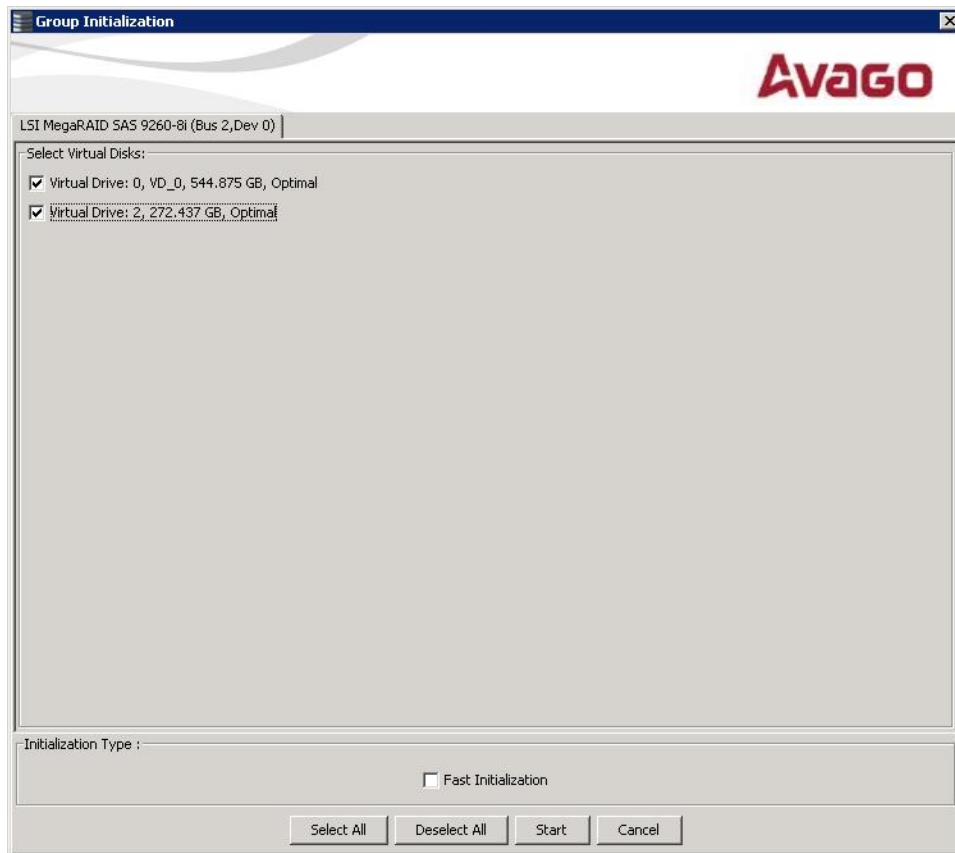
4. Select the **Fast Initialization** check box if you want to use this option.
If you leave the box unchecked, the MegaRAID Storage Manager utility runs a *Full Initialization* operation on the virtual drive. For more information, see [Section 8.1.1, Running a Group Initialization](#), on page 153.
5. Click **Start** to begin the initialization.
You can monitor the progress of the initialization. For more information, see [Section 7.8, Monitoring Rebuilds and Other Processes](#), on page 151.

8.1.1 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 initialization.

1. Click **Manage > Initialize**.
The following screen appears.

Figure 135 Group Initialization Dialog Box



15-1197_maintenance_groupInitialization

2. Check the virtual drives to run the initialization on or click **Select All** to select all the virtual drives.
3. Click **Start**.

You can monitor the progress of the group initialization. See [Section 7.8, Monitoring Rebuilds and Other Processes](#), on page 151, for more information.

8.2 Running a Consistency Check Operation

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

You should run a *consistency check* operation on fault-tolerant virtual drives periodically. You must run the consistency check if you suspect that the virtual drive data might be corrupted. Be sure to back up the data before you run 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* operation. This section explains how to set the properties, schedule the check, and run the consistency check.


8.2.1 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. Click **Manage > Check Consistency**.

The following screen appears.

Figure 136 Group Consistency Check Dialog Box



15-1197_monitoring_groupConsistencyCheckWindow

2. Check the virtual drives to run the consistency check on or click **Select All** to select all the virtual drives.
3. Click **Start**.
You can monitor the progress of the group consistency check. See [Section 7.8, Monitoring Rebuilds and Other Processes](#), on page 151, for more information.

8.3 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 utility.

The utility lets you 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 utility normally detects newly installed drives and displays icons for them in the **MegaRAID Storage Manager** main menu screen. If for some reason the MegaRAID Storage Manager utility 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** main menu screen.
2. Select **Go To > Controller > Scan for Foreign Configuration**.
If the MegaRAID Storage Manager utility 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.

8.4 Rebuilding a Drive

If a single drive in a RAID 1, RAID 5, or RAID 10 virtual drive fails, the system is protected from data loss. If hot spare disks are available, a failed drive is rebuilt automatically without any user intervention. A 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 hot spare drives 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 symbol 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 if you need 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 starts. 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 utility.
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* operation in the **Group Show Progress** window by selecting **Manage > Show Progress**.

8.5 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 utility commands to remove the drive from the first configuration and change the drive state to *Unconfigured Good*.

CAUTION After you perform this procedure, *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** main menu, click **Go To > Physical Drive > Make Drive (O)ffline**.

The drive status changes to *Offline*.

2. Click **Go To > Physical Drive > (M)ark Drive as Missing**.

The drive status changes to *Unconfigured Good*.

CAUTION 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 is then rebuilt. You can now use the removed drive for another configuration.

CAUTION If the MegaRAID Storage Manager utility 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 can make the drive so that another configuration cannot use it by using the **Mark physical disk as missing** command and the **Rescan** command.

Chapter 9: MegaRAID Storage Manager Screen and Menus

This chapter explains how to start MegaRAID Storage Manager and describes the MegaRAID Storage Manager main menu screen and menus.

9.1 Starting MegaRAID Storage Manager

Follow these steps to start the MegaRAID Storage Manager utility and view the main menu screen:

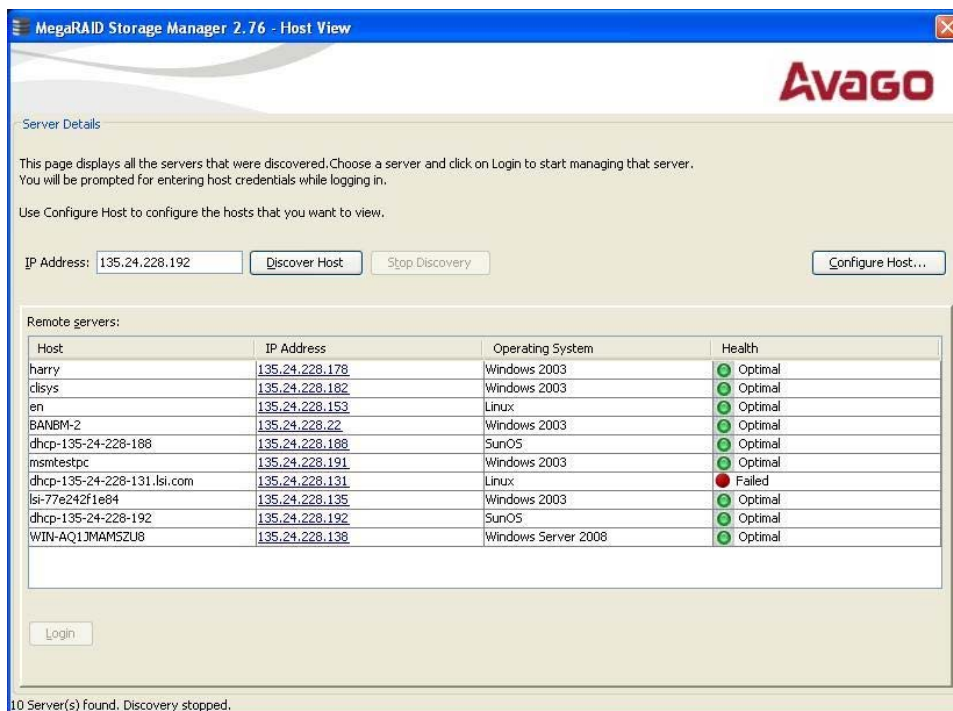
1. Start the program using the method required for your operating system environment:
 - To start the MegaRAID Storage Manager utility on a Microsoft Windows system, select **Start > Programs > MegaRAID Storage Manager > StartupUI**, or double-click the **MegaRAID Storage Manager** shortcut on the desktop.

NOTE If a warning appears stating that Windows Firewall has blocked some features of the program, click **Unblock** to allow the MegaRAID Storage Manager utility to start. (The Windows Firewall sometimes blocks the operation of programs that use Java® technology.)

- To start the MegaRAID Storage Manager utility on a Red Hat Linux system, select **Applications > System Tools > MegaRAID Storage Manager StartupUI**.
- To start the MegaRAID Storage Manager utility on a SuSE Linux or an SLES system, select **Start > System > More Programs > MegaRAID Storage Manager**.

When the program starts, the **Select Server** window appears, as shown in the following figure. The remote servers display, along with their IP address, operating system, and health status.

Figure 137 Select Server Window



15-1197_startup_hostViewWindow

If the circle in the server icon is orange instead of green, it means that the server is running in a degraded state, for example, because a drive used in a virtual drive has failed. If the circle is red, the storage configuration in the server has failed.

NOTE To access servers on a different subnet, type in the box at the bottom of the screen the IP address of a server in the desired subnet where the MegaRAID Storage Manager utility is running, and click **Update**. If you check the **Connect to remote server at: IP address** box, you can also access a stand-alone (remote) installation of the MegaRAID Storage Manager utility, if it has a network connection.

2. Double-click the icon of the server that you want to access.
The following screen appears.

Figure 138 Server Login Window



15-1197_startup_loginWindow

3. Enter your user name and password.
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.
4. 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 connected to a network, this procedure is the login to the computer itself, not the network login.

5. Enter the root/administrator user name and password to use the *Full Access* mode.


NOTE In a Linux system, users that belong to the root group can log in. You do not have to be the user *root*.

If your user name and password are correct for the Login mode you have chosen, the **MegaRAID Storage Manager** main menu appears.

9.2 MegaRAID Storage Manager Main Menu Screen

This section describes the **MegaRAID Storage Manager** main menu screen, which is shown in the following figure.

Figure 139 MegaRAID Storage Manager Main Menu Screen



15-1197_startup_mainMenuWindow

The following sections describe the panels and the menu options that appear on this screen.

9.2.1 Dashboard View/Physical View/Logical View

The left panel of the **MegaRAID Storage Manager** main menu screen displays the *Dashboard* view, the *Physical* view, or the *Logical* view of the system, and the related devices, depending on which tab is selected.

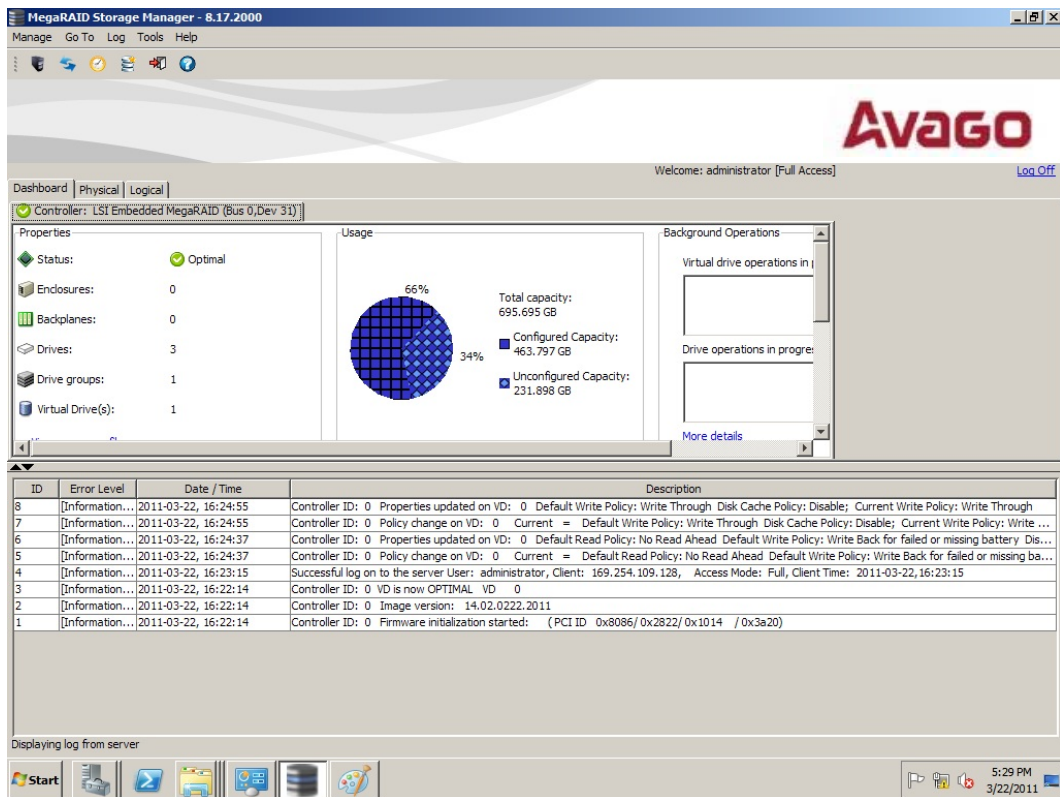
9.2.1.1 Dashboard View

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
- The MegaRAID Storage Manager utility features and their status (enabled or disabled)
- Actions you can perform
- Links to Online Help

The following screen appears.

Figure 140 MegaRAID Storage Manager Dashboard View



15-1197_startup_DashboardView

9.2.1.2 Physical View

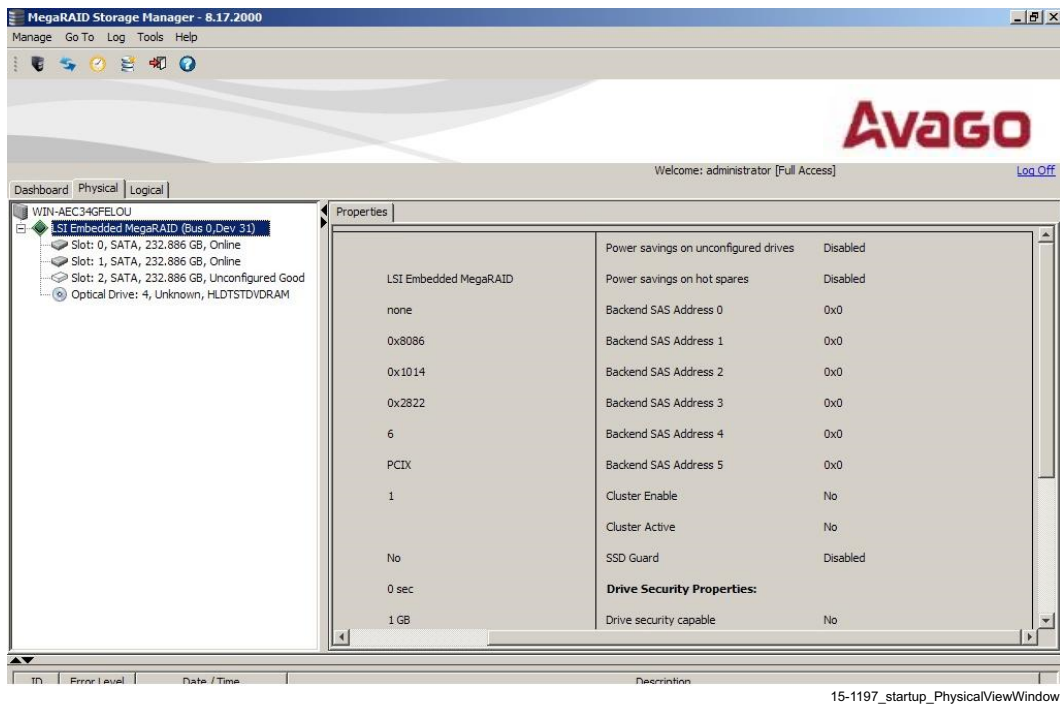
The *Physical* view shows the hierarchy of physical devices in the system. At the top of the hierarchy is the system itself, followed by the controller and the backplane.

One or more controllers are installed in the system. The controller label identifies the controller so that you can easily differentiate among multiple controllers. Each controller has one or more ports.

Drives and other devices are attached to the ports. The properties for each device appear in the right panel of the screen under the **Properties** tab.

The following screen appears.

Figure 141 MegaRAID Storage Manager Physical View

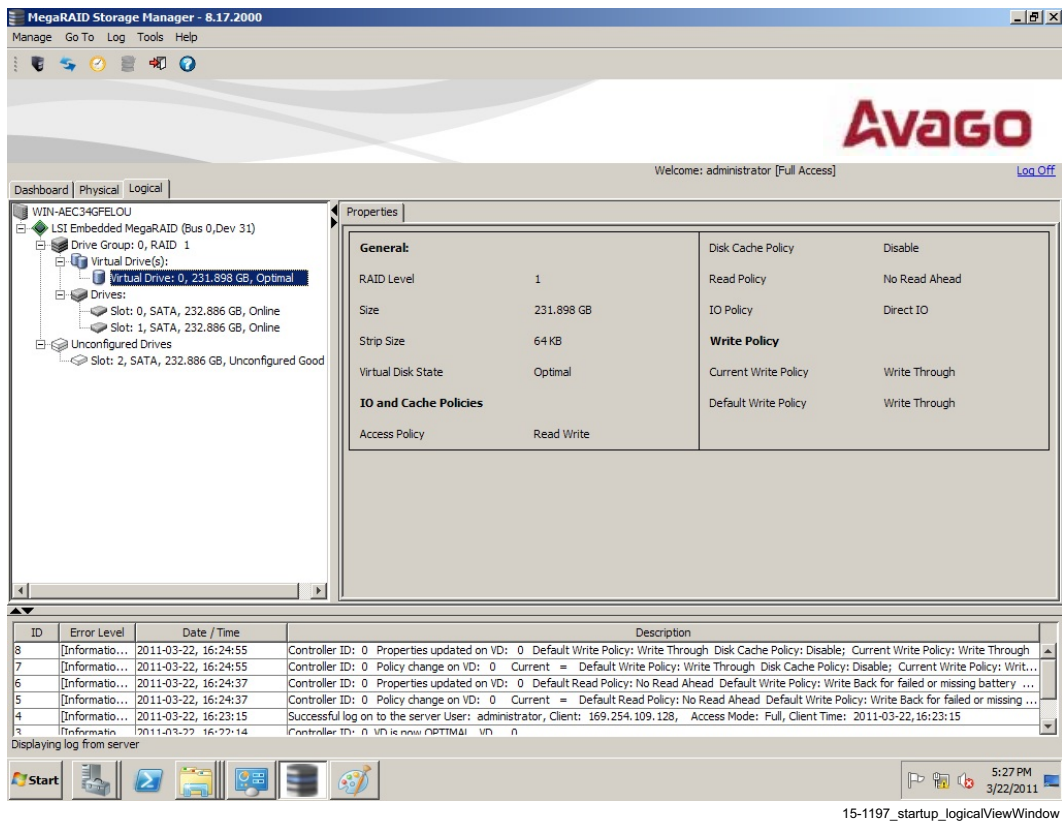


9.2.1.3 Logical View

The *Logical* view shows the hierarchy of controllers, virtual drives, and the drives and drive groups that make up the virtual drives. The properties for these components appear in the right panel under the **Properties** tab.

The following screen appears.

Figure 142 MegaRAID Storage Manager Logical View







The following icons in the left panel represent the controllers, drives, and other devices:


Table 10 Icons and Descriptions


Icon	Description
	Status
	System
	Controller
	Backplane
	Port
	Drive group
	Virtual drive
	Online drive
	Global hot spare
	Tape drive

Table 10 Icons and Descriptions (Continued)

Icon	Description
	CD-ROM
	Foreign Drive
	Unconfigured Foreign Drive
	Unconfigured Drive

NOTE The MegaRAID Storage Manager utility shows the icons for tape drive devices, but the utility does not support tape-related operations. 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 drive has failed: .

9.2.2 Event Log Panel

The lower part of the **MegaRAID Storage Manager** main menu screen 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.

For more information about the event log, see [Section 7.1, Monitoring System Events](#), on page 135. For more information about the event log entries, see [Appendix A, Events and Messages](#), on page 166.

9.2.3 Menu Bar

Here are brief descriptions of the main selections on the **MegaRAID Storage Manager** menu bar. Specific menu options are described in more detail in [Chapter 7, System Event and Storage Device Monitors](#), and this chapter.

9.2.3.1 Manage Menu

The **Manage** menu has a **Refresh** option for updating the display in the **MegaRAID Storage Manager** main menu screen (refresh is seldom required; the display normally updates automatically) and an **Exit** option to end your session on the MegaRAID Storage Manager utility. 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.

9.2.3.2 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, in which you access the Configuration Wizard to configure 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**.

9.2.3.3 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](#), on page 166.

9.2.3.4 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. For more information, see [Section 7.2, Configuring Alert Notifications](#), on page 137.

9.2.3.5 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 utility.

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® browser or the Mozilla® browser for the MegaRAID Storage Manager Online Help to display.

Appendix A: Events and Messages

The MegaRAID Storage Manager utility 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** main menu screen.

This appendix lists the MegaRAID Storage Manager utility events that can appear in the event log.

NOTE The MegaRAID Storage Manager utility can be used to manage a wide range of MegaRAID controllers. Some of the events and messages listed in this appendix are not applicable to the Embedded RAID configuration.

Each message that appears in the event log has an error level that indicates the severity of the event, as shown in the following table.

Table 11 Event Error Levels

Error Level	Description
Information	Informational message; no user action is necessary.
Warning	Some component may be close to a failure point.
Caution	A component has failed, but the system has not lost data.
Fatal	A component has failed, and data loss has occurred or will occur.
Dead	A catastrophic error has occurred, and the controller has died. This event is seen only after the controller has been restarted.

The following table lists all the MegaRAID Storage Manager utility event messages. The event message descriptions include placeholders for specific values that are determined when the event is generated. Some of the error messages are relevant only for hardware RAID.

Table 12 Event Messages

Number (Hex)	Number (Decimal)	Type	Event Text
0x0000	0	Information	Firmware initialization started (PCI ID %04x/%04x/%04x/%04x)
0x0001	1	Information	Firmware version %s
0x0002	2	Fatal	Unable to recover cache data from TBBU
0x0003	3	Information	Cache data recovered from TBBU successfully
0x0004	4	Information	Configuration cleared
0x0005	5	Warning	Cluster down; communication with peer lost
0x0006	6	Information	%s ownership changed from %02x to %02x
0x0007	7	Information	Alarm disabled by user
0x0008	8	Information	Alarm enabled by user
0x0009	9	Information	Background initialization rate changed to %d%%
0x000A	10	Fatal	Controller cache discarded due to memory/battery problems
0x000B	11	Fatal	Unable to recover cache data due to configuration mismatch
0x000C	12	Information	Cache data recovered successfully
0x000D	13	Fatal	Controller cache discarded due to firmware version incompatibility
0x000E	14	Information	Consistency Check rate changed to %d%%
0x000F	15	Dead	Fatal firmware error: %s

Table 12 Event Messages (Continued)

Number (Hex)	Number (Decimal)	Type	Event Text
0x0010	16	Information	Factory defaults restored
0x0011	17	Information	Flash downloaded image corrupt
0x0012	18	Caution	Flash erase error
0x0013	19	Caution	Flash timeout during erase
0x0014	20	Caution	Flash error
0x0015	21	Information	Flashing image: %s
0x0016	22	Information	Flash of new firmware image(s) complete
0x0017	23	Caution	Flash programming error
0x0018	24	Caution	Flash timeout during programming
0x0019	25	Caution	Flash chip type unknown
0x001A	26	Caution	Flash command set unknown
0x001B	27	Caution	Flash verify failure
0x001C	28	Information	Flush rate changed to %d seconds
0x001D	29	Information	Hibernate command received from host
0x001E	30	Information	Event log cleared
0x001F	31	Information	Event log wrapped
0x0020	32	Dead	Multi-bit ECC error: ECAR=%x, ELOG=%x, (%s)
0x0021	33	Warning	Single-bit ECC error: ECAR=%x, ELOG=%x, (%s)
0x0022	34	Dead	Not enough controller memory
0x0023	35	Information	Patrol Read operation complete
0x0024	36	Information	Patrol Read operation paused
0x0025	37	Information	Patrol Read Rate changed to %d%%
0x0026	38	Information	Patrol Read operation resumed
0x0027	39	Information	Patrol Read operation started
0x0028	40	Information	Rebuild rate changed to %d%%
0x0029	41	Information	Reconstruction rate changed to %d%%
0x002A	42	Information	Shutdown command received from host
0x002B	43	Information	Test event: %s
0x002C	44	Information	Time established as %s; (%d seconds since power on)
0x002D	45	Information	User entered firmware debugger
0x002E	46	Warning	Background Initialization aborted on %s
0x002F	47	Warning	Background Initialization corrected medium error (%s at %lx)
0x0030	48	Information	Background Initialization completed on %s
0x0031	49	Fatal	Background Initialization corrected medium error (%s at %lx, %s at %lx)
0x0032	50	Fatal	Background Initialization detected uncorrectable double medium errors (%s at %lx on %s)
0x0033	51	Caution	Background Initialization failed on %s
0x0034	52	Progress	Background Initialization progress on %s is %s
0x0035	53	Information	Background Initialization started on %s
0x0036	54	Information	Policy change on %s from %s to %s
0x0038	56	Warning	Consistency Check aborted on %s

Table 12 Event Messages (Continued)

Number (Hex)	Number (Decimal)	Type	Event Text
0x0039	57	Warning	Consistency Check corrected medium error (%s at %lx, %s at %lx)
0x003A	58	Information	Consistency Check done on %s
0x003B	59	Information	Consistency Check done with corrections on %s, (corrections=%d)
0x003C	60	Fatal	Consistency Check detected uncorrectable double medium errors (%s at %lx on %s)
0x003D	61	Caution	Consistency Check failed on %s
0x003E	62	Fatal	Consistency Check failed with uncorrectable data on %s
0x003F	63	Warning	Consistency Check found inconsistent parity on %s at strip %lx
0x0040	64	Warning	Consistency Check inconsistency logging disabled on %s (too many inconsistencies)
0x0041	65	Progress	Consistency Check progress on %s is %s
0x0042	66	Information	Consistency Check started on %s
0x0043	67	Warning	Initialization aborted on %s
0x0044	68	Caution	Initialization failed on %s
0x0045	69	Progress	Initialization progress on %s is %s
0x0046	70	Information	Fast initialization started on %s
0x0047	71	Information	Full initialization started on %s
0x0048	72	Information	Initialization complete on %s
0x0049	73	Information	Properties updated to %s (from %s)
0x004A	74	Information	Reconstruction complete on %s
0x004B	75	Fatal	Reconstruction of %s stopped due to unrecoverable errors
0x004C	76	Fatal	Reconstruct detected uncorrectable double medium errors (%s at %lx on %s at %lx)
0x004D	77	Progress	Reconstruction progress on %s is %s
0x004E	78	Information	Reconstruction resumed on %s
0x004F	79	Fatal	Reconstruction resume of %s failed due to configuration mismatch
0x0050	80	Information	Reconstructing started on %s
0x0051	81	Information	State change on %s from %s to %s
0x0052	82	Information	Clear aborted on %s
0x0053	83	Caution	Clear failed on %s (Error %02x)
0x0054	84	Progress	Clear progress on %s is %s
0x0055	85	Information	Clear started on %s
0x0056	86	Information	Clear completed on %s
0x0057	87	Warning	Error on %s (Error %02x)
0x0058	88	Information	Format complete on %s
0x0059	89	Information	Format started on %s
0x005A	90	Caution	Hot Spare SMART™ polling failed on %s (Error %02x)
0x005B	91	Information	Inserted: %s
0x005C	92	Warning	%s is not supported
0x005D	93	Warning	Patrol Read operation corrected medium error on %s at %lx
0x005E	94	Progress	Patrol Read operation progress on %s is %s
0x005F	95	Fatal	Patrol Read operation found an uncorrectable medium error on %s at %lx
0x0060	96	Caution	Predictive failure: %s

Table 12 Event Messages (Continued)

Number (Hex)	Number (Decimal)	Type	Event Text
0x0061	97	Fatal	Puncturing bad block on %s at %lx
0x0062	98	Information	Rebuild aborted by user on %s
0x0063	99	Information	Rebuild complete on %s
0x0064	100	Information	Rebuild complete on %s
0x0065	101	Caution	Rebuild failed on %s due to source drive error
0x0066	102	Caution	Rebuild failed on %s due to target drive error
0x0067	103	Progress	Rebuild progress on %s is %s
0x0068	104	Information	Rebuild resumed on %s
0x0069	105	Information	Rebuild started on %s
0x006A	106	Information	Rebuild automatically started on %s
0x006B	107	Caution	Rebuild stopped on %s due to loss of cluster ownership
0x006C	108	Fatal	Reassign write operation failed on %s at %lx
0x006D	109	Fatal	Unrecoverable medium error during rebuild on %s at %lx
0x006E	110	Information	Corrected medium error during recovery on %s at %lx
0X006F	111	Fatal	Unrecoverable medium error during recovery on %s at %lx
0x0070	112	Information	Removed: %s
0x0071	113	Warning	Unexpected sense: %s, CDB%s, Sense: %s
0x0072	114	Information	State change on %s from %s to %s
0x0073	115	Information	State change by user on %s from %s to %s
0x0074	116	Warning	Redundant path to %s broken
0x0075	117	Information	Redundant path to %s restored
0x0076	118	Information	Dedicated Hot Spare PD %s no longer useful due to deleted array
0x0077	119	Caution	SAS topology error: Loop detected
0x0078	120	Caution	SAS topology error: Unaddressable device
0x0079	121	Caution	SAS topology error: Multiple ports to the same SAS address
0x007A	122	Caution	SAS topology error: Expander error
0x007B	123	Caution	SAS topology error: SMP timeout
0x007C	124	Caution	SAS topology error: Out of route entries
0x007D	125	Caution	SAS topology error: Index not found
0x007E	126	Caution	SAS topology error: SMP function failed
0x007F	127	Caution	SAS topology error: SMP CRC error
0x0080	128	Caution	SAS topology error: Multiple subtractive
0x0081	129	Caution	SAS topology error: Table to table
0x0082	130	Caution	SAS topology error: Multiple paths
0x0083	131	Fatal	Unable to access device %s
0x0084	132	Information	Dedicated Hot Spare created on %s (%s)
0x0085	133	Information	Dedicated Hot Spare %s (%s) disabled
0x0086	134	Caution	Dedicated Hot Spare %s no longer useful for all arrays
0x0087	135	Information	Spare created on %s (%s)
0x0088	136	Information	Spare %s (%s) disabled

Table 12 Event Messages (Continued)

Number (Hex)	Number (Decimal)	Type	Event Text
0x0089	137	Caution	Spare %s does not cover all arrays
0x008A	138	Information	Created %s
0x008B	139	Information	Deleted %s
0x008C	140	Information	Marking %s inconsistent due to active writes at shutdown
0x008D	141	Information	Battery Present
0x008E	142	Warning	Battery Not Present
0x008F	143	Information	New Battery Detected
0x0090	144	Information	Battery has been replaced
0x0091	145	Caution	Battery temperature is high
0x0092	146	Warning	Battery voltage low
0x0093	147	Information	Battery started charging
0x0094	148	Information	Battery is discharging
0x0095	149	Information	Battery temperature is normal
0x0096	150	Fatal	Battery needs replacement - SOH Bad
0x0097	151	Information	Battery relearn started
0x0098	152	Information	Battery relearn in progress
0x0099	153	Information	Battery relearn completed
0x009A	154	Caution	Battery relearn timed out
0x009B	155	Information	Battery relearn pending: Battery is under charge
0x009C	156	Information	Battery relearn postponed
0x009D	157	Information	Battery relearn will start in 4 days
0x009E	158	Information	Battery relearn will start in 2 day
0x009F	159	Information	Battery relearn will start in 1 day
0x00A0	160	Information	Battery relearn will start in 5 hours
0x00A1	161	Information	Battery removed
0x00A2	162	Information	Current capacity of the battery is below threshold
0x00A3	163	Information	Current capacity of the battery is above threshold
0x00A4	164	Information	Enclosure (SES) discovered on %s
0x00A5	165	Information	Enclosure (SAFTE) discovered on %s
0x00A6	166	Caution	Enclosure %s communication lost
0x00A7	167	Information	Enclosure %s communication restored
0x00A8	168	Caution	Enclosure %s fan %d failed
0x00A9	169	Information	Enclosure %s fan %d inserted
0x00AA	170	Caution	Enclosure %s fan %d removed
0x00AB	171	Caution	Enclosure %s power supply %d failed
0x00AC	172	Information	Enclosure %s power supply %d inserted
0x00AD	173	Caution	Enclosure %s power supply %d removed
0x00AE	174	Caution	Enclosure %s EMM %d failed
0x00AF	175	Information	Enclosure %s EMM %d inserted
0x00B0	176	Caution	Enclosure %s EMM %d removed

Table 12 Event Messages (Continued)

Number (Hex)	Number (Decimal)	Type	Event Text
0x00B1	177	Warning	Enclosure %s temperature sensor %d below warning threshold
0x00B2	178	Caution	Enclosure %s temperature sensor %d below error threshold
0x00B3	179	Warning	Enclosure %s temperature sensor %d above warning threshold
0x00B4	180	Caution	Enclosure %s temperature sensor %d above error threshold
0x00B5	181	Caution	Enclosure %s shutdown
0x00B6	182	Warning	Enclosure %s not supported; too many enclosures connected to port
0x00B7	183	Caution	Enclosure %s firmware mismatch (EMM %d)
0x00B8	184	Warning	Enclosure %s sensor %d bad
0x00B9	185	Caution	Enclosure %s phy bad for slot %d
0x00BA	186	Caution	Enclosure %s is unstable
0x00BB	187	Caution	Enclosure %s hardware error
0x00BC	188	Caution	Enclosure %s not responding
0x00BD	189	Information	SAS/SATA mixing not supported in enclosure; %s disabled
0x00BE	190	Information	Enclosure (SES) hotplug on %s was detected, but is not supported
0x00BF	191	Information	Clustering enabled
0x00C0	192	Information	Clustering disabled
0x00C1	193	Information	PD too small to be used for auto-rebuild on %s
0x00C2	194	Information	Battery backup unit (BBU) enabled; changing Write Through virtual disks to Write Back
0x00C3	195	Warning	BBU disabled; changing Write Back virtual disks to Write Through
0x00C4	196	Warning	Bad block table on %s is 80% full
0x00C5	197	Fatal	Bad block table on %s is full; unable to log block %lx
0x00C6	198	Information	Consistency Check Aborted Due to Ownership Loss on %s
0x00C7	199	Information	BGI aborted due to ownership loss on %s
0x00C8	200	Caution	Battery/charger problems detected; SOH Bad
0x00C9	201	Warning	Single-bit ECC error: ECAR=%x, ELOG=%x, (%s); warning threshold exceeded
0x00CA	202	Caution	Single-bit ECC error: ECAR=%x, ELOG=%x, (%s); critical threshold exceeded
0x00CB	203	Caution	Single-bit ECC error: ECAR=%x, ELOG=%x, (%s); further reporting disabled
0x00CC	204	Caution	Enclosure %s Power supply %d switched off
0x00CD	205	Information	Enclosure %s Power supply %d switched on
0x00CE	206	Caution	Enclosure %s Power supply %d cable removed
0x00CF	207	Information	Enclosure %s Power supply %d cable inserted
0x00D0	208	Information	Enclosure %s Fan %d returned to normal
0x00D1	209	Information	BBU Retention test was initiated on previous boot
0x00D2	210	Information	BBU Retention test passed
0x00D3	211	Caution	BBU Retention test failed!
0x00D4	212	Information	NVRAM Retention test was initiated on previous boot
0x00D5	213	Information	NVRAM Retention test passed
0x00D6	214	Caution	NVRAM Retention test failed!
0x00D7	215	Information	%s test completed %d passes successfully
0x00D8	216	Caution	%s test FAILED on %d pass. Fail data: errorOffset=%x goodData=%x badData=%x

Table 12 Event Messages (Continued)

Number (Hex)	Number (Decimal)	Type	Event Text
0x00D9	217	Information	Self check diagnostics completed
0x00DA	218	Information	Foreign configuration detected
0x00DB	219	Information	Foreign configuration imported
0x00DC	220	Information	Foreign configuration cleared
0x00DD	221	Warning	NVRAM is corrupt; reinitializing
0x00DE	222	Warning	NVRAM mismatch occurred
0x00DF	223	Warning	SAS wide port %d lost link on PHY %d
0x00E0	224	Information	SAS wide port %d restored link on PHY %d
0x00E1	225	Warning	SAS port %d, PHY %d has exceeded the allowed error rate
0x00E2	226	Warning	Bad block reassigned on %s at %lx to %lx
0x00E3	227	Information	Controller hot plug detected
0x00E4	228	Warning	Enclosure %s temperature sensor %d differential detected
0x00E5	229	Information	Disk test cannot start; no qualifying disks found
0x00E6	230	Information	Time duration provided by host is not sufficient for self check
0x00E7	231	Information	Marked missing for %s on array %d row %d
0x00E8	232	Information	Replaced missing as %s on array %d row %d
0x00E9	233	Information	Enclosure %s temperature %d returned to normal
0x00EA	234	Information	Enclosure %s firmware download in progress
0x00EB	235	Warning	Enclosure %s firmware download failed
0x00EC	236	Warning	%s is not a certified drive
0x00ED	237	Information	Dirty cache data discarded by user
0x00EE	238	Information	PDs missing from configuration at boot
0x00EF	239	Information	VDs missing drives and will go offline at boot: %s
0x00F0	240	Information	VDs missing at boot: %s
0x00F1	241	Information	Previous configuration completely missing at boot
0x00F2	242	Information	Battery charge complete
0x00F3	243	Information	Enclosure %s fan %d speed changed
0x00F4	244	Information	Dedicated spare %s imported as global due to missing arrays
0x00F5	245	Information	%s rebuild not possible as SAS/SATA is not supported in an array
0x00F6	246	Information	SEP %s has been rebooted as a part of enclosure firmware download; SEP will be unavailable until this process completes.
0x00F7	247	Information	Inserted: %s Info: %s
0x00F8	248	Information	Removed: %s Info: %s
0x00F9	249	Information	%s is now OPTIMAL
0x00FA	250	Warning	%s is now PARTIALLY DEGRADED
0x00FB	251	Caution	%s is now DEGRADED
0x00FC	252	Fatal	%s is now OFFLINE
0x00FD	253	Warning	Battery requires reconditioning; please initiate a LEARN cycle
0x00FE	254	Warning	VD %s disabled because RAID-5 is not supported by this RAID key
0x00FF	255	Warning	VD %s disabled because RAID-6 is not supported by this controller

Table 12 Event Messages (Continued)

Number (Hex)	Number (Decimal)	Type	Event Text
0x0100	256	Warning	VD %s disabled because SAS drives are not supported by this RAID key
0x0101	257	Warning	PD missing: %s
0x0102	258	Warning	Puncturing of LBAs enabled
0x0103	259	Warning	Puncturing of LBAs disabled
0x0104	260	Critical	Enclosure %s EMM %d not installed
0x0105	261	Information	Package version %s
0x0106	262	Warning	Global affinity hot spare %s commissioned in a different enclosure
0x0107	263	Warning	Foreign configuration table overflow
0x0108	264	Warning	Partial foreign configuration imported, PDs not imported:%s
0x0109	265	Information	Connector %s is active
0x010A	266	Information	Board Revision %s
0x010B	267	Warning	Command timeout on PD %s, CDB:%s
0x010C	268	Warning	PD %s reset (Type %02x)
0x010D	269	Warning	VD bad block table on %s is 80% full
0x010E	270	Fatal	VD bad block table on %s is full; unable to log block %lx (on %s at %lx)
0x010F	271	Fatal	Uncorrectable medium error logged for %s at %lx (on %s at %lx)
0x0110	272	Information	VD medium error corrected on %s at %lx
0x0111	273	Warning	Bad block table on PD %s is 100% full
0x0112	274	Warning	VD bad block table on PD %s is 100% full
0x0113	275	Fatal	Controller needs replacement, IOP is faulty
0x0114	276	Information	CopyBack started on PD %s from PD %s
0x0115	277	Information	CopyBack aborted on PD %s and src is PD %s
0x0116	278	Information	CopyBack complete on PD %s from PD %s
0x0117	279	Progress	CopyBack progress on PD %s is %s
0x0118	280	Information	CopyBack resumed on PD %s from %s
0x0119	281	Information	CopyBack automatically started on PD %s from %s
0x011A	282	Critical	CopyBack failed on PD %s due to source %s error
0x011B	283	Warning	Early power off warning was unsuccessful
0x011C	284	Information	BBU field replaceable unit (FRU) is %s
0x011D	285	Information	%s FRU is %s
0x011E	286	Information	Controller hardware revision ID %s
0x011F	287	Warning	Foreign import shall result in a backward incompatible upgrade of configuration metadata
0x0120	288	Information	Redundant path restored for PD %s
0x0121	289	Warning	Redundant path broken for PD %s
0x0122	290	Information	Redundant enclosure EMM %s inserted for EMM %s
0x0123	291	Information	Redundant enclosure EMM %s removed for EMM %s
0x0124	292	Warning	Patrol Read operation cannot be started, because the PDs are either not ONLINE, or are in a VD with an active process, or are in an excluded VD
0x0125	293	Information	Copyback aborted by user on PD %s and src is PD %s
0x0126	294	Critical	Copyback aborted on hot spare %s from %s, as hot spare needed for rebuild

Table 12 Event Messages (Continued)

Number (Hex)	Number (Decimal)	Type	Event Text
0x0127	295	Warning	Copyback aborted on PD %s from PD %s, as rebuild required in the array
0x0128	296	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	297	Information	Copyback cannot be started as PD %s is too small for src PD %s
0x012A	298	Information	Copyback cannot be started on PD %s from PD %s, as SAS/SATA is not supported in an array
0x012B	299	Information	Microcode update started on PD %s
0x012C	300	Information	Microcode update completed on PD %s
0x012D	301	Warning	Microcode update timeout on PD %s
0x012E	302	Warning	Microcode update failed on PD %s
0x012F	303	Information	Controller properties changed
0x0130	304	Information	Patrol Read properties changed
0x0131	305	Information	CC schedule properties changed
0x0132	306	Information	Battery properties changed
0x0133	307	Warning	Periodic battery relearn is pending; please initiate manual learn cycle because Automatic Learn is not enabled
0x0134	308	Information	Drive security key created
0x0135	309	Information	Drive security key backed up
0x0136	310	Information	Drive security key from escrow, verified
0x0137	311	Information	Drive security key changed
0x0138	312	Warning	Drive security key, re-key operation failed
0x0139	313	Warning	Drive security key is invalid
0x013A	314	Information	Drive security key destroyed
0x013B	315	Warning	Drive security key from escrow is invalid
0x013C	316	Information	VD %s is now secured
0x013D	317	Warning	VD %s is partially secured
0x013E	318	Information	PD %s security activated
0x013F	319	Information	PD %s security disabled
0x0140	320	Information	PD %s is reprovisioned
0x0141	321	Information	PD %s security key changed
0x0142	322	Fatal	Security subsystem problems detected for PD %s
0x0143	323	Fatal	Controller cache pinned for missing or offline VD %s
0x0144	324	Fatal	Controller cache pinned for missing or offline VD: %s
0x0145	325	Information	Controller cache discarded by user for VD: %s
0x0146	326	Information	Controller cache destaged for VD %s
0x0147	327	Warning	Consistency check started on an inconsistent VD %s
0x0148	328	Warning	Drive security key failure, cannot access secured configuration
0x0149	329	Warning	Drive security password from user is invalid
0x014A	330	Warning	Detected error with the remote battery connector cable
0x014B	331	Information	Power state change on PD %s from %s to %s
0x014C	332	Information	Enclosure %s element (SES code 0x%x) status changed

Table 12 Event Messages (Continued)

Number (Hex)	Number (Decimal)	Type	Event Text
0x014D	333	Information	PD %s rebuild not possible because the HDD/CacheCade® software mix is not supported in a drive group
0x014E	334	Information	Copyback cannot be started on PD %s from %s, because the HDD/CacheCade software mix is not supported in a drive group
0x014F	335	Information	VD bad block table on %s is cleared
0x0150	336	Caution	SAS topology error: 0x%lx
0x0151	337	Information	VD cluster of medium errors corrected for %s at %lx (on %s at %lx)
0x0152	338	Information	Controller requests a host bus rescan
0x0153	339	Information	Controller repurposed and factory defaults restored
0x0154	340	Information	Drive security key binding updated
0x0155	341	Information	Drive security is in EKM mode
0x0156	342	Warning	Drive security failed to communicate with EKMS
0x0157	343	Information	%s needs key to be %s %s
0x0158	344	Warning	%s secure failed
0x0159	345	Critical	Controller encountered a fatal error and was reset
0x015A	346	Information	Snapshots enabled on %s (Repository %s)'
0x015B	347	Information	Snapshots disabled on %s (Repository %s) by the user
0x015C	348	Critical	Snapshots disabled on %s (Repository %s), due to a fatal error
0x015D	349	Information	Snapshot created on %s at %s
0x015E	350	Information	Snapshot deleted on %s at %s
0x015F	351	Information	View created at %s to a snapshot at %s for %s
0x0160	352	Information	View at %s is deleted, to snapshot at %s for %s
0x0161	353	Information	Snapshot rollback started on %s from snapshot at %s
0x0162	354	Fatal	Snapshot rollback on %s internally aborted for snapshot at %s'
0x0163	355	Information	Snapshot rollback on %s completed for snapshot at %s
0x0164	356	Information	Snapshot rollback progress for snapshot at %s, on %s is %s
0x0165	357	Warning	Snapshot space for %s in snapshot repository %s, is 80%% full
0x0166	358	Critical	Snapshot space for %s in snapshot repository %s, is full
0x0167	359	Warning	View at %s to snapshot at %s, is 80%% full on snapshot repository %s
0x0168	360	Critical	View at %s to snapshot at %s, is full on snapshot repository %s
0x0169	361	Critical	Snapshot repository lost for %s
0x016A	362	Warning	Snapshot repository restored for %s
0x016B	363	Critical	Snapshot encountered an unexpected internal error: 0x%lx
0x016C	364	Information	Auto snapshot enabled on %s (snapshot repository %s)
0x016D	365	Information	Auto snapshot disabled on %s (snapshot repository %s)
0x016E	366	Critical	Configuration command could not be committed to disk, please retry
0x016F	367	Information	COD on %s updated as it was stale
0x0170	368	Warning	Power state change failed on %s (from %s to %s)
0x0171	369	Warning	%s is not available
0x0172	370	Information	%s is available
0x0173	371	Information	%s is used for CacheCade with capacity 0x%lx logical blocks

Table 12 Event Messages (Continued)

Number (Hex)	Number (Decimal)	Type	Event Text
0x0174	372	Information	%s is using CacheCade %s
0x0175	373	Information	%s is no longer using CacheCade %s
0x0176	374	Critical	Snapshot deleted due to resource constraints for %s in snapshot repository %s
0x0177	375	Warning	Auto Snapshot failed for %s in snapshot repository %s
0x0178	376	Warning	Controller reset on-board expander
0x0179	377	Warning	CacheCade (%s) capacity changed and is now 0x%lx logical blocks
0x017A	378	Warning	Battery cannot initiate transparent learn cycles
0x017B	379	Information	Premium feature %s key was applied for - %s
0x017C	380	Information	Snapshot schedule properties changed on %s
0x017D	381	Information	Snapshot scheduled action is due on %s
0x017E	382	Information	Performance Metrics: collection command 0x%lx
0x017F	383	Information	Premium feature %s key was transferred - %s
0x0180	384	Information	Premium feature serial number %s
0x0181	385	Warning	Premium feature serial number mismatched. Key-vault serial num - %s
0x0182	386	Warning	Battery cannot support data retention for more than %d hours. Please replace the battery
0x0183	387	Information	%s power policy changed to %s (from %s)
0x0184	388	Warning	%s cannot transition to max power savings
0x0185	389	Information	Host driver is loaded and operational
0x0186	390	Information	%s mirror broken
0x0187	391	Information	%s mirror joined
0x0188	392	Warning	%s link %d failure in wide port
0x0189	393	Information	%s link %d restored in wide port
0x018A	394	Information	Memory module FRU is %s
0x018B	395	Warning	Cache-vault power pack is sub-optimal. Please replace the pack
0x018C	396	Warning	Foreign configuration auto-import did not import any drives
0x018D	398	Warning	Cache-vault microcode update required
0x018E	399	Warning	CacheCade (%s) capacity exceeds maximum allowed size, extra capacity is not used
0x018F	399	Warning	LD (%s) protection information lost
0x0190	400	Information	Diagnostics passed for %s
0x0191	401	Critical	Diagnostics failed for %s
0x0192	402	Information	Server power capability diagnostic test started
0x0193	403	Information	Drive cache settings enabled during rebuild for %s
0x0194	404	Information	Drive cache settings restored after rebuild for %s
0x0195	405	Information	Drive %s commissioned as Emergency spare
0x0196	406	Warning	Reminder: Potential non-optimal configuration due to drive %s commissioned as emergency spare
0x0197	407	Information	Consistency check suspended on %s
0x0198	408	Information	Consistency check resumed on %s
0x0199	409	Information	Background initialization suspended on %s
0x019A	410	Information	Background initialization resumed on %s
0x019B	411	Information	Reconstruction suspended on %s

Table 12 Event Messages (Continued)

Number (Hex)	Number (Decimal)	Type	Event Text
0x019C	412	Information	Rebuild suspended on %
0x019D	413	Information	Copyback suspended on %s
0x019E	414	Information	Reminder: Consistency check suspended on %
0x019F	415	Information	Reminder: Background initialization suspended on %s
0x01A0	416	Information	Reminder: Reconstruction suspended on %s
0x01A1	417	Information	Reminder: Rebuild suspended on %s
0x01A2	418	Information	Reminder: Copyback suspended on %s
0x01A3	419	Information	Reminder: Patrol Read operation suspended
0x01A4	420	Information	Erase aborted on %s
0x01A5	421	Critical	Erase failed on %s (Error %02x)
0x01A6	422	Progress	Erase progress on %s is %s
0x01A7	423	Information	Erase started on %s
0x01A8	424	Information	Erase completed on %s
0x01A9	425	Information	Erase aborted on %s
0x01AA	426	Critical	Erase failed on %s
0x01AB	427	Progress	Erase progress on %s is %s
0x01AC	428	Information	Erase started on %s
0x01AD	429	Information	Erase complete on %s
0x01AE	430	Warning	Potential leakage during erase on %s
0x01AF	431	Warning	Battery charging was suspended due to high battery temperature
0x01B0	432	Information	NVCache firmware update was successful
0x01B1	433	Warning	NVCache firmware update failed
0x01B2	434	Fatal	%s access blocked as cached data in CacheCade is unavailable
0x01B3	435	Information	CacheCade disassociate started on %s
0x01B4	436	Information	CacheCade disassociate completed on %s
0x01B5	437	Critical	CacheCade disassociate failed on %s
0x01B6	438	Progress	CacheCade disassociate progress on %s is %s
0x01B7	439	Information	CacheCade disassociate aborted by user on %s
0x01B8	440	Information	Link speed changed on SAS port %d and PHY %d
0x01B9	441	Warning	Advanced software options was deactivated for - %s
0x01BA	442	Information	%s is now accessible
0x01BB	443	Information	%s is using CacheCade
0x01BC	444	Information	%s is no longer using CacheCade
0x01BD	445	Information	Patrol Read operation aborted on %s

Glossary

A

AHCI

Advanced Host Controller Interface

AMI

American MegaTrends, Inc

array

a group of one to eight physical disks (recognized by the host computer system as one large disk drive) in a RAID drive group

B

BBS

BIOS boot specification

BBU

battery backup unit

BGI

background initialization
BGI is used in RAID 5 configurations only.

BIOS

basic input/output system

BP

broadcast processor

C

CLI

command line interface

CPU

central processing unit

CT

command tool

CU

configuration utility

D

DAT

digital audio tape

DHP

driver health protocol

DUD

driver update diskette

E

EDBA

extended BIOS data area

F

FGI

foreground initialization

FRU

field replaceable unit

G

GA

general availability

GUI

graphical user interface

H

HBA

host bus adapter

HII

Human Interface Infrastructure

I

IDE

integrated drive electronics

I/O

input/output

L

LD

logical disk

LED

light-emitting diode

LRU

least recently used

N

NVRAM

nonvolatile random access memory

O

OAF

open address frame

OS

operating system

P

PCI

Peripheral Component Interconnect

PCIe

PCI Express

PD

physical drive

PMM

post memory management

POST

power-on self-test

R

RAID

Redundant Array of Independent Disks

RHEL

Red Hat Enterprise Linux

S

SAS

Serial SCSI

SATA

Serial Advanced Technology Attachment

SBBM

soft bad block management

SCSI

Small Computer System Interface

SLES

SuSE Linux Enterprise Server

SSD

solid state drive

strip

the size of each data segment

stripe size

A virtual drive property that indicates the length of the interleaved data segments that the RAID controller writes across multiple drives, does not include 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.

SuSE

Gesellschaft für Software-und Systementwicklung MBH

SWR

software RAID

T

TBBU

transportable battery backup unit

U

UA

unit attention

UEFI

Unified Extensible Firmware Interface

USB

Universal Serial Bus

V

VD

virtual drive

VM

virtual machine

Revision History

pub-005085, Revision 2.2, March 19, 2018

Updated [Section 2.2.10.2, Updating or Loading Software RAID Driver Kernel RPM on SuSE 12](#)

pub-005085, Revision 2.1, January 15, 2015

- Updated [Section 2.2.3, Installing Red Hat Linux 6.x.](#)
- Updated [Section 2.2.4, Installing Red Hat Linux 7.x.](#)
- Updated [Section 2.2.5, Installing the Red Hat Linux Driver on a New System.](#)
- Updated [Section 2.2.8, Known Restrictions for the Driver Installation Process.](#)
- Updated [Section 2.2.9, Installing the SuSE Linux Enterprise Server Operating System Driver.](#)
- Updated [Section 2.2.9.1, Installing SuSE Linux Enterprise Server 10, SP-4 Drivers.](#)
- Updated [Section 2.2.9.2, Installing the SLES-10, SP-4 XEN Drivers.](#)
- Updated [Section 2.2.10, Installing the SuSE Linux Enterprise Server 11 and 12 Driver.](#)
- Updated [Section 2.2.10.1, Loading the Xen Kernel RPM on SLES-11.](#)
- Updated [Section 2.2.11, Installing the SLES Operating System on Grantley Servers.](#)

pub-005085, Revision 2.0, September 1, 2015

- Changed document number from Agile number to new Avago number.
- Removed Chapter 5, MegaCLI Command Tool.
- Replaced MegaCLI with reference to StorCLI document.
- Updated [Section 2.1, Windows Driver Installation.](#)
- Updated [Section 2.1.1, Windows 7, Windows 2008, Windows Vista, Windows 10, and Windows 2012 Driver Installation.](#)
- Updated [Section 2.2.1, Obtaining the Driver Image File.](#)
- Updated [Section 2.2.2.2, Preparing Installation Disks with the Red Hat Linux 4 and 5 and SuSE Linux Operating Systems.](#)
- Added [Section 2.2.3, Installing Red Hat Linux 6.x.](#)
- Added [Section 2.2.4, Installing Red Hat Linux 7.x.](#)
- Updated [Section 2.2.7, Enabling RAID Mode during Red Hat Linux 5 Driver Installation.](#)
- Updated [Section 2.2.9, Installing the SuSE Linux Enterprise Server Operating System Driver.](#)
- Added [Section 2.2.9.1, Installing SuSE Linux Enterprise Server 10, SP-4 Drivers.](#)
- Added [Section 2.2.9.2, Installing the SLES-10, SP-4 XEN Drivers.](#)
- Added [Section 2.2.10, Installing the SuSE Linux Enterprise Server 11 and 12 Driver.](#)
- Added [Section 2.2.10.1, Loading the Xen Kernel RPM on SLES-11.](#)
- Added [Section 2.2.11, Installing the SLES Operating System on Grantley Servers.](#)
- Updated [Section 3.1, Performing a Quick Configuration.](#)
- Updated [Section 3.5.1, Selecting the Configuration Method.](#)
- Updated [Section 3.9, Rebuilding a Drive.](#)

- Updated [Section 3.12.3, Viewing or Changing Virtual Drive Properties](#).
- Updated [Section 7.5, Running a Patrol Read Operation](#).
- Updated graphics.
- Performed minor edits for clarity and consistency.

48712-00, Rev. G, July 2014

- Updated [Section 4.1.1, Viewing Controller Properties](#).
- Updated [Section 4.2, Managing Virtual Drives](#).

48712-00, Rev. F, February 2014

- Added [Section 4.4, UDK2010 Support](#).
- Updated [Section 4.2.3, Selecting Virtual Drive Operations](#).

48712-00, Rev. E, October 2013

- Updated the initial section of [Chapter 4: Human Interface Infrastructure Configuration Utility](#).
- Updated [Section 1.1, Embedded RAID Software Features](#).

48712-00, Rev. D, April 2013

- Added [Section 1.1.6, Unified Extensible Firmware Interface Features](#).
- Added [Chapter 4: Human Interface Infrastructure Configuration Utility](#).

48712-00, Rev. C, February 2013

Added up to 4 TB of support in [Section 1.1.1, Device Support](#).

48712-00, Rev. B, March 2012

Revised the guide to document changes to the driver installation procedures, configuration utilities, and new event messages.

48712-00, Rev. A, June 2011

Initial release of the document.

