Firmware Operation Manual

for

Fibre-to-SAS/SATA RAID

Version 1.8 (Mar., 2008)

Firmware Revision: 3.53.x/3.61.x/3.63.x

Contact Information

Asia Pacific
(International Headquarters)
Infortrend Technology, Inc.
8F, No. 102 Chung-Shan Rd., Sec. 3
Chung-Ho City, Taipei Hsien, Taiwan
Tel: +886-2-2226-0126
Fax: +886-2-2226-0020
sales.ap@infortrend.com
support.ap@infortrend.com
http://esupport.infortrend.com.tw
http://www.infortrend.com.tw

Americas
Infortrend Corporation
2200 Zanker Road, Unit D,
San Jose, CA. 95131
USA
Tel: +1-408-988-5088
Fax: +1-408-988-6288
sales.us@infortrend.com
http://esupport.infortrend.com
http://www.infortrend.com

China
Infortrend Technology, Limited
Room 1210, West Wing, Tower One, Junefield Plaza, No. 6 Xuanwumen Street, Xuanwu District, Beijing, China
Post code: 100052
Tel: +86-10-6310-6168
Fax: +86-10-6310-6188
sales.cn@infortrend.com
support.cn@infortrend.com
http://esupport.infortrend.com.tw
http://www.infortrend.com.cn

Europe (EMEA)
Infortrend Europe Limited
1 Cherrywood, Stag Oak Lane
Chineham Business Park
Basingstoke, Hampshire
RG24 8WF, UK
Tel: +44-1256-707-700
Fax: +44-1256-707-889
sales.eu@infortrend.com
support.eu@infortrend.com
http://esupport.infortrend-europe.com/
http://www.infortrend.com

Japan
Infortrend Japan, Inc.
6F, Okayasu Bldg.,
1-7-14 Shibaura Minato-ku,
Tokyo, 105-0023 Japan
Tel: +81-3-5730-6551
Fax: +81-3-5730-6552
sales.jp@infortrend.com
support.jp@infortrend.com
http://esupport.infortrend.com.tw
http://www.infortrend.co.jp

Germany
Infortrend Deutschland GmbH
Werner-Eckert-Str.8
81829 Munich Germany
Tel: +49 (0)89 45 15 18 7 - 0
Fax: +49 (0)89 45 15 18 7 - 65
sales.de@infortrend.com
support.eu@infortrend.com
http://www.infortrend.com/germany

Copyright 2007

This Edition First Published 2007

Disclaimer

Infortrend Technology makes no representations or warranties with respect to the contents hereof and specifically disclaims any implied warranties of merchantability or fitness for any particular purpose. Furthermore, Infortrend Technology reserves the right to revise this publication and to make changes from time to time in the content hereof without obligation to notify any person of such revisions or changes. **Product specifications are also subject to change without notice.**

Trademarks

Infortrend, Infortrend logo, EonStor, RAIDWatch, and SANWatch are all registered trademarks of Infortrend Technology, Inc. Other names prefixed with "IFT" and "ES" are trademarks of Infortrend Technology, Inc.

PowerPC® is a trademark of International Business Machines Corporation and Motorola Inc.

Solaris and Java are trademarks of Sun Microsystems, Inc.

All other names, brands, products or services are trademarks or registered trademarks of their respective owners.

Table of Contents

Chapter Four Optimization & Preparation Tasks

Chapter Five Creating RAID Arrays & Host LUN Mapping (Using LCD Keypad Panel)

Chapter Six Creating RAID Arrays & Host LUN Mapping (Using Terminal Access)

Chapter Seven Fibre Channel Options

Chapter Eight Host-side and Drive-side Parameters

Chapter Nine Enclosure Management

Chapter Ten Data Integrity

Chapter Eleven Array Expansion

Chapter Twelve S.M.A.R.T. Support

Chapter Thirteen AV Optimization

Chapter Fourteen Redundant Controller

Appendix A Firmware Functionality Specifications

Appendix B System Functions: Upgrading Firmware

Appendix C Event Messages

About This Manual

This manual provides all of the necessary information that a system administrator needs to configure and maintain EonStor Fibre-to-SAS or Fibre-to-SATA RAID subsystems.

There are several other manuals that came with your RAID subsystem:

- **Installation and Hardware Reference Manual**:
  For hardware-related information.

- **RAIDWatch User's Manual**
  For the Java-based, GUI RAIDWatch manager providing remote and concurrent management of RAID subsystems.

- **SANWatch User's Manual**
  For the latest Java-based, SANWatch GUI manager that provides all standard RAID configuration, monitoring, and the Snapshot data protection functionality.

- **Embedded RAIDWatch User's Manual**
  For the firmware-embedded, browser-based, GUI RAIDWatch manager providing remote and concurrent management of RAID subsystems without the need to install software on host.

- **RAID Planning Guide**
  For RAID basic concepts, configurations and subsystem management. Similar concerns are also included in Chapter 4 of this manual.

The order of the chapters is arranged in accordance with the steps necessary for creating RAID configurations and tuning the characteristics of a RAID array.

The terminal screen displays and the LCD messages may vary when using subsystems running different firmware revisions.

Chapter 1	Shows users how to make RS-232 serial port or Ethernet connections for using terminal emulation or telnet access to the firmware-embedded configuration utility.
Chapter 2	Helps users understand screen elements and messages on the LCD display.

Chapter 3	Teaches users how to interpret the information found on the RS-232C terminal emulation.
Chapter 4	Teaches users how to setup the optimal settings, caching parameters, and other preparation tasks before creating RAID arrays.
Chapter 5	Provides step-by-step instructions on creating a RAID using the LCD keypad panel.
Chapter 6	Gives step-by-step instructions on creating a RAID via a Telnet or RS-232C serial port session.
Chapter 7	Includes all the Fibre Channel-related functions implemented since the firmware release 3.12.
Chapter 8	Provides the advanced options for the host and drive channel configurations including I/O queuing, number of LUNs, in-band, and disk access characteristics settings.
Chapter 9	Shows how to supervise system operating status.
Chapter 10	Details various methods, either manual or automated, for maintaining data integrity and preventing faults.
Chapter 11	Shows how to expand a configured RAID array.
Chapter 12	Shows how to use the SMART-related drive error preventative methods.
Chapter 13	Describes the details for the AV optimization mode implementations.
Chapter 14	Provides configurable options and working theories for redundant controller configurations.
Appendix A	Lists the important firmware features supported by the firmware revision, arranged in accordance with the latest firmware release as of press date.
Appendix B	Shows how to upgrade firmware and boot record.
Appendix C	Lists all of the controller event messages.

Revision History

- Rev. 1.0: Initial release.
- Rev. 1.1: Added new functions such as the Drive Delayed Write.
- Rev. 1.2: Revised descriptions of subordinate firmware functionality.
- Rev. 1.3: Added new functions such as Auto Shutdown.
- Rev. 1.4: Updated information about different host/drive interfaces
- Rev. 1.5: Revised Chapter 12 and added controller events.
- Rev. 1.6: Revised all chapters,
    - Added Chapter 1 for management access connections.
    - Added FW 3.61 new functions such as Clear Core Dump.
    - Moved pre-array preparation tasks to Chapter 4.
    - Revised Chapter 5 and Chapter 6 to include only the array configuration procedures.
    - Revised other chapters such as Chapter 13 and Chapter 14 to reflect changes such as LUN mapping methods and AV optimization theories.
    - Revised Appendix C Event Messages; added Data Service (snapshot) related events.
    - Updated screen information using LCD and terminal interfaces.
    - Added event viewing keys such as the ENT button on the LCD keypad panel and the Space key using terminal emulation.
    - Removed Drive Check Time, updated Auto Rebuild on Drive Swap, and removed array optimization Random mode.
    - Removed 4K and 8K stripe sizes.
    - Added SAS expansion enclosure check drive screens.
    - Added terminal mode Gauge range tuning details.
- Rev. 1.7: Added and corrected information about the cross-controller mapping methodology in Chapter 4, 5, 6, and 14.
- Rev. 1.8: Revised for firmware revision 3.63
    - Added a new option for handling a degraded drive through the SMART predictive monitoring. An unreliable drive can be

manually disbanded from an array in order to immediately commence a rebuild.

- Firmware supports 4KBytes block size for allocating RAID volumes larger than 2TB in 32-bit Windows XP/Windows 2000 through the Embedded RAIDWatch utility. Please refer to Infortrend's Q&A section on web.

  There are no configuration items for the array block size in firmware terminal console or LCD keypad interface.

- A CPU threshold is added for triggering a raise in thermal fan rotation speed. The threshold values are configured through OEM append files.

- Added a flow chart in Chapter 5 for restoring firmware defaults using the LCD keypad panel.

- Added definitions for the I^2C slave device failure events in the Appendix C.

Who should read this manual?

This manual assumes that its readers are experienced with computer hardware installation and are familiar with storage enclosures.

Related Documentation

- Installation and Hardware Reference Manual
- SANWatch User's Manual
- RAIDWatch User's Manual
- Firmware release note (can be accessed through Infortrend's VIP Room or eSupport websites)

These documents can be found in the product utility CD included with your subsystem package.

Conventions

Naming

From this point on and throughout the rest of this manual, the EonStor series is referred to as simply the "subsystem" or the "system" and EonStor is frequently abbreviated as "ES."

Lists

Bulleted Lists: Bulleted lists are statements of non-sequential facts. They can be read in any order. Each statement is preceded by a square or round black dot "•."

Numbered Lists: Numbered lists are used to describe sequential steps you should follow in order.

Important information that users should be aware of is indicated with the following icons:

NOTE:

These messages inform the reader of essential but non-critical information. These messages should be read carefully as any directions or instructions contained therein can help you avoid making mistakes.

CAUTION!

Cautionary messages should also be heeded to help you reduce the chance of losing data or damaging the system.

IMPORTANT!

Important messages remind users of critical configuration parameters and design concerns that are directly or indirectly related to successful installation.

WARNING!

Warnings appear where overlooked details may cause damage to the equipment or result in personal injury. Warnings should be taken seriously.

Software and Firmware Updates

Please contact your system vendor or visit Infortrend's VIProom, FIP site, or eSupport websites for the latest software or firmware updates.

However, problems that occur during the updating process may cause irrecoverable errors and system down time. Always consult technical personnel before proceeding with any firmware upgrade.

IMPORTANT!

The firmware version installed on your system should provide the complete functionality listed in the specification sheet/user's manual. We provide special revisions for various application purposes. Therefore, please **DO NOT** upgrade your firmware unless you fully understand what a firmware revision will do.

Chapter

1

Accessing the Array:

Serial Port and Ethernet Management Console

This chapter describes how to establish the management access to your RAID subsystem. The main topics include the following:

1.1 RS-232C Serial Port

Infortrend's controllers and subsystems can be configured via a PC running a VT-100 terminal emulation program or a VT-100-compatible terminal. RAID enclosures usually provide one or more DB-9 or audio jack RS-232C serial ports. Simply use an included serial cable to connect the subsystem/controller's serial port to the serial (COM) port on your management computer.

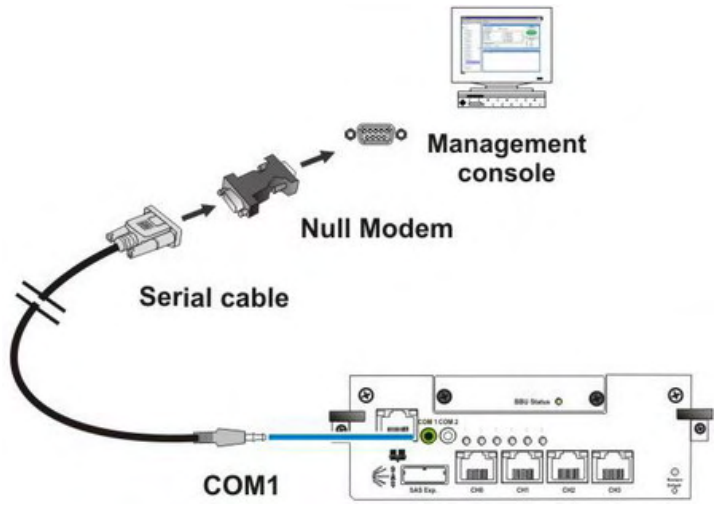

Figure 1-1: Management Consol via COM1 – Single RAID Controller

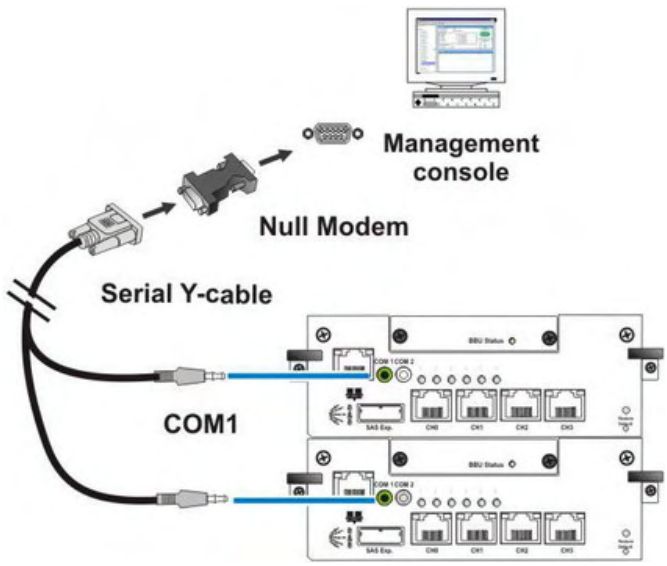

Figure 1-2: Management Consol via COM1 – Dual Controllers Using a Y-cable

Make sure you use the included null modem (IFT-9011) to convert the serial port signals. A null modem is always provided with your RAID subsystem/controller. The null modem has the serial signals swapped for connecting to a standard PC serial interface.

The following are guidelines on using the serial port:

- The serial port's defaults are:

Table 1-1 Serial Port Defaults

Baud rate	38400
Data bit	8
Parity	none
Stop bit	1

Flow control	Hardware

- Use the COM1 serial port for a terminal management session.
- In most cases, connecting RxD, TxD, and signal ground is enough to establish communication with a terminal.
- If you are using a PC as a terminal, any VT-100 terminal emulation software will suffice. Microsoft® Windows includes a terminal emulation program in the **Start** -> **Programs** -> **Accessories** -> **Communications** menu.
- For other details on connecting to the serial port, please refer to the **Installation and Hardware Reference Manual** that came with your subsystem.

1.1.1 Starting a Terminal Program (Windows XP Hyper Terminal)

Step 1. Locate and start a hyper terminal program in the Windows program menu: **Start** -> **Accessories** -> **Communications** -> **Hyper Terminal**.

Step 2. The program starts with an input field requiring you to enter a name for the connection.

Step 3. The next screen requires you to select a serial port on your PC.

Step 4. Select appropriate baud rate and data/stop bit values (identical to those set for the COM1 port on your RAID subsystem). Click OK, and you should then be able to establish a management console.

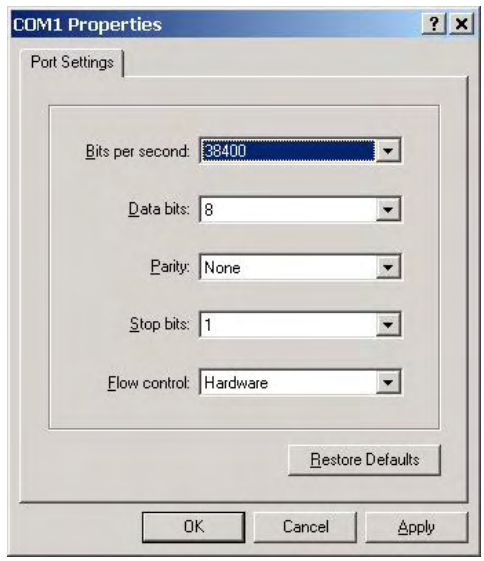

1.1.2 Starting RS-232C Terminal Emulation

The keys used when operating via the terminal are as follows:

←→↑↓	To select options
[Enter]	To go to a sub-menu or to execute a selected option
[Esc]	To cancel an option or return to the previous menu
[Ctrl]+[L]	To refresh the screen information

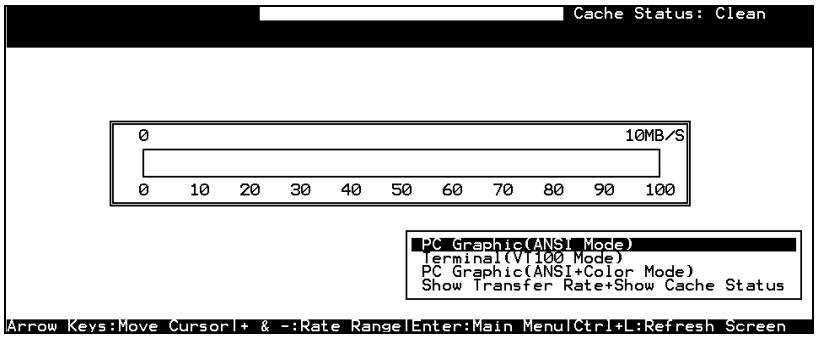

RS232C Initial Screen

NOTE:

If you connect the RS-232C cable while the RAID subsystem is powered on, press [Ctrl] + [L] to refresh the screen information. Key combinations may vary depending on the terminal software you use.

The initial screen appears when the controller finishes its self-test and is properly initialized. Use ↑ ↓ arrow keys to select terminal emulation mode, then press **[ENTER]** to enter the Main Menu.

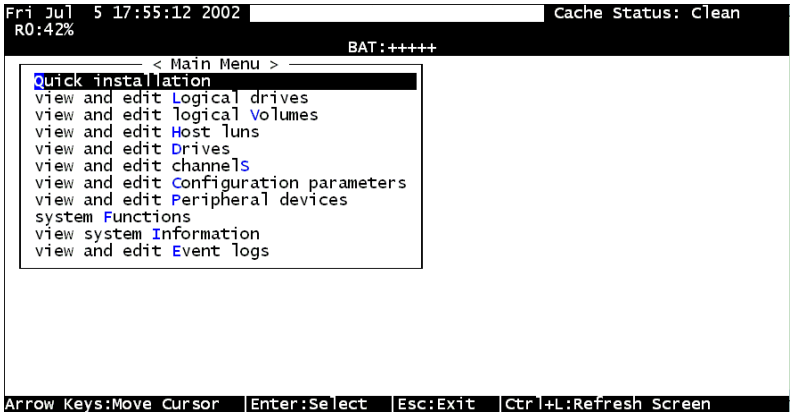

Functional Main Manu

Choose a functional item from the Main Menu to begin configuring your RAID.

1.2 Communication Parameters

The **Communication Parameters** is the first sub-menu under the "**View and Edit Configuration Parameters**" menu. In addition to the baud rate and terminal emulation options which have been discussed earlier, the sub-menu contains other options to prepare your management session using the Ethernet connection.

To access the sub-menu, use your arrow keys to select "**View and Edit Configuration Parameters**," and press **[ENTER]** on "**Communication Parameters**." This provides menu accesses to "**RS-232 Port Configuration**" and "**Internet Protocol <TCP/IP>**."

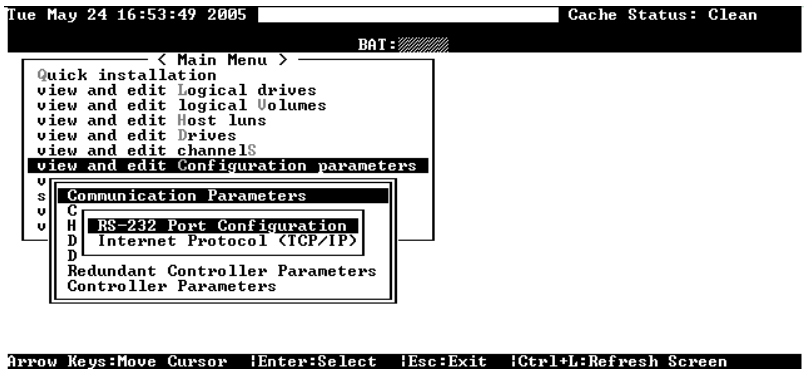

1.2.1 RS-232 Port Configuration

The "**RS-232 Port Configuration**" provides access to change the serial port operating parameters. Each COM port (COM1 or COM2) selection menu features two communication parameters: "**Baud Rate**" and "**Terminal Emulation**."

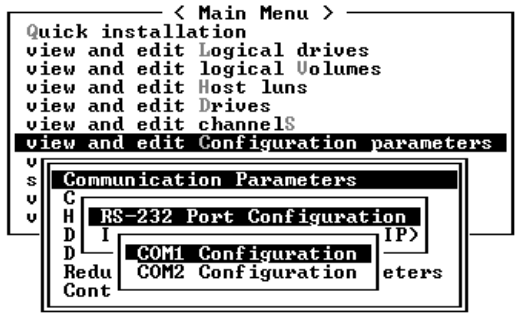

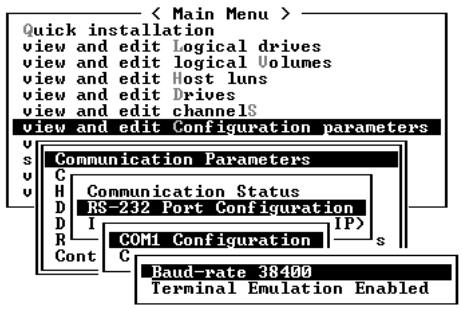

NOTE:

COM2 is used for UPS connection on the EonStor subsystems. Unlike previous models, the COM2 port on the EonStor system can no longer be used for management access.

1.2.2 Terminal Emulation:

The Terminal Emulation setting on the COM1 port is enabled by default. Usually there is no need to change this setting.

1.2.3 Baud Rate

To change the Baud Rate setting, select the highlighted menu item by pressing **[ENTER]**. Available options will be displayed in a pull-down menu. Select by pressing **[ENTER]** and press **ESC** several times to return to the previous configuration screen.

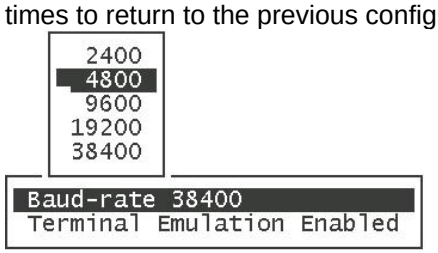

1.2.4 Internet Protocol <TCP/IP>

The Internet Protocol menu allows you to prepare the management access through the subsystem/controller's RJ-45 Ethernet port.

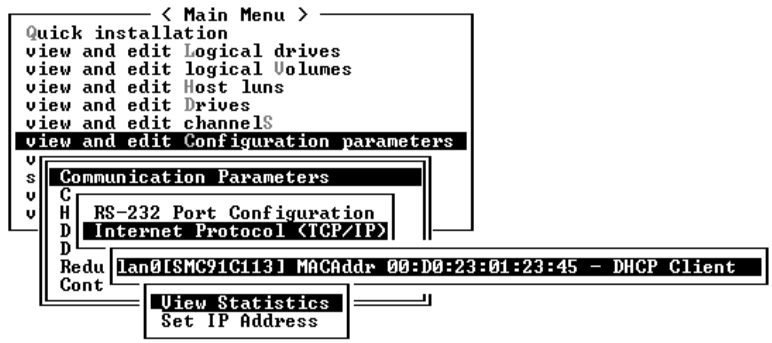

To access the configuration options, press **[ENTER]** on "**Internet Protocol <TCP/IP>**" to display the information of Ethernet port. Press **[ENTER]** on the chip information to display the "**View Statistics**" and the "**Set IP Address**" options.

1.2.5 View Statistics

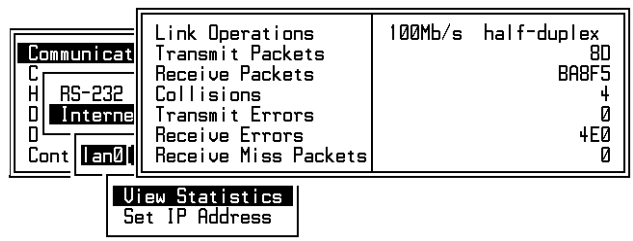

This window displays the current Ethernet link status.

1.2.6 Set the IP Address

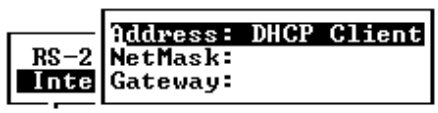

Provide a valid IP address for your subsystem/controller's Ethernet port. Consult your network administrator for an IP address and the associated NetMask and Gateway values. You may also key in "DHCP" if your local network supports automatic IP configuration.

NOTE:

The IP default is "DHCP client." However, if DHCP server can not be found within several seconds, a default IP address "10.10.1.1" will be loaded. This feature is available in the EonStor ASIC400 models.

One drawback of using DHCP is that if cable disconnection or other unpredictable network faults occur, your Ethernet port may be assigned with a different IP. This may cause problems for the management sessions using the RAIDWatch® Manager. You may not be able to receive important event messages before you

access the array by re-entering the new IP address assigned to the array.

RAIDWatch/SANWatch and its sub-module, the Configuration Client (the event notification utility), access a configured RAID system through its Ethernet port IP.

It may take several minutes to obtain an IP address from the DHCP server.

1.3 Out-of-Band via Ethernet

1.3.1 Connecting Ethernet Port:

Use a LAN cable to connect the Ethernet port(s) on the subsystem's RAID controller unit(s). Connect the cables between the subsystem/controller's Ethernet port and an Ethernet port on your local network.

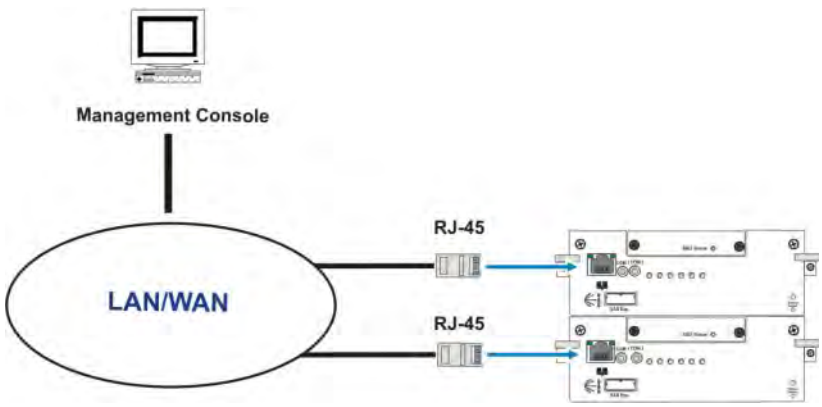

Figure 1-3: Ethernet Physical Connections

Single-controller RAID: Connect the Ethernet interface.

Dual-controller RAID: Connect the Ethernet interfaces from both controllers to your Ethernet network. The Ethernet port on the Secondary controller stays idle and becomes active in the event of Primary controller failure. The Ethernet port IP on the Primary's Ethernet port will be inherited by the Secondary controller during a controller failover process.

1.3.2 Configuring the Controller

To prepare the subsystem/controller for Ethernet connection:

1. Use a Terminal Emulator to Begin Configuration

Step 1. Connect the subsystem's serial port to a PC running a VT-100 terminal emulation program or a VT-100-compatible terminal using the included serial cable and null modem.

Step 2. Make sure the included null modem is already attached to the enclosure serial port or the management computer's COM port. The null modem converts the serial signals for connecting to a standard PC serial interface.

2. Assign an IP Address to the Ethernet Port

Step 1. Select "**View and Edit Configuration Parameters**" from the main menu. Select "**Communication Parameters**" -> "**Internet Protocol (TCP/IP)**" -> press **[ENTER]** on the chip hardware address -> and then select "**Set IP Address**."

Step 2. You may also use an auto discovery protocol such as DHCP. Simply key in "**DHCP**" in the IP address field.

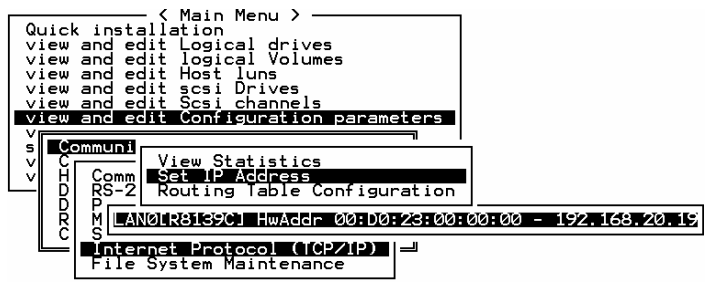

Step 3. Provide the IP address, NetMask, and Gateway values accordingly.

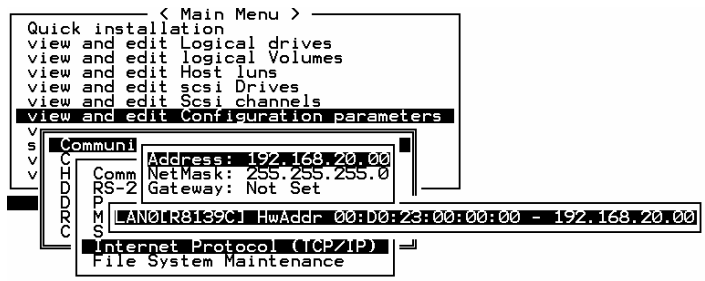

Step 4. Reset your RAID subsystem for the configuration to take effect.

Step 5. PING the IP address from your management computer to make sure the link is up and running.

1.4 Telnet Connection

Step 1. Use an Ethernet cable with RJ-45 phone jacks to connect the Ethernet port on the subsystem/controller module.

Step 2. Connect the other end of the Ethernet cable to your local area network. An IP address should be acquired for the subsystem's Ethernet port. The subsystem firmware also supports automatic client configuration such as DHCP.

Step 3. Consult your network administrator for an IP address that will be assigned to the subsystem/controller Ethernet port.

Step 4. Select "**View and Edit Configuration Parameters**" from the Main Menu on the terminal screen. Select "**Communication Parameters**" -> "**Internet Protocol (TCP/IP)**" -> press **ENTER** on the chip hardware address -> and then select "**Set IP Address**."

Step 5. Provide the IP address, NetMask, and Gateway values accordingly.

Step 6. PING the IP address from your management computer to make sure the link is valid.

Step 7. Open a command prompt window and key in "telnet xxx.xxx.xx.xxx (IP address)" to access the embedded firmware utility.

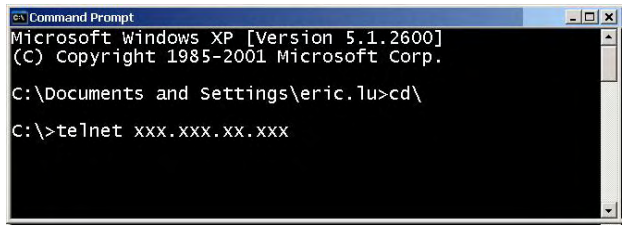

The default port number is 23.

1.5 Secure Link over SSH

Firmware supports remote management over the network connection and the security under SSH (Secure Shell) protection. SSH is widely used for its ability to provide strong authentication and secure communications over insecure channels. SSH secure

access can also be found as an option in the **RAIDWatch** management software.

SSH is more readily supported by Linux- or Unix-based systems. The support for SSH on Microsoft Windows platforms can be limited.

For making SSH link using Windows, there are SSH tools such as the “PuTTY” shareware.

To make an SSH link, use “root” as the default user name. If you have configured a **controller name** and **password** for your EonStor system, use it as your login name and password.

If a shareware is used, it may be necessary to configure the display options, e.g., the “Character set translation on received data” and “font type” setting in order for the terminal screen to be correctly displayed. The appearance settings may vary on different SSH tools. The default port number is 22.

Character set translation setting:

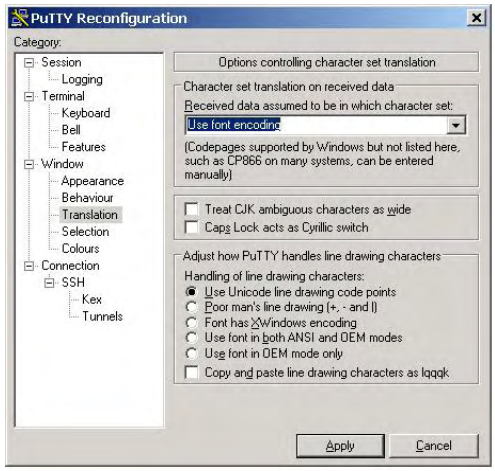

Appearance menu:

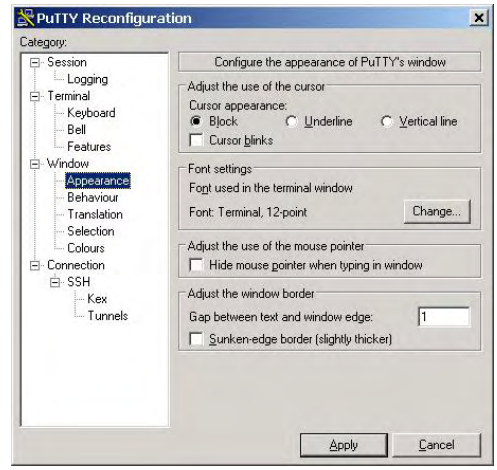

Font type menu:

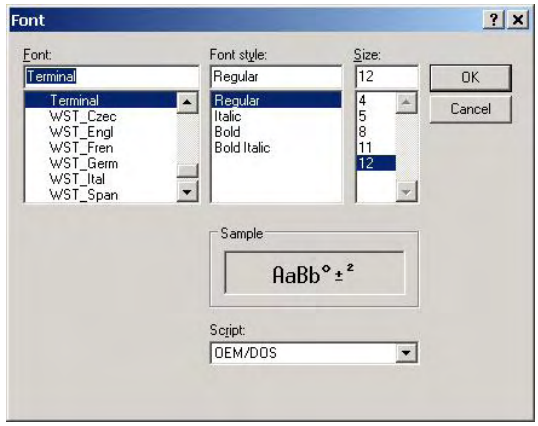

This page is intentionally left blank.

Chapter

2

LCD Screen Messages

2.1 The Initial Screen

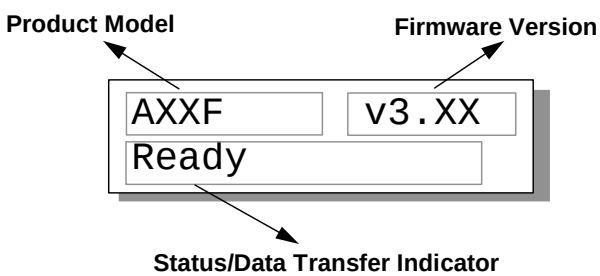

Figure 2-1: LCD Initial Screen

Status/Data Transfer Indicator:

Ready	There is at least one logical drive or logical volume mapped to a host ID/LUN combination.
No Host LUN	No logical drive created or the logical drive has not yet been mapped to any host ID/LUN.
▮▮▮▮▮▮▮	Indicates the percentage of internal processing resources being consumed, not the host bus throughput. Each block indicates megabytes of data that is currently being processed.

IMPORTANT!

For LCD menu structure, please refer to the LCD navigation map which can be acquired through the eSupport or VIProom document distribution areas.

2.2 Quick Installation Screen

```
Quick  Logical
Drive  Install ↓
```

Figure 2-2: Quick Installation Screen

Press **[ENT]** to create a logical drive. The subsystem/controller will start initialization of one logical drive to include all connected disk drives and automatically map the logical drive to the first ID and LUN number of the first host channel. The "Quick Installation" can only be performed when no logical drive has been configured.

IMPORTANT!

- In a large enclosure with many drive bays or a configuration that spans across multiple enclosures, including all disk drives into a logical drive **may not be** a good idea. A logical drive with too many members may cause difficulties with maintenance, e.g., rebuilding a failed drive will take a long time.

- RAID arrays deliver a high I/O rate by having all disk drives spinning and returning I/O requests simultaneously. If the combined performance of a large array exceeds the maximum transfer rate of a host channel, you will not be able to enjoy the performance gain by simultaneous disk access.

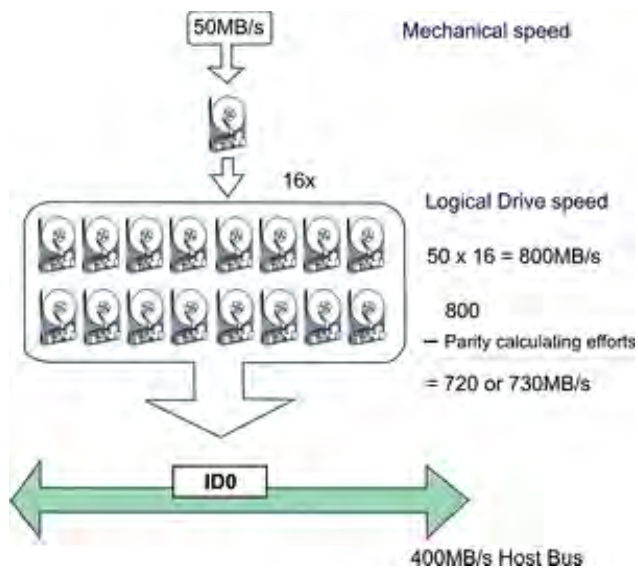

The diagram shows a logical drive consisting of 16 members

and associated with a host ID as the result of using the **Quick Installation** function in a 16-bay enclosure. Although host applications may not always achieve the theoretical numbers shown here, the host bus bandwidth apparently becomes a bottleneck, and the benefit of simultaneous disk access will be seriously reduced.

2.3 Logical Drive Status

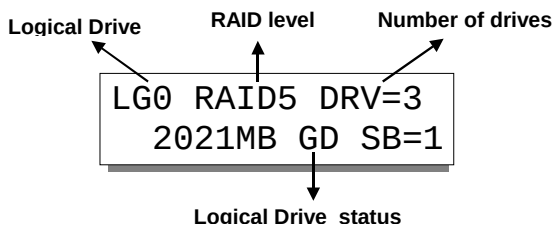

Figure 2-3: Logical Drive Status Screen

LG#:	The Logical Drive index number.
RAID#:	The RAID level applied for this logical drive.
DRV:	The number of physical drives included in this configuration.
Logical Drive Status:	
XxxxMB	The capacity of this logical drive.
SB=x	Standby drives available for this logical drive (including Local, Global, and Enclosure Spares). Except the Local spares specifically assigned to other logical configurations, all available spare drive(s) will be counted in this field, including Global and Enclosure-specific Spares.
xxxxMB INITING	The logical drive is now initializing.

xxxxMB INVALID	Fatal failure or incomplete array means that the LD has lost the protection by RAID configuration. If system cannot find some member disks for a specific LD at boot time, the LD will be considered as incomplete. If some member disks of a specific LD fail during operation, the LD will be considered as fatally failed.
xxxxMB GD SB=x	The logical drive is in good condition.
xxxxMB FL SB=x	One member drive failed in this logical drive.
xxxxMB RB SB=x	Logical drive is rebuilding.
xxxxMB DRVMISS	One of the member drives is missing.
INCOMPLETE ARRAY	One or more drives failed in this logical drive.
FATAL FAIL	Two or more member drives failed at the same time, the array is inaccessible
DRV MISS	A member drive is missing, could result from insecure installation
OFF LINE	A logical drive has fatally failed or manually shutdown. This state can result from other faults such as CRC error checksum

2.4 Logical Volume Status

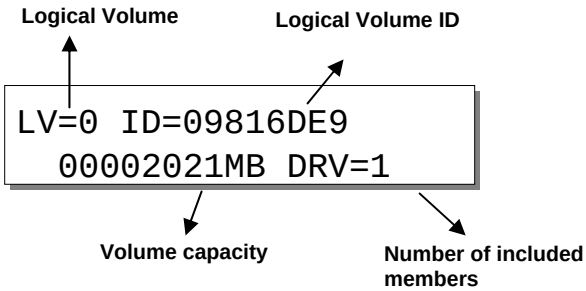

Figure 2-4: Logical Volume Status Screen

Logical Volume:	The Logical Volume number.
DRV=x:	The number of logical drive(s) contained in this logical volume.
Logical Volume ID:	This unique ID is randomly generated by the firmware. In EonPath applications, this ID can be used to identify a RAID volume accessed through two separate host links. Logical drives also have a similar unique ID for ease of identification across a storage network.
Logical Volume Status:	
xxxMB	The capacity of this logical volume.

2.5 Physical Drive Status

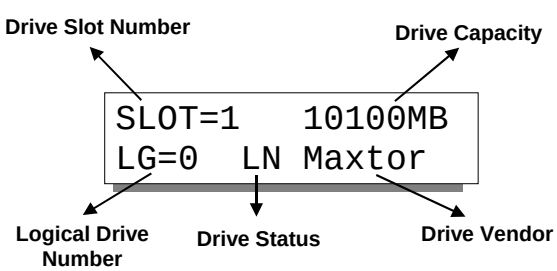

Figure 2-5: Drive Status Screen

Drive Status:

SLOT	The location of this disk drive
LG=*	This drive is a member of logical drive *
LG=x IN	Initializing
LG=x LN	On-line (already a member of a logical configuration)
LG=x RB	Rebuilding
LG=x SB	Local Spare drive

ABSENT	The disk drive does not exist
ADDING	The drive is about to be included in a logical drive through the ADD-Drive procedure
CEDING	When migrating from RAID6 to RAID5, the drive is about to be dismissed from a logical drive. Drive status will be indicated as a formatted drive
COPYING	The drive is copying data from a member drive it is about to replace
CLONE	The drive is a clone drive holding the replica of data from a source drive
CLONING	The drive is cloning data from a source drive

2.6 Channel Status

Host Channel

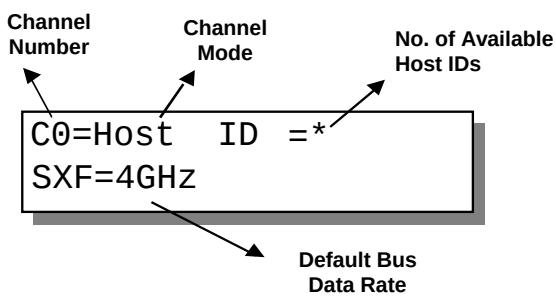

Figure 2-6: Host Channel Status Screen

Drive Channel (Single Controller Configuration)

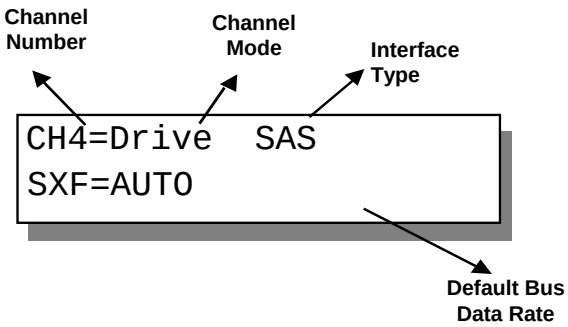

Figure 2-7: Drive Channel Status Screen

Channel Mode:

Host	Host channel mode
Drive	Drive channel mode
RCC	Dedicated inter-controller communication channel

Default Channel Bus Sync Clock:

AUTO	The default setting is set to the auto-negotiate mode
1GHz, 1.5GHz, 2GHz, 3GHz, 4GHz, …	Manually configured channel speed

Available Host IDs:

*	Multiple IDs on the channel (Host channel mode only)
(ID number)	IDs are defined as AIDs (Slot A controller IDs) or BIDs (Slot B controller IDs). Slot A is the default location for the Primary RAID controller. Host Channel: AIDs or BIDs facilitate the distribution of system workload between RAID controllers that reside in enclosure Slot A and Slot B. An AID and a BID can be associated with the same RAID volume. Drive Channel: A drive channel within a dual-controller configuration will carry both an AID and a BID that are preserved for the channel chip processors on Slot A and Slot B controllers.
NA	No ID applied

IMPORTANT!

For a single controller configuration, no IDs will be shown for a drive channel status screen. For a dual-controller configuration, drive channels come with preset IDs and both IDs will be displayed.

2.7 Controller Voltage and Temperature

Press **ENT** for two seconds to enter the Main Menu. Press the up or down arrow keys to select "**View and Edit Peripheral Dev**," then press **ENT**.

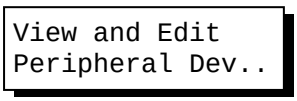

Press the up or down arrow keys to select "**Ctlr Peripheral Device Config.**." Press **ENT,** choose "**View Ctlr Periph Device Status..**", then press **ENT**.

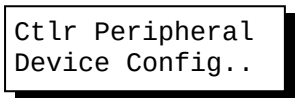

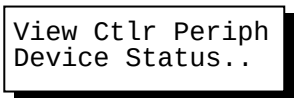

Press the up or down arrow keys to choose either "Voltage Monitor" or "**Temperature Monitor.**"

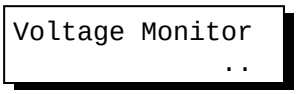

Select "**Temperature and Voltage Monitor**" by pressing **ENT**. Press the up or down arrow keys to browse through the various voltage and temperature statuses.

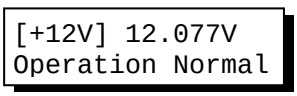

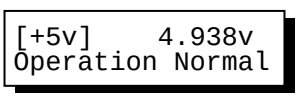

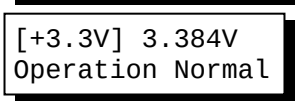

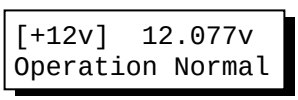

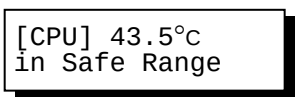

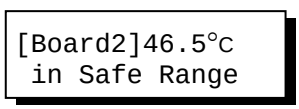

2.8 View and Edit Event Logs

Press **ENT** for two seconds to enter the Main Menu. Press the up or down arrow keys to select "**View and Edit Event Logs**," then press **ENT**.

```
View and Edit
Event Logs      ↑
```

Press the up or down arrow keys to browse through the existing event log items.

```
UPS Power
Failure Detected
```

To see more details about a specific event, use your arrow keys to move to an event, press **ENT** for 2 seconds to display the first page of event details, then use the arrow keys to move to the next page. When finished reading an event, press the **ESC** key to return to the event index.

For the limited space on the LCD screen, details of a system event will be displayed in several pages.

To delete a specified item and all events prior to this event, press the **ENT** key lightly to display the "delete event" confirm message, and then press **ENT** for 2 seconds to clear the events.

```
Press  2 Seconds
to Clear Events
```

IMPORTANT!

- The event log will be cleared after the controller is powered off or reset.
- For models running firmware revision 3.48 or later, events will be written to the drive reserved space and resetting the subsystem will not erase the previous event messages.

Chapter

3

Terminal Screen Messages

3.1 The Initial Screen

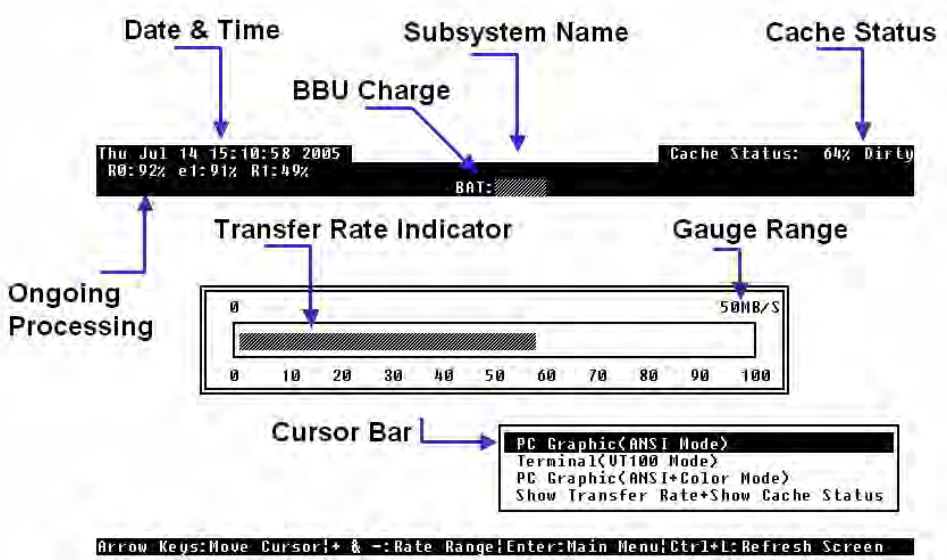

Cursor Bar: Highlights the current selection. Move the cursor bar to a desired item, then press **[ENTER]** to select

Subsystem Name: Identifies the type of controller/subsystem or a preset name

Transfer Rate Indicator: Indicates the current data transfer rate

Gauge Range: Move your cursor bar to "**Show Transfer Rate+Show Cache Status**." Press **[ENTER]** on it to activate the control options, and then use the "**Shift**" and "**+**" or "**-**" key combinations to change the gauge range in order to view the transfer rate indicator. The I/O transfer rate will be indicated in percentage against the gauge range.

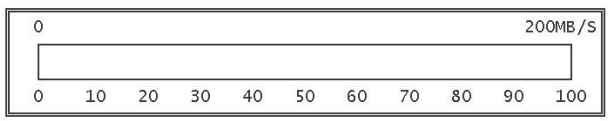

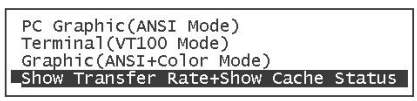

Cache Status: Indicates current cache status

Write Policy: Indicates current write-caching policy

Date & Time: Current system date and time, generated by controller real-time clock

PC Graphic (ANSI Mode): Enters the Main Menu and operates in ANSI mode

Terminal (VT-100 Mode): Enters the Main Menu and operates in VT-100 mode

PC Graphic (ANSI+Color Mode): Enters the Main Menu and operates in ANSI color mode

Show Transfer Rate+Show Cache Status: Press **[ENTER]** on this item to show the cache status and transfer rate

Ongoing Processing:

e#: logical drive # is being expanded

i#: logical drive # is being initialized

R#: logical drive # is being rebuilt

P#: logical drive # Parity Regeneration completion ratio

S#: logical drive # Media Scan completion ratio

For more details, please refer to *the* ***Logical Drive Status*** section in the following discussion.

3.2 Main Menu

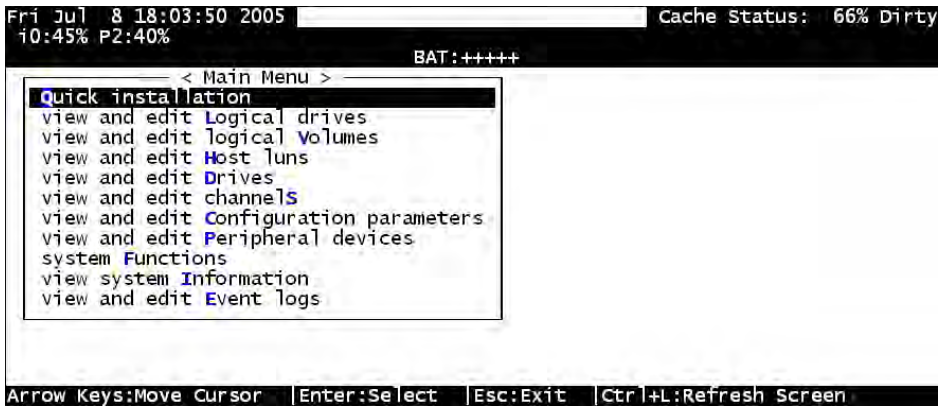

Use the arrow keys to move the cursor bar through the menu items, then press **[ENTER]** to choose a menu, or **[ESC]** to return to the previous menu/screen.

In a subsystem or controller head where battery status can be detected, battery status will be displayed at the top center. Status will be stated as Good, Bad, several "+" (plus) signs (VT-100 mode), or color blocks (ANSI mode) will be used to indicate battery charge. A battery fully-charged status will be indicated by five plus signs (+++++) or color blocks.

When initializing or scanning an array, the controller displays progress percentage on the upper left corner of the configuration screen. An "i" indicates array initialization. An "s" stands for scanning process. The number(s) next to them indicate the logical drive number (e.g., logical drive 0).

3.3 Quick Installation

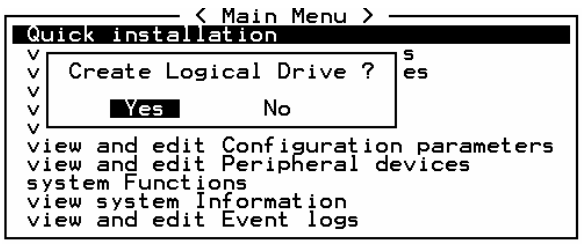

Type **Q** or use the ↑ ↓ keys to select "Quick installation", then press **[ENTER]**. Choose **Yes** to create a logical drive.

All possible RAID levels will be displayed. Use the ↑ ↓ keys to select a RAID level, then press **[ENTER]**. The assigned spare drive will be a Local Spare Drive, not a Global Spare Drive.

All available disk drives in the enclosure will be included in one logical drive. The subsystem will start initialization and automatically map the logical drive to LUN 0 of the first host channel available, usually Channel 0.

The "Quick Installation" can only be performed when no logical drive has been previously configured.

IMPORTANT!

- In a large enclosure with many drive bays or a configuration that spans across multiple enclosures, including all disk drives into a logical drive may not be a good idea. A logical drive with too many members may cause difficulties with maintenance, e.g., rebuild will take a longer time.

- RAID arrays deliver a high I/O rate by having all disk drives spinning and returning I/O requests simultaneously. If the combined performance of a large array exceeds the maximum transfer rate of a host channel, you will not be able to enjoy the performance gain by simultaneous disk access.

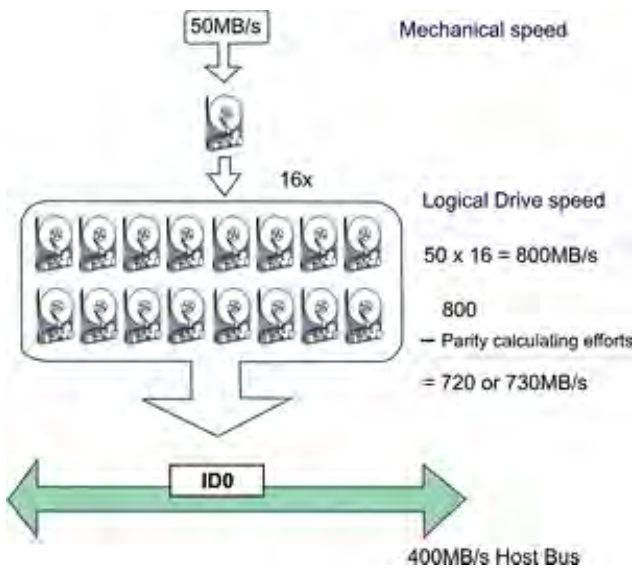

The diagram below shows a logical drive consisting of 16 members and associated with a host ID as the result of using the **Quick Installation** function in a 16-bay enclosure. Although host applications may not always realize the theoretical numbers shown here, the host bus bandwidth apparently becomes a bottleneck, and the benefit of simultaneous disk access will be seriously reduced.

3.4 Logical Drive Status

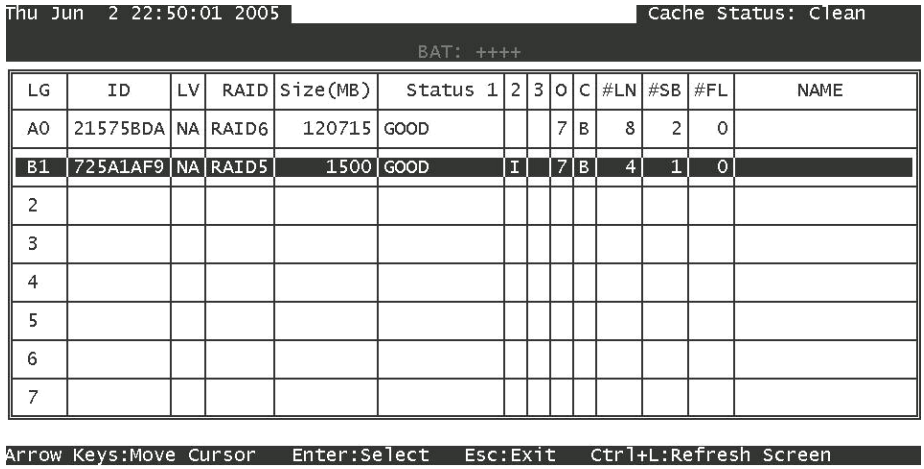

NOTE:

A logical drive in a single-controller subsystem is always managed by one controller, and the "A" or "B" indicator will not appear.

LG Logical Drive number

- A Managed by Slot A controller
- B Managed by Slot B controller

LV The Logical Volume to which this logical drive belongs

ID Firmware-generated unique array ID

RAID RAID level

SIZE (MB) Capacity of the Logical Drive

Status 1 Logical Drive Status – Column 1

- **GOOD** The logical drive is in good condition
- **DRV FAILED** A drive member failed in the logical drive
- **DRV INITING** Logical drive is being initialized
- **INCOMPLETE** One of the causes of the Incomplete state can be one or more member drives are missing or failed in the logical drive

	INVALID	The logical drive was created but has not been fully initialized when another version of firmware is being loaded. After the subsystem resets, the array status should return to normal. Fatal failure or incomplete array means that the LD has lost the protection by RAID configuration. If system cannot find some member disks for a specific LD at boot time, the LD will be considered as incomplete. If some member disks of a specific LD fail during operation, the LD will be considered as fatally failed.
	FATAL FAIL	Two or more member drives failed at the same time, the array is inaccessible
	DRV MISS	A member drive is missing; could result from insecure installation
	REBUILDING	The logical drive is being rebuilt
	OFF LINE	A logical drive has fatally failed or manually shut down. This state can result from other faults such as CRC error checksum

Status 2 Logical Drive Status – Column 2

I	Initializing drives
A	Adding drive(s)
E	Expanding logical drive
H	Add drive operation on hold

Status 3 Logical Drive Status – Column 3

R	Rebuilding the logical drive
P	Regenerating array parity

Column O Logical Drive Status – Stripe size

N/A	Default		
4	16KB	**6**	64KB
5	32KB	**7**	128KB
A	1024KB	**8**	256KB
		9	512KB

Column C Logical Drive Status – Write Policy setting

B	Write-back
T	Write-through

#LN Total number of drive members in the logical drive

#SB	Standby drives available for the logical drive. This includes all the spare drives (local spare, global spare) available for the specific logical drive
#FL	Number of Failed member(s) in the logical drive
Name	Logical drive name (user configurable)

3.5 Logical Volume Status

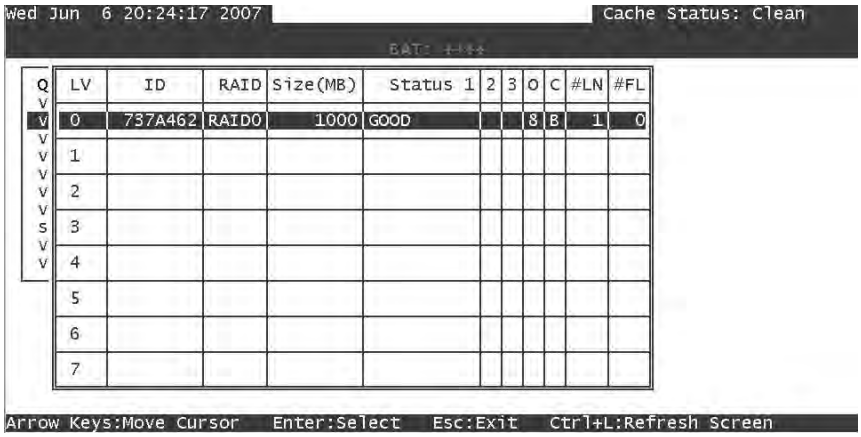

NOTE:

A logical volume in a single-controller subsystem is always managed by one controller, and the "A" or "B" indicator will not appear.

LV	Logical Volume number.
ID	Logical Volume ID number (randomly generated by firmware)
RAID	RAID0 means the members of the logical volume are striped together.
Size(MB)	Capacity of the Logical Volume
#LN	The number of Logical Drive(s) included in this Logical Volume
#FL	The number of failed member(s) within the logical volume.

* For other statuses, please refer to the logical drive information on the previous page.

3.6 Drive Status

SATA Drives

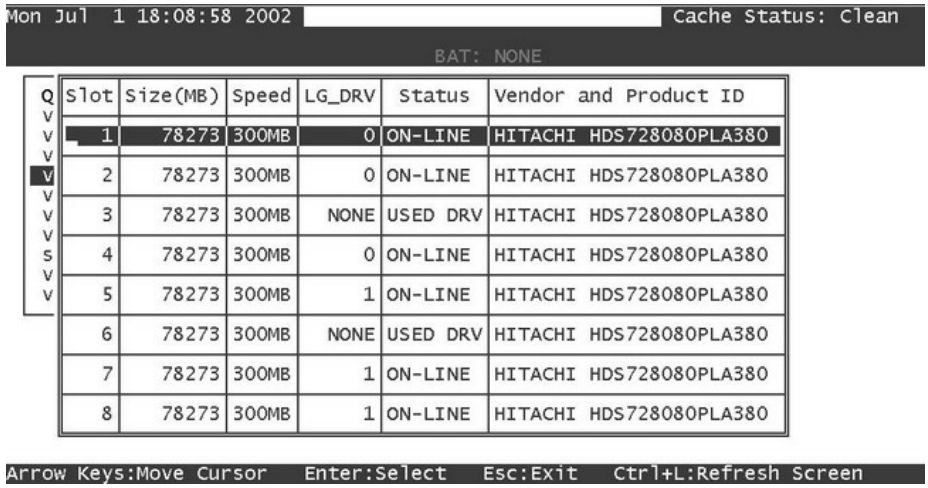

Slot	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
1	78273	300MB	0	ON-LINE	HITACHI HDS728080PLA380
2	78273	300MB	0	ON-LINE	HITACHI HDS728080PLA380
3	78273	300MB	NONE	USED DRV	HITACHI HDS728080PLA380
4	78273	300MB	0	ON-LINE	HITACHI HDS728080PLA380
5	78273	300MB	1	ON-LINE	HITACHI HDS728080PLA380
6	78273	300MB	NONE	USED DRV	HITACHI HDS728080PLA380
7	78273	300MB	1	ON-LINE	HITACHI HDS728080PLA380
8	78273	300MB	1	ON-LINE	HITACHI HDS728080PLA380

SAS Drives

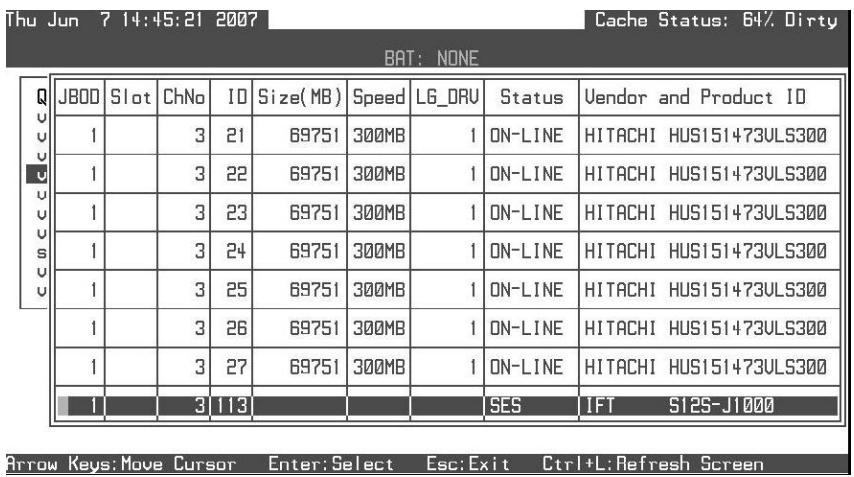

JBOD	Slot	ChNo	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
1		3	21	69751	300MB	1	ON-LINE	HITACHI HUS151473VLS300
1		3	22	69751	300MB	1	ON-LINE	HITACHI HUS151473VLS300
1		3	23	69751	300MB	1	ON-LINE	HITACHI HUS151473VLS300
1		3	24	69751	300MB	1	ON-LINE	HITACHI HUS151473VLS300
1		3	25	69751	300MB	1	ON-LINE	HITACHI HUS151473VLS300
1		3	26	69751	300MB	1	ON-LINE	HITACHI HUS151473VLS300
1		3	27	69751	300MB	1	ON-LINE	HITACHI HUS151473VLS300
1		3	113				SES	IFT S12S-J1000

Slot Drive slot in which a disk drive resides

Size (MB) Drive capacity

Speed **XXMB** Maximum transfer rate of the drive channel interface

ChNo Channel number. For drives within a SAS expansion enclosure, ChNo will be displayed as n1<n2>, showing two SAS domains connecting to a dual-ported SAS drive.

LG_DRV **X** The disk drive is a member of logical drive "X."

If the Status column shows "STAND-BY", the drive is a Local Spare belonging to logical drive "X."

ID A logical device ID assigned to the SAS drive.

Status	**Global**	The disk drive is a Global Spare Drive
	INITING	Proceeding with array initialization
	ON-LINE	The drive is in good condition
	REBUILD	Proceeding with array Rebuild process
	STAND-BY	Local Spare Drive or Global Spare Drive. The Local Spare Drive's LG_DRV column will show the logical drive number. The Global Spare Drive's LG_DRV column will show "Global".
	NEW DRV	A new drive has not been included in any logical drive or configured as a spare drive
	USED DRV	An used drive that is not a member of any logical drive or configured as a spare
	FRMT DRV	Formatted drive (drive formatted with a reserved section)
	BAD	Failed drive
	ABSENT	The disk drive does not exist
	ADDING	The drive is about to be included in a logical drive through the ADD-Drive procedure
	CEDING	When migrating from RAID6 to RAID5, the drive is about to be dismissed from a logical drive. Drive status will be indicated as a formatted drive
	COPYING	The drive is copying data from a member drive it is about to replace
	CLONE	The drive is a clone drive holding the replica of data from a source drive
	CLONING	The drive is cloning data from a source drive
	MISSING	Drive missing (a member drive was once here). This status is shown after boot-up and before I/Os are distributed to the hard drive or accessed by firmware. A missing drive may be corrected by re-inserting the improperly-installed drive tray, etc. If I/Os are distributed and this drive fails to respond, the status will become "failed."
	SB-MISS	Spare drive missing
Vendor and Product ID	The vendor and product model information of the disk drive	

JBOD For disk drives in the expansion enclosures, the number shown in the JBOD column indicates which enclosure the disk drives come from. The JBOD ID is configured via DIP switches or a rotary ID switch on enclosure.

JBOD	Slot	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
1	1	78273	300MB	0	ON-LINE	HITACHI HDS728080PLA380
1	2	78273	300MB	0	ON-LINE	HITACHI HDS728080PLA380
1	3	78273	300MB	NONE	USED DRV	HITACHI HDS728080PLA380
1	4	78273	300MB	0	ON-LINE	HITACHI HDS728080PLA380
3	5	78273	300MB	1	ON-LINE	HITACHI HDS728080PLA380
3	6	78273	300MB	NONE	USED DRV	HITACHI HDS728080PLA380
3	7	78273	300MB	1	ON-LINE	HITACHI HDS728080PLA380
3	8	78273	300MB	1	ON-LINE	HITACHI HDS728080PLA380

3.7 Channel Status

Fibre-to-SATA Configuration

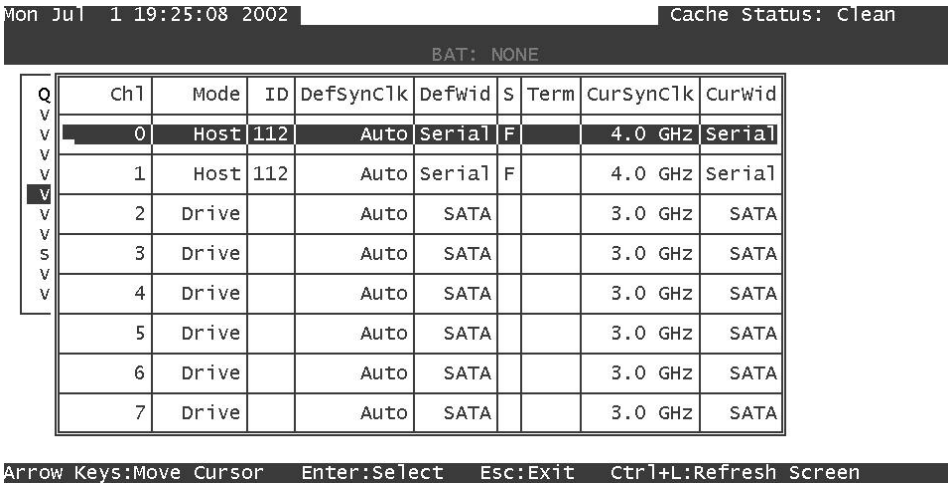

Fibre-to-SAS Configuration

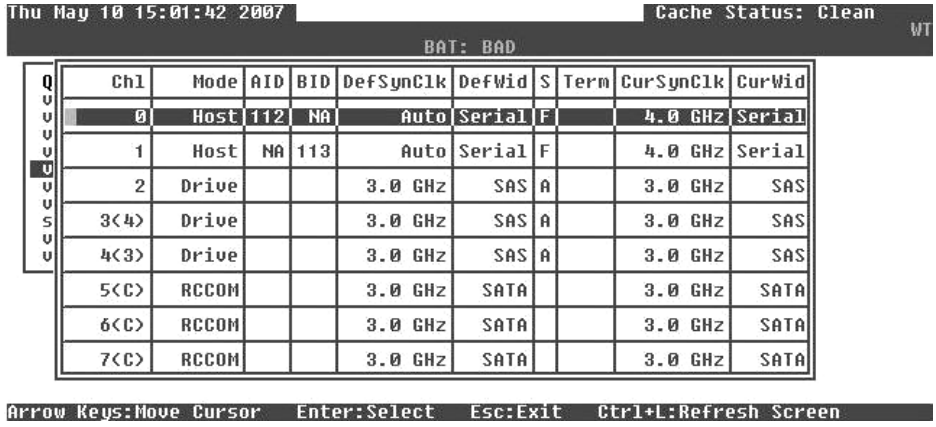

Chl Channel number; expansion links are also defined as drive channels yet with a bracketed number showing the counterpart SAS domain (in a dual-controller configuration).

Mode Channel mode

RCCOM	Redundant controller communication channel
Host	Host Channel mode
Drive	Drive Channel mode

AID IDs managed by the Slot A controller

*	Multiple IDs were applied (Host Channel mode only)
(ID number)	Host Channel: Specific IDs managed by the Slot A controller for host LUN mapping Drive Channel: Specific ID reserved for the channel processor on the Slot

		A controller
BID	IDs managed by the Slot B controller	
	*	Multiple IDs were applied (Host Channel mode only)
	(ID number)	Host Channel: Specific IDs managed by the Slot B controller for host LUN mapping Drive Channel: Specific ID reserved for the channel processor on the Slot B controller; used in redundant controller mode
	NA	No channel ID applied
AUTO	Channel bus data rate set to auto speed negotiation	
DefSynClk	Default bus synchronous clock:	
	*??.?*M	The default setting of the channel is ??.? MHz in Synchronous mode.
	Async.	The default setting of the channel is Asynchronous mode.
DefWid	Default bus width:	
	Serial	Serial transfer protocol; for Fibre Channel or SAS Channel
S	Signal:	
	F	Fibre
	A	SAS
Term	Terminator Status: (not applied here in Fibre-to-SAS/SATA solutions)	
	On	Terminator is enabled.
	Off	Terminator is disabled.
	Diff	The channel is a Differential channel. The terminator can only be installed/removed physically.
	Empty	Non-SCSI bus
CurSynClk	Current bus synchronous clock:	
	*??.?*GHz	The default setting of the channel bus is ??.? GHz
	Async.	The default setting of the channel bus is Asynchronous mode.
	(empty)	The default bus synchronous clock has changed. Reset the controller for the changes to take effect.
CurWid	Current Bus Width:	
	Serial	Serial transfer protocol; Fibre Channel, SAS Channel, SATA Channel.

3.8 Controller Voltage and Temperature

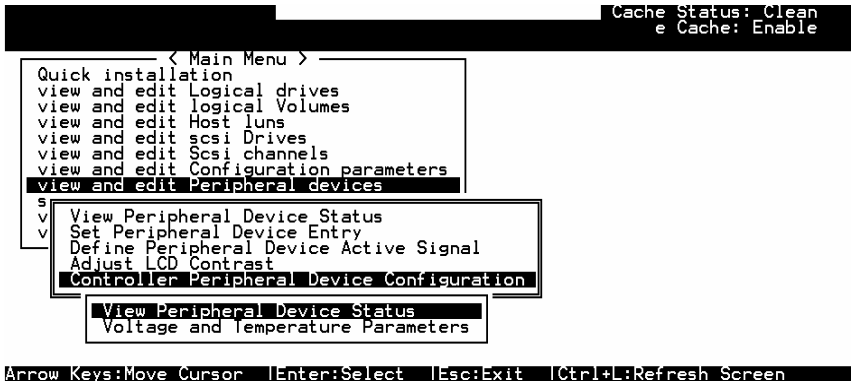

From Main Menu, choose "**View and Edit Peripheral Devices**," and press **[ENTER]**. From the sub-menu, choose "**Controller Peripheral Device Configuration**," "**View Peripheral Device Status**", then press **[ENTER]**.

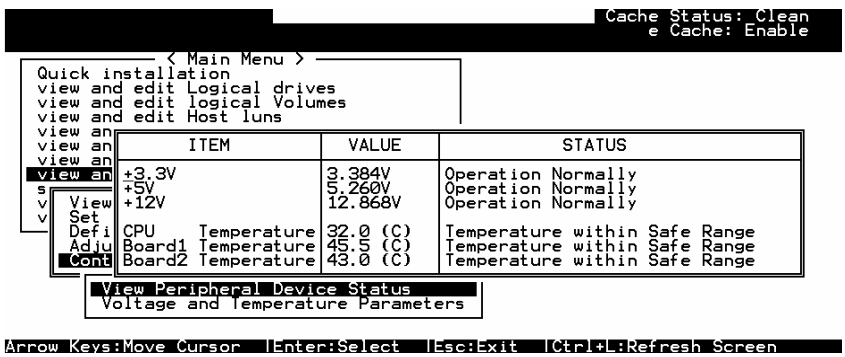

The current status of voltage and temperature detected by the controller will be displayed on-screen and will be stated as normal, out of order, or within the safety range.

3.9 Viewing Event Logs on the Screen

When errors occur, you may want to trace the records to see what has happened to your system. The controller's event log management records all events starting from the time when the system is powered on, recording up to 1,000 events. Powering off or resetting the controller will automatically delete all of the recorded event logs.

With firmware revisions 3.48 and later, the event logs are stored in **disk reserved space**, and hence the event logs are available after system reset. Disk reserved space is automatically created when composing a logical drive. With no logical drives, event logs can not be preserved.

To view the events log on-screen, select "**View and Edit Event Logs**" from the Main Menu by pressing **[ENTER]**.

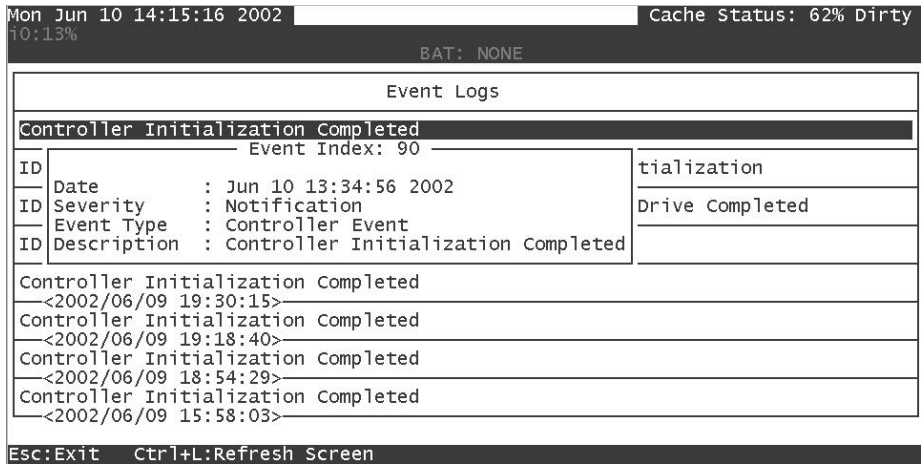

To check for more details about a specific event, move the cursor bar to highlight a specific event and press the **[Space]** key to display the complete event information.

To clear the saved event logs, scroll the cursor down to select an event and press **[ENTER]** to delete the event and the events below**.**

Choose **Yes** to clear the recorded event logs.

This page is intentionally left blank.

Chapter

4

Optimization & Preparation Tasks

There are preference parameters that cannot be easily altered after the creation of logical arrays. Reconfiguration takes time and inappropriate configurations prevent you from getting the best performance from your Infortrend arrays. It is therefore highly recommended to thoroughly consider preferences such as stripe sizes, caching parameters, etc. before creating your logical arrays.

4.1 Caching Parameters

4.1.1 Stripe Size

Each RAID level has a preset value for the array stripe size. If you prefer a different stripe size for a RAID array (a logical drive), you must backup or move the stored data elsewhere and re-create the array.

Listed below are the default stripe sizes implemented with different RAID levels. These values should be adequate for optimal performance with most applications.

Table 4-1: RAID Levels and Stripe Sizes

	Stripe Size:
RAID0	128KB
RAID1	128KB
RAID3	16KB

RAID5	128KB
RAID6	128KB
NRAID	128KB

Stripe sizes different from the above defaults can be manually applied to individual logical drives during the initial configuration stage to match the access sizes conducted by your host applications.

NOTE:

The **Stripe size** here refers to the "**Inner Stripe Size**" specifying the **chunk size** allocated on each individual data drive for parallel access instead of the "**Outer Stripe Size**" which is the sum of chunks on all data drives.

Configuration Process

Although stripe size can be adjusted on a per logical drive basis, users are not encouraged to make a change to the default values.

Smaller stripe sizes are ideal for I/Os that are transaction-based and randomly accessed. However, using the wrong stripe size can cause problems. For example, when an array set at 16KB stripe size receives files of 128KB size, each drive will have to spin and write many more times to conduct small fragment 16KB writes to hard disks.

4.1.2 Write-Back Cache Enable/Disable

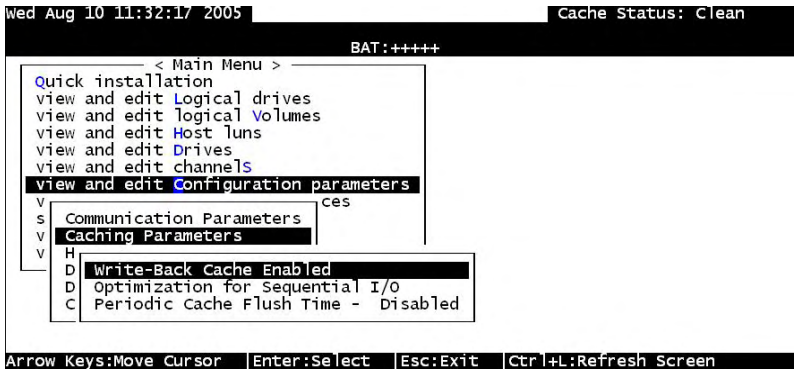

As one of the sub-menus in "**Caching Parameters**," this option controls the cached write policy.

Select "**Caching Parameters**" by moving the cursor bar, then press **[ENTER].** Select "**Write-Back Cache**," then press **[ENTER].** When "Write-back" is "Enabled," the write requests from the host will be held in cache memory and distributed to disk drives later. When "Write-back" is "Disabled" (i.e., the Write-through is adopted,) host writes will be directly distributed to individual disk drives. Select **Yes** in the dialog box that follows to confirm the setting.

- The **Write-through** mode is safer if your controller is not configured in a redundant pair and there is no battery backup or UPS device to protect cached data.

- **Write-back** caching can dramatically improve write performance by caching the unfinished writes in memory and letting them be committed to drives in a more efficient manner. In the event of power failure, a battery backup module can hold cached data for days.

- **Configuration Variables**

  The Write-back options can be found either here in the **Configuration Parameters** menu or in the "**View and Edit Logical Drive**" sub-menu (logical drive or logical volume). One found here is the system general setting; the other applies to a specific logical array.

LCD Keypad

Step 1. From the Main Menu, select "**View and Edit Config Parms**," "**Caching Parameters**," and press **ENT.**

```
Caching
Parameters     ..
```

Step 2. As one of the sub-menus in "**Caching Parameters**," this option controls the cached write function. Press **ENT** to enable or disable "Write-back Cache."

```
Write-Back Cache
Enabled        ..
```

Step 3. Press **ENT** for two seconds to confirm. The current status will be displayed on the LCD.

```
Disable Write
-Back Cache    ?
```

The Write caching options also appear in array-specific (logical drive and logical volume) configuration menu and should look like the screens shown on the right.

```
Change Logical
Drive Parameter?
```

```
Write Policy
Default        ..
```

```
Write Policy
Default        ?
```

Terminal

System General Setting

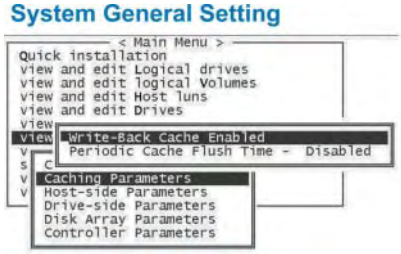

Array-specific Setting

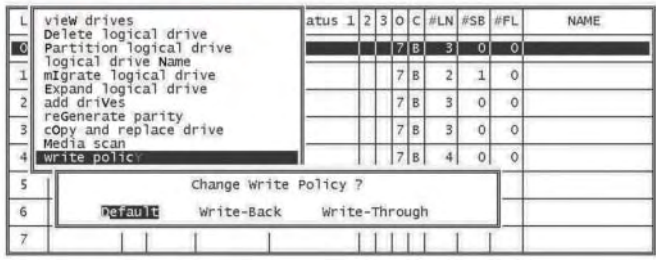

The configuration options are related to the **Event Triggered Operation** feature.

1. The **Event Triggered Operation** feature allows the firmware to automatically enable or disable Write-back caching in the event of component failure or critical system alarms.

   As shown below, a relatively unsafe condition will force the controller to assume a conservative "Write-through" caching mode.

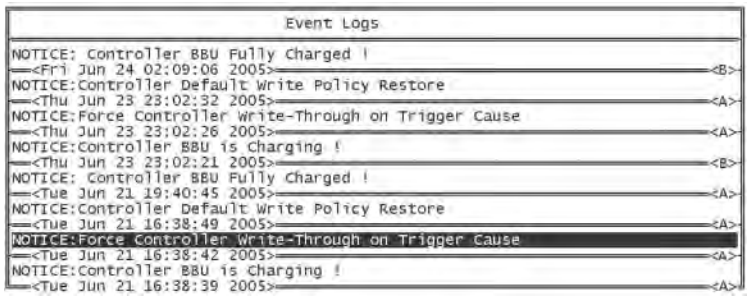

2. A "**Default**" Write-back option is available with individual logical arrays.

   If a logical array's Write-back mode is set to "Default," the caching mode of that particular array will be dynamically controlled by the firmware.

3. If the Write-back mode is manually specified as "**Enabled**" or "**Disabled**" in a particular logical array, then I/Os directed to that array will be handled in accordance with the setting regardless of the system's general setting.

4. The "**Event Trigger**" options can be found in "**View and Edit Peripheral devices**" -> "**Set Peripheral Device Entry**" -> "**Event Trigger Operations**."

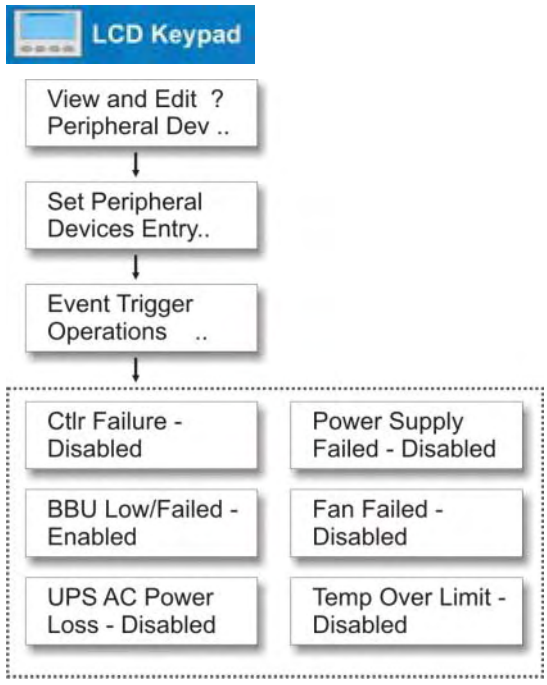

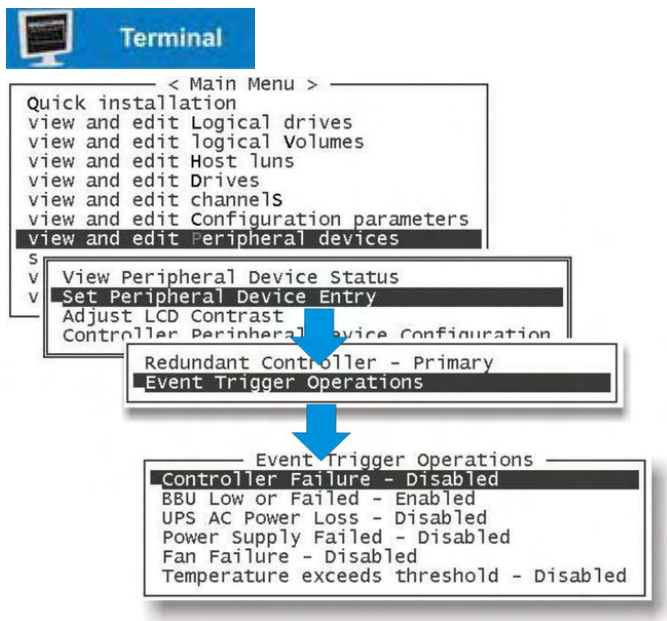

Enable one or more preferred options on the list to protect your array from hardware faults.

4.1.3 Periodic Cache Flush

If Write-back caching is preferred for better performance yet data integrity is also a concern, e.g., a configuration without battery protection or synchronized cache between partner controllers, the system can be configured to flush the cached writes at preset intervals.

LCD Keypad

Step 1. From the Main Menu, select "**View and Edit Config Parms**," "**Caching Parameters**," and press **ENT.**

```
Caching
Parameters   ..
```

Step 2. Use the arrow keys to scroll through the options and select "**Periodic CachFlush Time**", and then press **ENT** to proceed.

```
Period CachFlush
Time - Disable
```

Step 3. The "**Set Cache Flush Time – Disable**" appears. The default is "Disable." Use your arrow keys to select an option from "**ConSync**," "**30sec**," to "**600** sec." "**ConSync**" stands for "continuously synchronized."

```
Set Cache Flush
Time - Disable
```

Step 4. Press **ENT** to select and press **ESC** to exit and the setting will take effect immediately.

Terminal

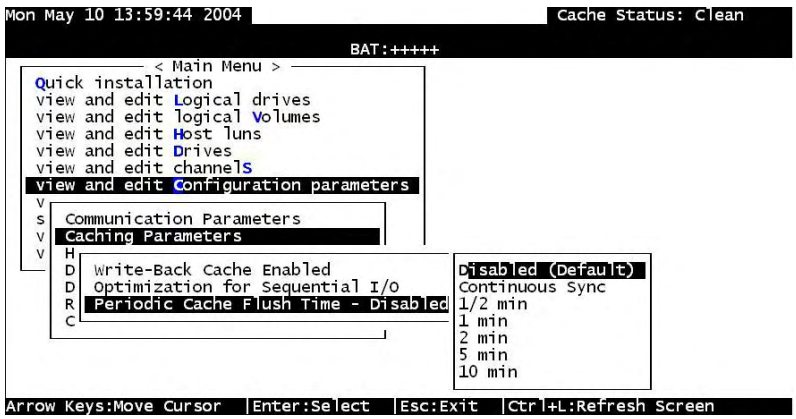

Note that the "**Continuous Sync**" option holds data in cache for as long as necessary to complete a write operation and immediately commits a write request to hard drives if it is not followed by a series of sequential writes.

IMPORTANT!

- Every time you change the Caching Parameters you must reset the controller for the changes to take effect.

4.2 Preparing Channels and Channel IDs

4.2.1 Channel Mode Settings

Subsystems: The EonStor subsystems come with preset data paths and there is no need to modify channel mode configurations.

NOTE:

1. For different channel assignments, please refer to the Hardware manual that came with your subsystem.
2. Technical terms like Slot A, Slot B, RCC (Redundant Controller Communications), and DRVRCC will only appear in a dual-controller configuration.
3. The latest Infortrend **ASIC400** EonStor models come with dedicated RCC (Redundant Controller Communications) chipsets that provide communication paths strung between partner RAID controllers. The "**Drive+RCC**" and "**RCC**" options will not appear on the list of available channel modes.
4. You can still find these RCC channels on the channel list, only that there are no configurable options with these dedicated RCC paths.

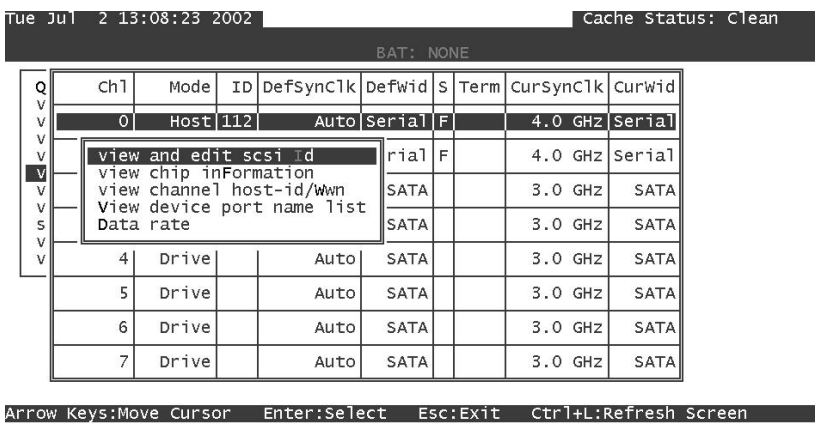

NOTE:

In addition to the 1U FF series controller head, most EonStor RAID subsystems have preset host or drive channels interfaced through a backplane. The channel mode options are not available on these

models.

4.2.2 Channel ID Configuration

Each host channel comes with a default ID, which will not be sufficient if your subsystem comes in a dual-active controller configuration.

In a dual-active controller configuration, you need to manually create **Controller B IDs** to distribute the workload between partner RAID controllers. The idea is diagrammed below:

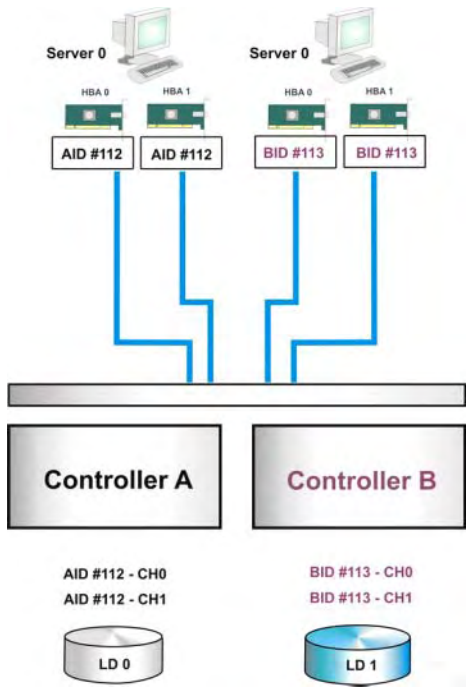

Configuration:

1. 2 logical drives (LD)
2. LD LUN mapping associations
   LD0: CH0 AID112 & CH1 AID112
   LD1: CH0 BID113 & CH1 BID113
3. Controller B IDs need to be manually created.

- Note that in this example a multi-pathing software is required to manage the fault-tolerant links to a RAID storage volume.

Figure 4-1: ID Mapping for Cross-controller Load Sharing

A logical group of physical drives can be associated either with **Controller A** IDs or **Controller B** IDs through the host LUN mapping process. These A or B IDs then appear to the application servers as storage capacity volumes. As a rule of thumb, a logical drive associated with A IDs is managed by Controller A. One that is associated with B IDs is managed by Controller B.

Depending on how many RAID capacity volumes you wish visible to your application servers, create one or more Controller A or Controller B IDs. In firmware menus, these IDs are specified as the **slot A** or **slot B** IDs.

CAUTION!

Firmware revisions 3.48 and later also support the **cross-controller** ID mapping. The cross-controller mapping allows you to associate a logical drive with BOTH controller A and controller B IDs. However, mapping to both controllers' IDs is only beneficial when it is difficult making the fault-tolerant host links between RAID controllers and host HBAs, e.g., using SAS-to-SAS RAID systems. Currently, there is no external SAS switch available on the market. For Fibre-host systems, fault-tolerant links can easily be made with the help of external bypass such as Fibre Channel switches.

For details of fault-tolerant link connections, please refer to your system Hardware Manual.

LCD Keypad

Step 1. Press **ENT** for two seconds to enter the Main Menu. Press the up or down arrow keys to select "**View and Edit Channels**," then press **ENT**.

```
View and Edit
Channels    ..
```

Step 2. Channel information will be displayed. Press **ENT** on the host channel you wish the ID changed.

```
CH0=Host ID=112
SXF=4GHz
```

Step 3. Press the up or down arrow keys to select "**Set Channel ID**," then press **ENT**.

```
Set Channel ID
              ..
```

Step 4. Use the up or down arrow keys to browse through the existing host IDs. Press **ENT** on any ID combination to continue.

```
CHL=0 ID=112
Slot A       ..
```

Terminal

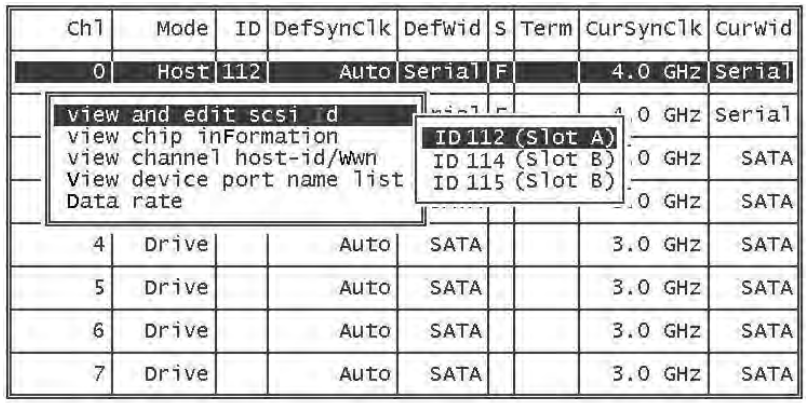

Step 1. Select a host channel, press **[ENTER]** to display the command list.

Step 2. Select "**View and Edit ID**." A list of existing ID(s) will be displayed on the screen. As a default, the subsystem comes with only a **Slot A** controller ID.

Step 3. Select one of the existing IDs and press **[ENTER]**. You may then add a new ID or delete an existing ID.

Adding a Host ID

Press **ENT** on a host channel, on "**Set Channel ID**", and then on an existing ID.

LCD Keypad

Step 1. Use the up and down arrow keys to select "**Set Channel ID**", then press **ENT**.

```
Set Channel
ID           ..
```

Step 2. An existing ID displays.

```
CHL=0 ID=*
Slot A       ..
```

Step 3. Press **ENT** to display "**Add Channel ID**." Press **ENT** again to display the question mark.

```
Add Channel ID
             ..
```

```
Add Channel ID
             ?
```

Step 4. In a dual-controller configuration, once you

```
Controller
A            ?
```

enter the Add ID process, use the up and down arrow keys to select either the Slot A or Slot B controller.

Step 5. An ID next to the existing ID will display on the screen. Use arrow keys to select an ID. When the preferred ID is selected, press **ENT** for two seconds to complete the process.

```
Add CHL=* ID=*
Channel A     ?
```

Step 6. A prompt will remind you to reset the subsystem for the configuration change to take effect. You may press **ENT** to reset the subsystem immediately or you may press **ESC** to continue adding other host IDs and reset the subsystem later.

```
Change Setting
Do Reset Ctlr  ?
```

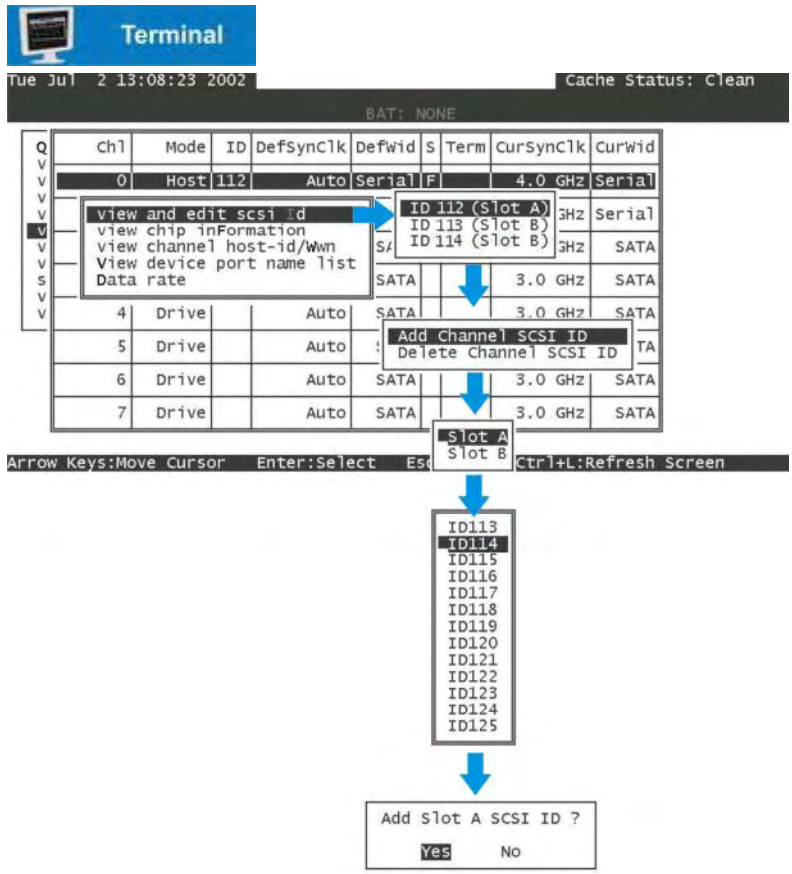

Step 1. Press **[ENTER]** on one of the existing IDs.

Step 2. Select **Add Channel ID.**

Step 3. Specify the host ID either as the **Slot A** or **Slot B** ID. Press **[ENTER]** to proceed.

Step 4. Available IDs will appear in a pull-down list. Select by pressing **[ENTER]** and then select **Yes** to confirm.

Step 5. A confirmation box will prompt to remind you to reset the controller for the configuration to take effect. You may select **Yes** for an immediate reset or **No** to reset later.

```
NOTICE: Change made to this setting will NOT take effect
        until the controller is RESET. Prior to resetting the
        controller, operation may not proceed normally.

          Do you want to reset the controller now ?

                    Yes                    No
```

Deleting an ID

LCD Keypad

Step 1. Press **ENT** for two seconds to enter the Main Menu. Press the up or down arrow keys to select "**View and Edit Channels**," then press **ENT**.

```
View and Edit
Channels      ↕
```

```
CHL=Host ID=113
SXF=AUTO      ..
```

Step 2. The first host channel should appear. Press **ENT** to select a host channel.

Step 3. Press **ENT** on "**Set Channel ID..**"

```
Set Channel
ID            ..
```

Step 4. A list of host channel and host ID combinations will appear. Use the up or down arrow keys to select the ID you wish to remove. Press **ENT** to select a channel ID combination.

```
CHL=0 ID=112
              ..
```

Step 5. You will then be prompted by the "**Add Channel ID**" option. Press the down arrow key to proceed.

```
Add Channel ID
              ..
```

Step 6. The "**Delete Channel ID**" option will appear. Press **ENT** to display the confirmation box. Press **ENT** for two seconds to remove the ID.

```
Delete Channel
ID            ..
```

```
Delete ID=113
              ?
```

Step 7. A prompt will remind you to reset the subsystem for the configuration change to take effect. You may press **ENT** to reset the subsystem immediately or you may press **ESC** to continue adding other host IDs and reset the subsystem later.

```
Change Setting
Do Reset Ctlr  ?
```

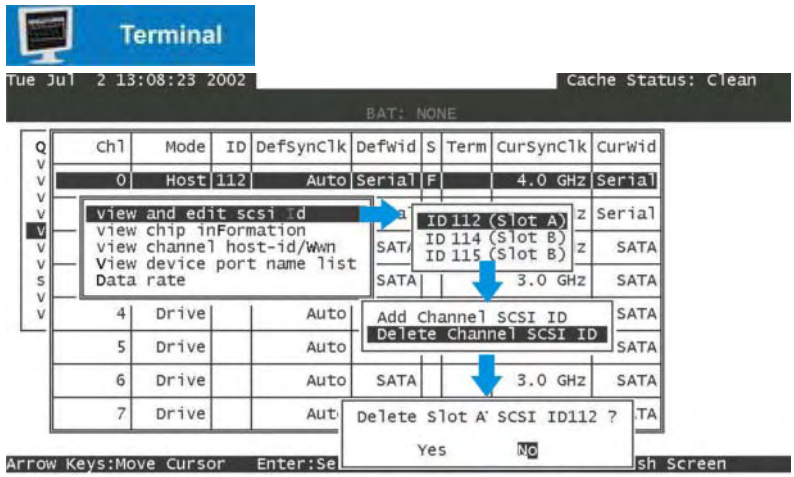

Step 1. Select the host bus ID you wish to delete from the "**View and Edit IDs**" menu.

Step 2. Press **[ENTER]** to display the **Delete Channel ID** option. The dialog box "**Delete ID#?**" will appear.

Step 3. Select **Yes**, then press **[ENTER]** to confirm.

IMPORTANT!

- Every time you change a channel ID, you must reset the subsystem/controller for the changes to take effect.

```
NOTICE: Change made to this setting will NOT take effect
        until the controller is RESET. Prior to resetting the
        controller, operation may not proceed normally.

          Do you want to reset the controller now ?

                    Yes                    No
```

- At least one controller's ID should be present on each channel bus.

4.2.3 Data Rate (Host Channel Bus)

The data rate default is "AUTO" and should work fine with most configurations. In some cases, you may want to install a 4Gbps interface subsystem in a storage network consisting of 2Gbps devices. Please note that mixing 4G and 2G devices in a storage

network may not be supported with all kinds of HBAs or Fibre switches.

Note that the data rate setting on drive channel menus is the maximum transfer rate of the channel bus in that mode. It does not mean a single disk drive can actually carry out that amount of sustained read/write performance.

Host Channel:

Step 1. From Main Menu, select "**View and Edit Channels**," and then the host channel you wish to change.

```
C0=Host AID=112
BID=NA SXF=AUTO
```

Step 2. Press **ENT** on the channel and use the arrow keys to find the "**Data Rate**" option.

```
Data Rate
                ..
```

Step 3. Press **ENT** on the Data Rate option to display "**Set Chl=X Data Rate To AUTO?**", where "**X**" stands for the channel number.

```
Set Chl=X Data
Rate To AUTO ?
```

Step 4. Use your arrow keys to display the desired data rate. Press **ENT** to confirm the selection.

Drive Channel:

Step 1. From Main Menu, select "**View and Edit Channels**," and then the drive channel you wish to change.

```
CHX=Drive AID=7
BID=8 SXF=AUTO
```

```
View Chip
 Information  ..
```

Step 2. Press **ENT** on the channel and use the arrow keys to find the "**Data Rate**" option.

```
Data Rate
                ..
```

Step 3. Press **ENT** on the Data Rate option to display "**Set Chl=X Data Rate To AUTO?**", where "**X**" stands for the channel number.

Step 4. Use your arrow keys to display a data rate value

```
Set Chl=X Data
Rate to AUTO ?
```

which ranges from 33 to 300MBps (SATA drive channels). Press **ENT** to confirm a selection.

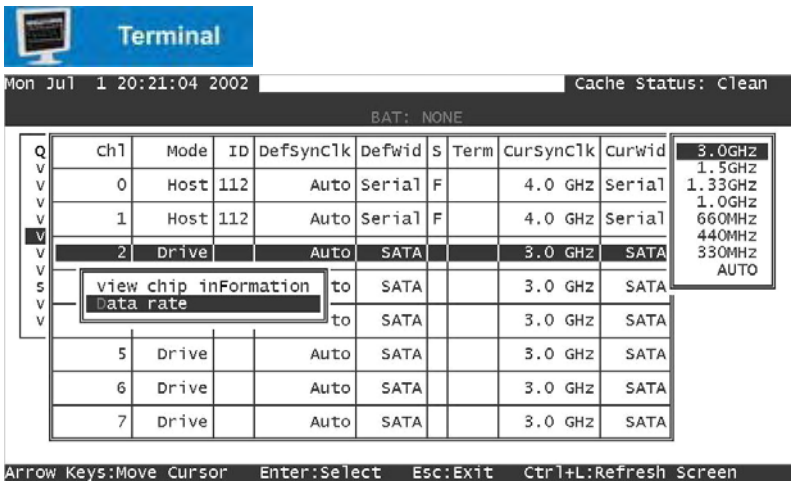

Press **[ENTER]** on a drive channel to display the data rate option. Press **[ENTER]** again to display a pull-down list.

4.3 Controller Date and Time

This sub-menu only appears in subsystem/controller models equipped with a real-time clock. Setting the correct date and time is important especially when tracing system faults or applying automated maintenance utilities such as Media Scan scheduler.

Infortrend's latest RAIDWatch supports time synchronization with SNTP time server and it is recommended to specify your time zone.

Time Zone

The controller uses GMT (Greenwich Mean Time), a 24-hour clock. To change the clock to your local time zone, enter the numbers of hours earlier or later than the Greenwich Mean Time after the plus (+) or minus (-) sign. For example, "+9" is Japan's time zone.

Step 1. Choose "**View and Edit Configuration Parameters**," "**Controller Parameters**," then press **ENT**.

Step 2. Press the up or down arrow keys to scroll down and select "**Set Controller Date and Time**", then press **ENT**.

Step 3. Choose "**Time Zone**" by pressing **ENT**.

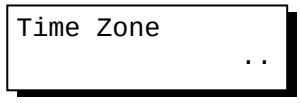

Step 4. Use the **down** key to enter the **plus** sign and the **up** key to enter **numbers**.

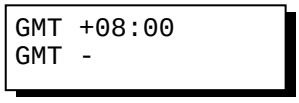

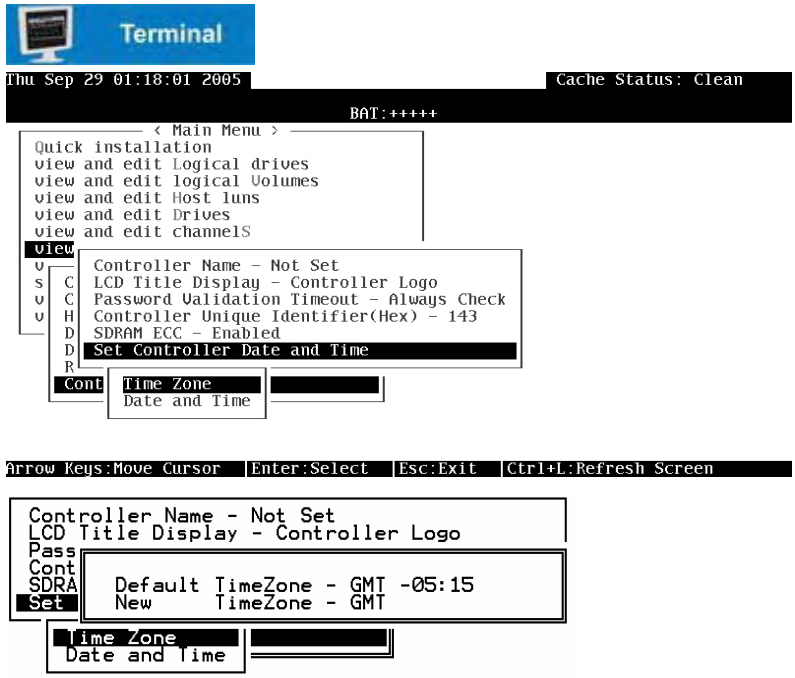

Date and Time

Step 1. Use your arrow keys to scroll down and select "**Date and Time**" by pressing ENT.

```
Date and Time
             ..
```

Step 2. Use the arrow keys to select and enter the numeric representatives in the following order: month, day, hour, minute, and the year. Use up/down arrow keys to change the number displayed on screen, press **ENT** to shift to the next number.

```
[MMDDhhmm[YYYY]]
```

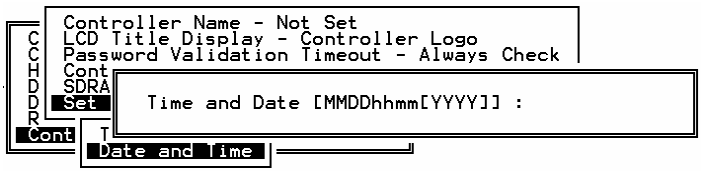

Enter time and date in its numeric representatives in the following order: month, day, hour, minute, and the year.

4.4 Faulty Drive Detection

There is one configurable options directly related to detecting faulty drives and the insertion of a replacement drive. Without enabling this option, a failed drive is only detected when host I/Os are distributed to it or when the firmware scans the drive busses during system initialization. By then controller will return timeout and transaction delays may be experienced.

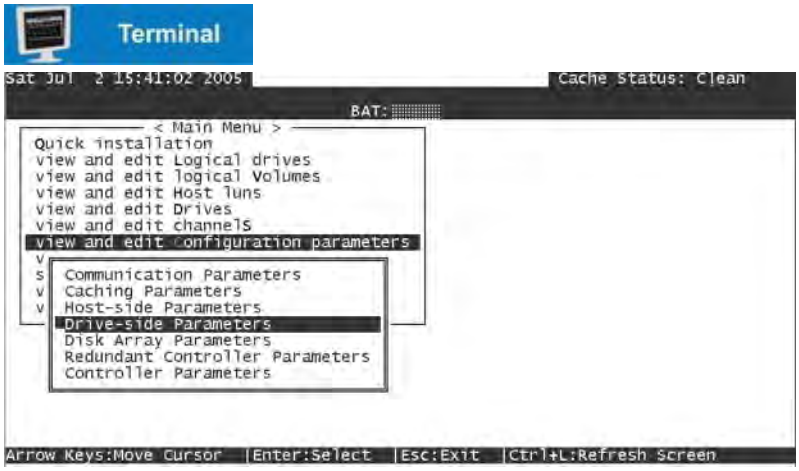

Auto Rebuild on Drive Swap Check Time

When enabled, the **Auto Rebuild** check time scans the drive bus/channel on which a failed drive resides. If the drive swap check detects a replacement drive, the system firmware will automatically proceed with the array rebuild process.

Without the **Auto Rebuild** check time, the rebuild process can be manually initiated through a "**rebuild**" command under the "**View and Edit Logical Drive**" sub-menu. This check time mechanism is specifically applicable in a configuration where no hot-spare is available.

LCD Keypad

Step 1. Select "**View and Edit Config Parms**" from the terminal Main Menu. Enter its sub-menus by pressing **ENT**.

```
View and Edit
Channels      ↕
```

Step 2. Use arrow keys to select "**Drive-side Parameters**." press **ENT** to enter its sub-menus.

```
Drive-side
Parameters    ..
```

Step 3. There are a dozen configurable options under Drive-side parameters. Use arrow keys to select "**Auto Rebuild on Drv Swap**." Press **ENT** on it to change the setting. The options range from Disabled and 5 to 60 seconds.

```
Auto Rebuild on
Drv Swap-15 sec
```

Terminal

The option can be found under "**View and Edit Configuration Parameters**" -> "**Drive-side Parameters**." Press **[ENTER]** on it to display the pull-down menu.

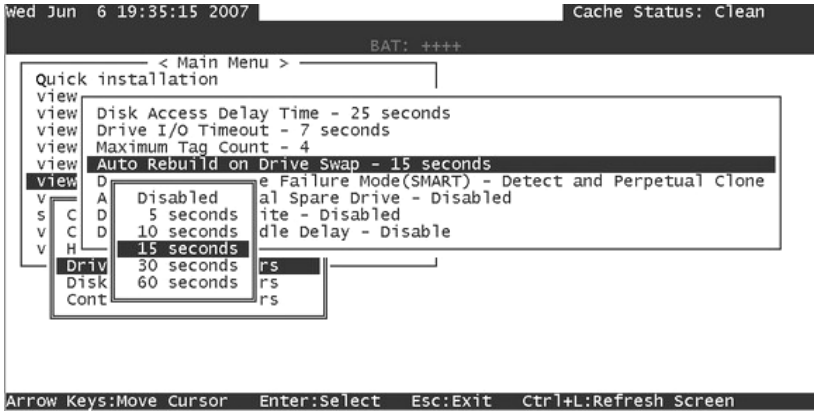

4.5 Spare Drive Implementations

Shown below are two Spare drive policies designed to prevent configuration errors: **Auto-assign Global Spare** and **Enclosure Spare Drive**.

Auto-assign Global Spare

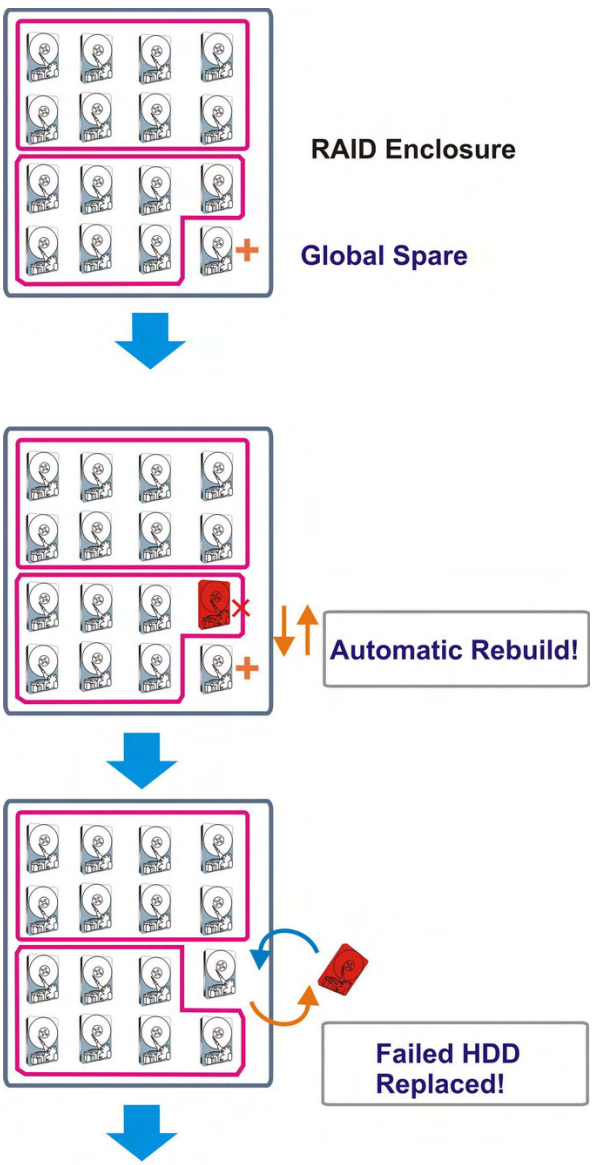

The **Auto-Assign Global Spare** feature is designed to reduce the chance of down time by operator's negligence. Shown on the left is a RAID enclosure with its drives configured into two arrays and a Global Spare. One logical drive consists of 8 members; the other consists of 7.

Diagrams on the left shows how the Auto-assign mechanism helps prevent downtime:

1 A member drive in one of the two logical drives fails. The Global Spare immediately participates in the rebuild.

2 The Failed drive is then replaced by a replacement drive. The original Global Spare becomes a member of the 7-drive array.

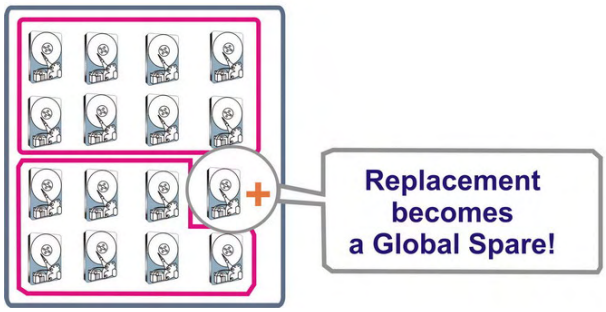

3 With the **Auto-Assign** feature, firmware automatically configures the replacement drive as a **Global Spare**.

The Auto-Assign feature prevents the situation when a failed drive is replaced and the operator **forgets** to configure the replacement drive as another Global Spare leaving the array vulnerable to the occurrence of another drive failure.

Configuration Process:

LCD Keypad

Step 1. Select "**View and Edit Config Parms**" from the terminal Main Menu. Enter its sub-menus by pressing **ENT**.

```
View and Edit
Channels      ↕
```

Step 2. Use arrow keys to select "**Drive-side Parameters**." press **ENT** to enter its sub-menus.

```
Drive-side
Parameters    ..
```

Step 3. There are a dozen of configurable options. Use the arrow keys to select "**Periodic SAF-TE ChkTime -**." Press **ENT** on it to change the setting. The options ranges from Disabled, 50ms,… to 60 seconds.

```
Periodic SAF-TE
ChkTime-Disabled
```

Terminal

The configuration option can be found under "**View and Edit Configuration Parameters**"-> "**Drive-side Parameters**"-> "**Auto-Assign Global Spare Drive**."

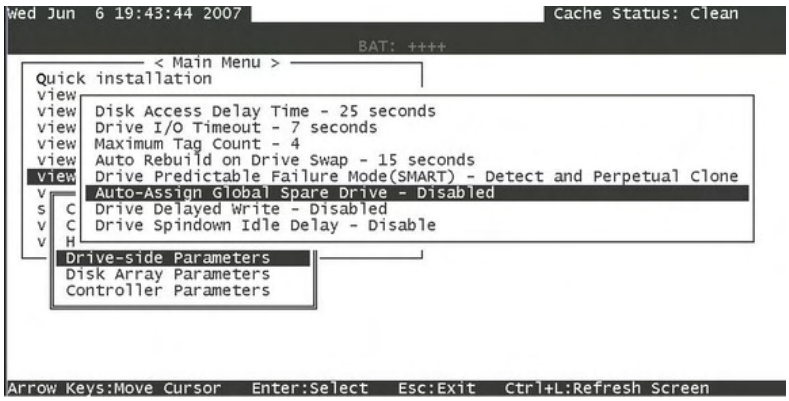

Enclosure Spare

In addition to the traditional "**Local**" and "**Global**" hot spares, another hot-spare type is added in firmware release 3.48. Global hot-spare may cause a problem as diagrammed below in a storage application consisting of multiple enclosures:

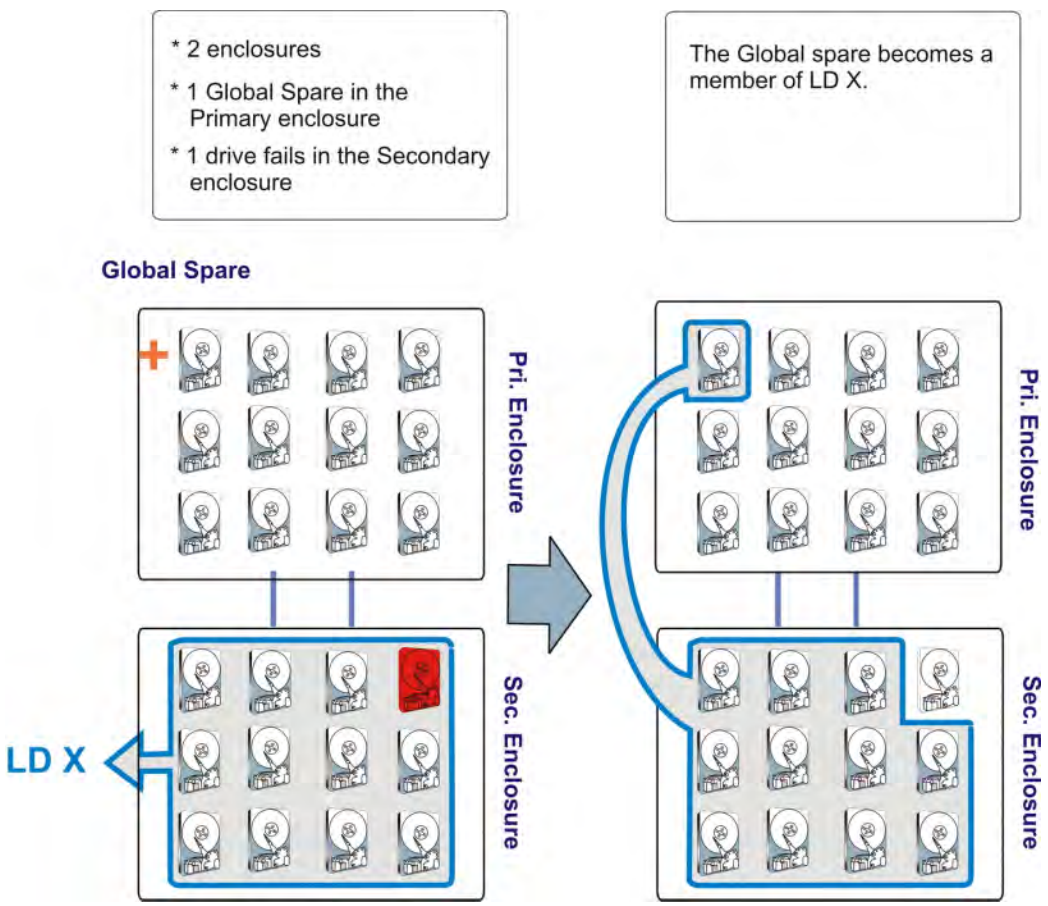

Figure 4-2: Fault Condition: Global Spare and Cross-enclosure Rebuild

A Global spare participates in the rebuild of any failed drive. When a Global spare participates in the rebuild of a logical drive in another enclosure, it will become the member of that logical drive. Although the logical drive can work properly, however, spanning a logical configuration across different enclosures increases the chance of removing the wrong drive, mixing SAS and SATA drives of different RPM's, etc.

The **Enclosure Spare** helps prevent the situation from causing inconvenience. An **Enclosure Spare** only participates in the rebuild of drives that reside in the same enclosure.

Configuration Process:

Step 1. Select "**View and Edit Drives**" from the terminal Main Menu. Enter its sub-menus by pressing **ENT**.

```
View and Edit
Drives       ↕
```

Step 2. Use arrow keys to select a new or formatted drive. Press **ENT** on it to display drive-specific functions.

```
C=* I=* XXXXMB
NEW DRV Vendor
```

Step 3. Use arrow keys to find "**Add Enclosure Spare Drive**." Press **ENT** on it for two seconds to confirm.

```
Add Enclosure
Spare Drive  ..
```

Step 4. A message prompts to confirm a successful configuration. Press **ESC** to skip the message

```
Add Spare Drive
Successful !!
```

Step 5. The disk drive should now be indicated as an Enclosure spare.

```
C=* I=* XXXXXMB
ENCLOS Vendor
```

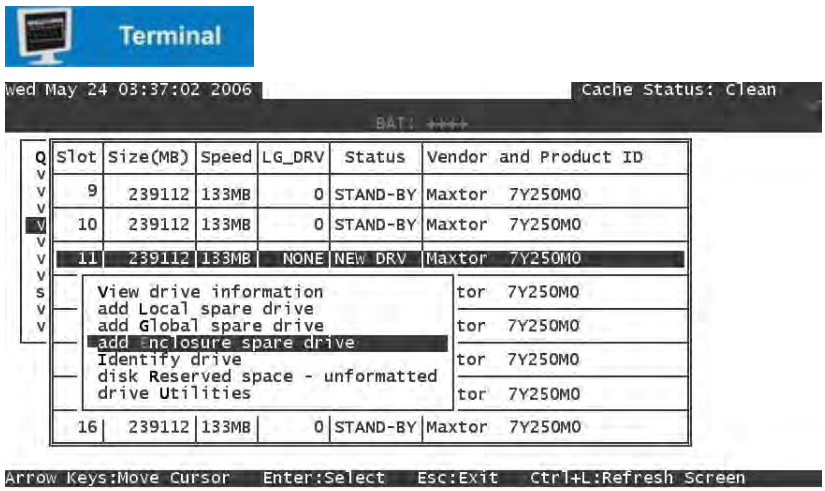

Step 1. In "**View and Edit Drives**" menu, press **[ENTER]** on a "**NEW**" or "**FORMATTED**" drive to display options related to individual disk drives.

Step 2. Move cursor bar to "**add Enclosure spare drive**." Select **Yes** on the confirm box to complete the process.

4.6 Drive Delayed Write

This option applies to disk drives that come with embedded **read-ahead or writer buffers**. When enabled, the embedded buffer can improve read/write performance. However, this option should be disabled for mission-critical applications. In the event of power outage or drive failures, data cached in drive buffers may be lost, and data inconsistency will occur.

Following are the defaults for different storage configurations:

- On **dual-controller** models that come with BBUs, the default is "Disabled."
- On **single-controller** models that come without BBUs, the default is "Enabled."

LCD Keypad

Step 1. Select "**View and Edit Config Parms**" from the terminal Main Menu. Enter its sub-menus by pressing **ENT**.

```
View and Edit
Channels      ↕
```

Step 2. Use arrow keys to select "**Drive-side Parameters**." press **ENT** to enter its sub-menus.

```
Drive-side
Parameters    ..
```

Step 3. There are a dozen of configurable options. Use the arrow keys to select "**Drive Delayed Write -**." Press **ENT** on it to change the setting.

```
Drive Delayed
Write - Disabled
```

Terminal

The configuration option can be found under "**View and Edit Configuration Parameters**"-> "**Drive-side Parameters**"-> "**Drive Delayed Write**."

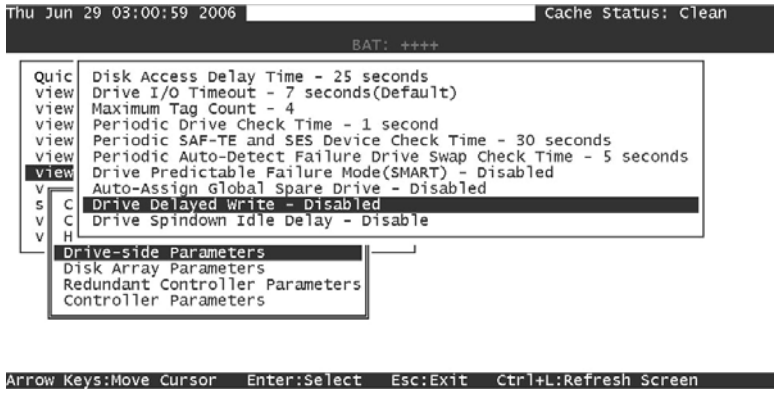

4.7 System Functions

Choose "**System Functions**" in the Main Menu, then press **ENT**. Press the up or down arrow keys to select a submenu, then press **ENT**.

4.7.1 Mute Beeper

LCD Keypad

When the controller's beeper has been activated, choose "Mute Beeper," then press **ENT** to turn the beeper off temporarily for the current event. The beeper will still activate on the next event.

```
Mute Beeper
                ..
```

Terminal

```
< Main Menu >
Quick installation
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
view and edit Drives
view and edit channelS
view and edit Configuration parameters
view and edit Peripheral devices
system Functions
V
V  Mute beeper
C
R     Mute Beeper ?
S
C     Yes      No
```

When the subsystem's beeper (onboard alarm) is activated, choose "**Mute Beeper**," then press **[ENTER].** Choose **Yes** and press **[ENTER]** in the next dialog box to turn the beeper off temporarily for the current event. The beeper will still be activated by the next event.

4.7.2 Password

Use the controller's password to protect the system from unauthorized entry. Once the controller's password is set, regardless of whether the front panel, the RS-232C terminal interface or the RAIDWatch Manager is used, the user can only configure and monitor the RAID controller by providing the correct password.

IMPORTANT!

- The controller requests a password whenever a user is entering the main menu from the initial screen or a configuration change is made. If the controller is going to be left unattended, the "Password Validation Timeout" should be set to "Always Check."

- The controller password and controller name share a 32-character space. The maximum number of characters for a controller password is 32. If 31 characters are used for a controller name, there will be only one character left for the controller password and vice versa. Since revision 3.42J09, later firmware revisions support a 32-character space.

4.7.2.1 Change Password

Step 1. To set or change the controller password, press the up or down arrow keys to select "Change Password," then press **ENT**.

```
Change Password
                ..
```

If the password has previously been set, the controller will ask for the old password first. If password has not yet been set, the controller will directly ask for the new password. The password cannot be replaced unless the correct old password is provided.

```
Old     Password
                ..
```

Step 2. Press the up or down arrow keys to select a character, then press **ENT** to move to the next space.

```
Re-Ent  Password
                ..
```

Step 3. After entering all the characters (alphabetic or numeric), press **ENT** for two seconds to confirm. If the password is correct, or there is no preset password, it will ask for the new password.

```
Change Password
Successful
```

Enter the password again to confirm.

To disable or delete the password, press **ENT** on the first flashing digit for two seconds when requested to enter a new password. The existing password will be deleted. No password checking will occur when entering the Main Menu from the initial terminal screen or making configuration changes.

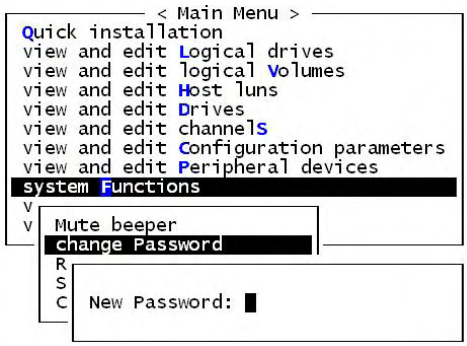

To set or change the controller password, move the cursor bar to "**Change Password**," then press **[ENTER]**.

If a password has previously been set, the controller will ask for the old password first. If the password has not yet been set, the controller will directly ask for the new password. The password cannot be replaced unless the correct old password is provided.

Key-in the old password, then press **[ENTER]**. If the password is incorrect, it will not allow you to change the password. Instead, it will display the message "**Password incorrect!**," then retrun to the previous menu.

If the password is correct, or there is no preset password, it will request for a new password.

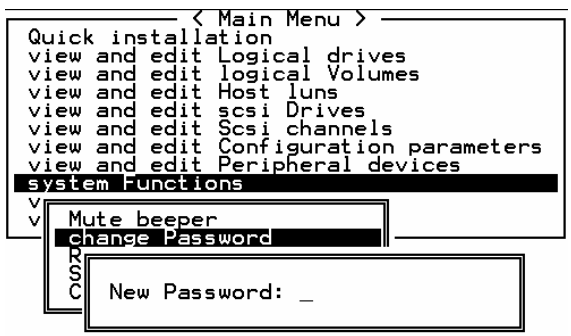

Enter the desired password in the column, then press **[ENTER].** The next dialog box will display "**Re-Enter Password**." Enter the password again to confirm and press **[ENTER]**.

The new password will now become the controller's password. Providing the correct password is necessary when entering the Main Menu from the initial screen.

4.7.3 Reset Controller

Step 1. To reset the controller without powering off the system, press the up or down arrow keys to "Reset Controller," then press **ENT**.

```
Reset This
Controller    ..
```

```
Reset This
Controller    ?
```

Step 2. Press **ENT** again for two seconds to confirm. The controller will now reset.

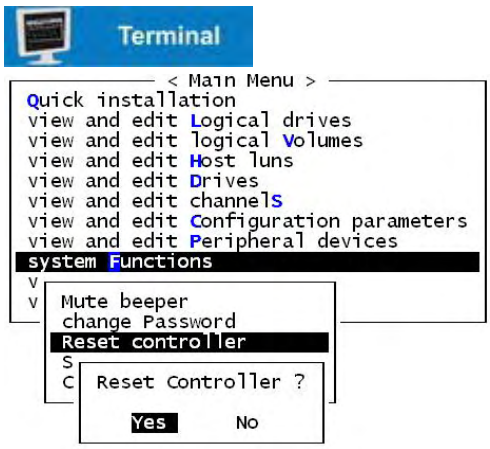

To reset the controller without powering off the system, move the cursor bar to "**Reset Controller**," then press **[ENTER]**. Choose ***Yes*** in the dialog box that follows, then press **[ENTER].** The controller will now reset.

NOTE:

Before resetting or powering off the RAID controller (subsystem) it is advised you execute the Shutdown Controller function to flush the cache contents in the memory in order to reduce the chance of encountering data inconsistency.

4.7.4 Shutdown Controller

LCD Keypad

Before powering off the controller, unwritten data may still reside in cache memory. Use the "Shutdown Controller" function to flush the cache content.

```
Shutdown This
Controller    ?
```

```
Shutdown This
Controller    ..
```

Step 1. Press the up or down arrow keys to "Shutdown Controller," then press **ENT**. Press **ENT** again for two seconds to confirm.

Step 2. The controller will now flush the cache memory. Press **ENT** for two seconds to confirm and to reset or power off the subsystem.

```
ShutdownComplete
   Reset Ctlr?
```

Terminal

Before powering off the controller, unwritten data may still reside in cache memory. Use the "Shutdown Controller" function to flush the cache content. Move the cursor bar to "**Shutdown Controller,**" then press [**ENTER**]. Choose ***Yes*** in the dialog box that follows, then press [**ENTER**]. The controller will now flush the cache memory.

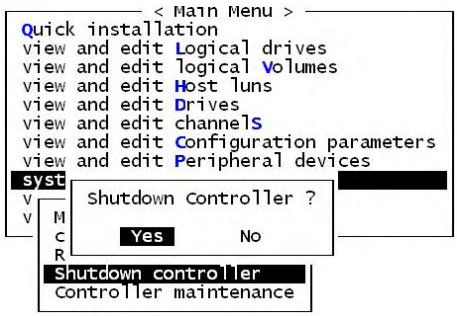

For Controller Maintenance functions, such as "**Download Firmware**," please refer to ***Appendix B***.

4.7.5 Controller Maintenance

For Controller Maintenance functions, please refer to ***Appendix C***.

4.7.6 Saving NVRAM to Disks

You can choose to backup your controller-dependent configuration information to disks. We strongly recommend using this function to save the configuration profile whenever a configuration change is made. The information will be distributed to every logical drive in the RAID system. If using the RAIDWatch manager, you can save your configuration details as a file to a computer system drive.

A RAID configuration of drives must exist for the controller to write NVRAM content onto it.

LCD Keypad

Step 1. From the Main Menu, choose "System Functions." Use arrow keys to scroll down and select "Controller Maintenance," "Save NVRAM to Disks," then press **ENT.**

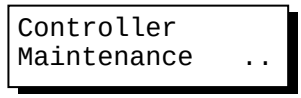

Step 2. Press **ENT** for two seconds on the message prompt, "Save NVRAM to Disks?".

```
Restore with
Password    ?
```

A prompt will inform you that NVRAM information has been successfully saved.

Terminal

At least a RAID configuration must exist for the controller to write your configuration data onto it.

From the Main Menu, choose "**System Functions**." Use the arrow keys to scroll down and select "**Controller Maintenance**," "**Save NVRAM to Disks**," then press **[ENTER].**

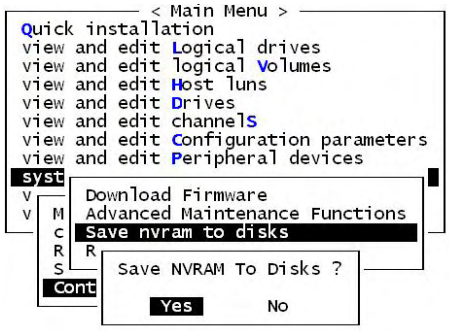

Choose **Yes** to confirm.

A prompt will inform you that NVRAM information was successfully saved.

4.7.7 Restore NVRAM from Disks

If you want to restore your NVRAM information that was previously saved onto the array, use this function to restore the configuration setting.

LCD Keypad

Step 1. From the Main Menu, choose "System Functions." Use arrow keys to scroll down and select "Controller Maintenance," "Restore NVRAM from Disks..," and then press **ENT**.

```
Restore NVRAM
 from Disks    ?
```

Step 2. Press **ENT** for two seconds to confirm.

Step 3. In case your previous password (reserved at the time you saved your NVRAM configuration contents) is different from your current password, you are provided with the options whether to restore the password you previously saved with your configuration profile.

```
Restore without
Password    ?
```

A prompt will inform you that the controller NVRAM data has been successfully restored from disks.

Terminal

From the Main Menu, choose "**System Functions**." Use the arrow keys to scroll down and select "**Controller Maintenance**," "**Restore NVRAM from disks**," and then press **[ENTER].**

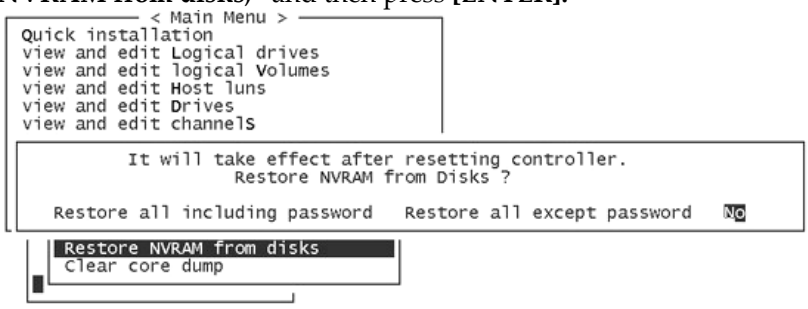

CAUTION!

In case your previous password (reserved at the time you saved your NVRAM configuration contents) is different from your current password, you are provided with the options whether to restore the password you previously saved.

Press **[ENTER]** on one of the selections to confirm.

A prompt will notify you that the controller NVRAM data was successfully restored from disks.

4.7.8 Clear Core Dump

WARNING!

Upon seeing core dump events, power down and reboot your

system after checking system events and correcting system faults. It is highly recommended to contact technical support immediately.

Please **DO NOT** clear the core dump data before causes of failures can be verified and corrected.

Clear Core Dump is a last resort option that helps debug critical issues in the event of serious system faults. When system firmware detects critical errors (such as multi-bit errors, PCI Bus Parity errors, etc.), it distributes configuration and error codes in cache memory into a core file in disk reserved space. Chances are you can refer to these error codes from the core file conducted onto drive media if system finally crashes.

If system is recovered from serious faults later, you can execute the **Clear Core Dump** function to release disk space.

4.7.9 Adjust LCD Contrast

The controller LCD contrast is set at the factory to a level that should be generally acceptable. The controller is equipped with an LCD contrast adjustment circuit in case the factory-preset level needs to be adjusted either via the RS-232 terminal emulation menus or using the LCD keypad panel.

LCD Keypad

Step 1. From the main menu, choose "View and Edit Peripheral Dev."

```
View and Edit
Peripheral Dev..
```

Step 2. Press **ENT** on it, press arrow keys to scroll down, and select "Adjust LCD Contrast," press **ENT** to proceed, and then use the arrow keys to find an optimal setting.

```
Adjust LCD
Contrast     ..
```

```
LCD Contrast :
■■■■■■■------
```

Step 3. Press **ESC** to return to the previous menu.

4.8 Controller Parameters

4.8.1 Controller Name

The controller name represents a RAID subsystem in a topology that consists of numerous RAID subsystems. With dual-controller configurations, only one controller name is applied and will pass down to the surviving controller in the event of single controller failure.

LCD Keypad

Step 1. Select "View and Edit Config Parms" from the Main Menu.

```
Controller Name:
_ _ _ _
```

Step 2. Choose "View and Edit Configuration Parameters," "Controller Parameters," then press **ENT**.

```
Enter Ctlr Name:
□
```

Step 3. The current name will be displayed. Press **ENT** for two seconds and enter the new controller name by using the up or down arrow keys. Press **ENT** to move to another character and then press **ENT** for two seconds on the last digit of the controller name to complete the process.

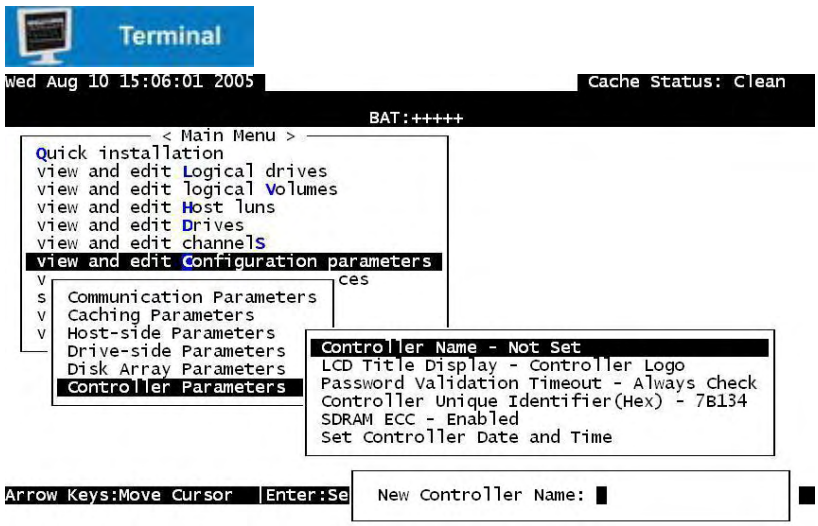

Choose "**View and Edit Configuration Parameters**," "**Controller Parameters**," then press **[ENTER]**. The current name displays. Press **[ENTER]**. Enter a name in the dialog box that prompts, then press **[ENTER]**.

4.8.2 LCD Title Display Controller Name

LCD Keypad

Step 1. Choose "View and Edit Configuration Parameters," "Controller Parameters," then press **ENT**.

```
LCD Title Disp –
Controller Logo?
```

Step 2. Use the up or down arrow keys to choose to display the embedded controller logo or any given name on the LCD initial screen.

```
LCD Title Disp –
Controller Name?
```

Terminal

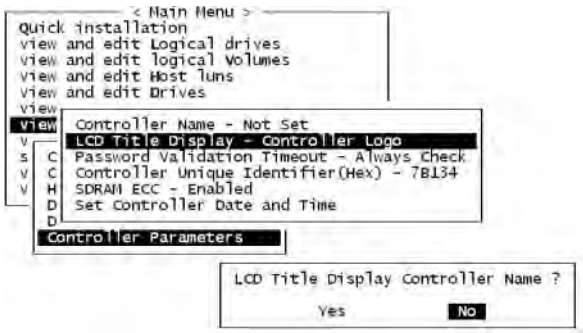

Choose "**View and Edit Configuration Parameters**," "**Controller Parameters**," then press **[ENTER]**. Choose to display the embedded controller model name or any given name on the LCD. Giving a specific name to each controller will make them easier to identify if you have multiple RAID systems that are monitored from a remote station.

4.8.3 Password Validation Timeout

TIPS:

The Always Check timeout will disable any configuration change made without entering the correct password.

LCD Keypad

Step 1. Choose "View and Edit Configuration Parameters,"

```
PasswdValidation
Timeout-5 mins..
```

"Controller Parameters," then press **ENT**.

Step 2. Select "Password Validation Timeout," and press **ENT**. Press the up or down arrow keys to choose to enable a validation timeout from one to five minutes, or to "Always Check."

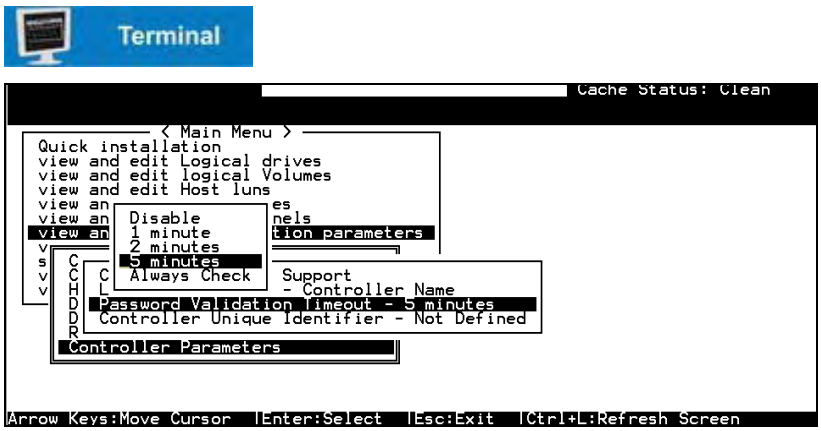

Choose "**View and Edit Configuration Parameters**," "**Controller Parameters**," then press **[ENTER]**. Select "**Password Validation Timeout**," and press **[ENTER]**. Choose to enable a validation timeout from "**1 minute**" to "**Always Check**."

4.8.4 Controller Unique Identifier

What is the Controller Unique Identifier?

1. A specific identifier helps RAID controllers to identify their counterpart in a dual-active configuration.

2. The unique ID is generated into a Fibre Channel WWN node name for RAID controllers or RAID subsystems using Fibre Channel host ports. The node name prevents host computers from misaddressing the storage system during the controller failover/failback process in the event of single controller failure.

3. The unique ID is also generated into a MAC address for the controller's Ethernet port. The MAC address will be taken over by a surviving controller in the event of single RAID controller failure.

4. When a controller fails and a replacement is combined as the secondary controller, the FC port node names and port names will be passed down to the replacement controller. The host will not acknowledge any differences so that controller failback is totally transparent.

LCD Keypad

```
Ctlr Unique (Hex)-
ID-
```

Step 1. Choose "View and Edit Configuration Parameters," "Controller Parameters," then press **ENT**.

Step 2. Press the up or down arrow keys to select "Ctlr Unique ID-," then press **ENT**.

Step 3. Enter any hex number between "0" and "FFFFF" and press **ENT** to proceed.

NOTE:

Usually every RAID subsystem/controller comes with a default ID. In rare occasions should this identifier be changed.

Terminal

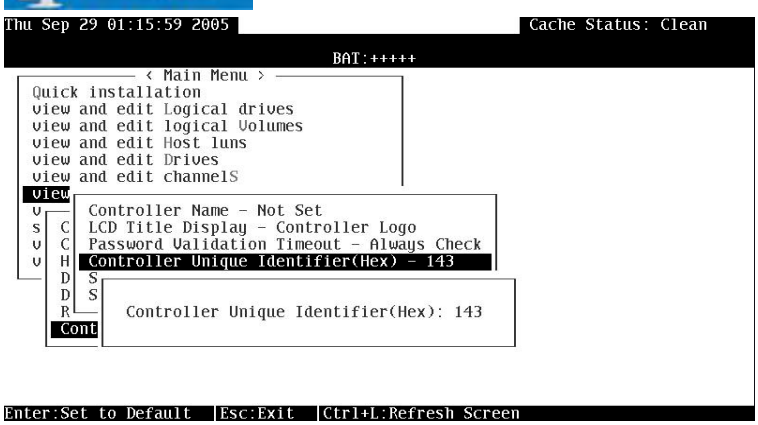

Enter any hex number between "**0**" and "**FFFFF**" for the unique identifier. The value you enter ***MUST*** be different for each controller.

Every EonStor subsystem comes with a default ID. This ID should be sufficient for avoiding WWNN and WWPN conflicts.

4.9 System Defaults

Table 4-2: System Firmware Defaults

Configurable Options	Default
Event Triggered Operations:	
Controller failure	Disabled
BBU low or failed	Enabled
UPS AC power loss	Disabled
Power supply failure	Disabled
Fan failure	Disabled
Temperature exceeds threshold	Disabled
Host-side Parameters	
Maximum Queued IO Count	1024 for ASIC400 models
LUNs per host ID	8
Max. Number of Concurrent Host-LUN Connections	4
Number of Tags Reserved for Each Host-LUN Connection	32
Fibre Connection option	Loop only
Peripheral Device Parameters (for in-band management access only)	
Peripheral device type	Enclosure Service Device (0xD)
Peripheral device qualifier	Connected
Device support removable media	Disabled
LUN applicability	All undefined LUNs
Cylinder/Head/Sector- variables	N/A
Drive-side Parameters	
Disk Access Delay Time	Per product interface; will be a larger values in multi-enclosure applications
Drive I/O Timeout	7 seconds
Max. Tag Count	8: Fibre Channel
Periodic SAF-TE and SES Check Time	30 seconds
Auto Rebuild on Drive Swap check time	15 seconds
Drive Predictable Failure Mode (S.M.A.R.T.)	Disabled
Drive Delayed Write	Enabled (single-controller, w/o BBU) Disabled (dual-controller, w/ BBU)
Drive Spindown Idle Delay	Disabled

Voltage & Temperature Parameters	
+3.3V thresholds	3.6V – 2.9V
+5V thresholds	5.5V – 4.5V
+12V thresholds	13.2V - 10.8V
CPU temperature	90 - 5°C
Board temperature (RAID controller board)	80 - 5°C
* The thresholds for other sensors within the chassis are not user-configurable. It is user's responsibility to maintain a reasonable ambient temperature, e.g., 35°C at the installation site.	
Disk Array Parameters	
Rebuild Priority	Normal
Verification on Write	
Verification on LD Initialization	Disabled
Verification on LD Rebuild	Disabled
Verification on Normal Drive Writes	Disabled
Max. Drive Response Timeout	Disabled
AV Optimization	Disabled

Chapter

5

Creating RAID Arrays & Host LUN Mapping

(Using LCD Keypad Panel)

A navigation roadmap for the configuration menu options through LCD keypad is separately available as a PDF file. You may check your Product Utility CD or contact technical support for the latest update.

5.1 Power On RAID Enclosure

Before you start to configure a RAID system, make sure that hardware installation is completed before any configuration takes place. Power on your RAID subsystem.

NOTE:

- If your EonStor RAID subsystem comes with dual-redundant RAID controllers, your subsystem's LCD panel can provide access to the operating status screen of the Secondary controller. However, in a dual-controller configuration, only the Primary controller responds to user's configuration.

- Each controller's operating mode is indicated by the flashing digit on the upper right of the LCD screen as "A" or "B." If the LCD displays "B," that means the LCD screen is currently displaying Slot B controller messages. Press both the Up and Down arrow keys for one second to switch around the access

to different RAID controllers.

5.2 Individual Disk Menu:

Drive Utilities

Step 1. From the "View and Edit Drives" menu, select the drive that the utility is to be performed on; then press **ENT**.

Step 2. Select "Drive Utilities; then press **ENT**. Choose "Read/Write Test".

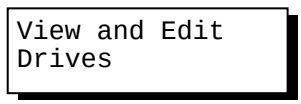

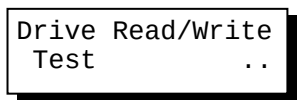

These options are not available for drives already configured in a logical configuration, and can only be performed before a reserved space is created on a disk drive.

5.2.1 Drive Read/Write Test

From the "View and Edit Drives" menu, select a new or used drive that the utility is to be performed on; then press **ENT**. Select "Drive Utilities;" then press **ENT**.

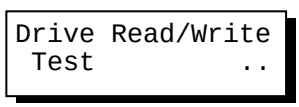

Step 3. Choose "Read/Write Test" and press **ENT**.

Tips:

This function will not be available with a "formatted" or "used" drive. You may manually remove the 256MB reserved space to re-configure a "used" drive into a "new" drive.

Step 4. Press the up or down arrow keys to select and choose to enable/disable the following options:

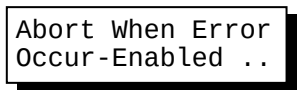

1. **Read-Write Test Type** – "Read Only" or "Read and Write."

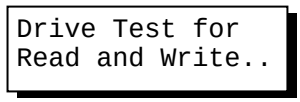

2. **Action on Detection of Error**

- Continue Test
- Abort on Detection of Any Errors (including hardware, media errors, etc.)
- Abort on Detection of Hard Errors only (only hardware errors in 8-bit encoding)

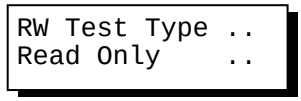

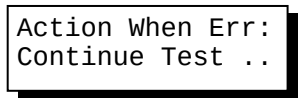

3. **Recovery Process**

- No Recovery
- Mark Blk Bad (mark the affected blocks as bad)
- Reassignment (try to reassign data residing on the affected sectors to other drive sectors)
- Frst Try Reassi (Recovery by attempting to reassign and mark bad if reassignment fails)

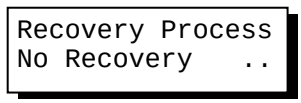

4. **Execute Drive Testing**

Step 5. When finished with configuration, select "Execute Drive Testing" and press **ENT** to proceed.

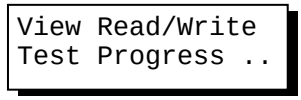

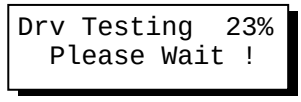

Step 6. The Read/Write test progress will be indicated as a percentage.

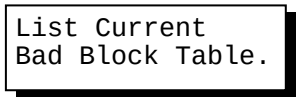

Step 7. You may press **ESC** and select "Read/Write Test" later and press the up or down arrow keys to select to "**View Read/Write Testing Progress**" or to "**List Current Bad Block Table**." If you

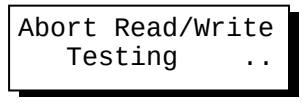

want to stop testing the drive, select "**Abort Drive Testing**" and press **ENT** to proceed.

View Physical Drives

A RAID subsystem consists of many physical drives. These drives can be grouped into one or several logical configurations.

Step 1. Press the front panel **ENT** button for two seconds to enter the Main Menu. Use the up or down arrow keys to navigate through the menus. Choose "View and Edit Drives," then press **ENT**.

```
View and Edit
Drives
```

Step 2. Use the up or down arrow keys to scroll down the list of connected drives' information screens.

```
Slot=* 1010MB
New DRV SEAGATE
```

You may first examine whether there is any drive installed but not shown here. If there is a drive installed but not listed, the drive may be defective or not installed correctly. Please check your enclosure installation and contact your system vendor.

Step 3. Press **ENT** on a drive. Choose "View Drive Information" by pressing **ENT**. Use the up or down arrow keys to navigate through the screens.

```
View Drive
Information   ..
```

Step 4. The Revision Number of the selected drive will be shown. Press the down arrow key to see other information.

```
Revision Number:
0274
```

Other information screens include "**Serial Number**," "**Disk Capacity**" (displayed in blocks; each block equals 512K Bytes), and "**Node Name (WWNN)**."

IMPORTANT!

- Drives of the same brand/model/capacity might not feature the same block number.
- The basic read/write unit of a hard drive is logical block. If members of a logical drive have different block numbers (capacity), the smallest block number will be taken as the maximum capacity to be used in every drive for constructing a logical array. Therefore, use drives of the same capacity so your investment is not wasted.
- You may assign a Local Spare Drive to a logical drive whose members have a block number equal to or smaller than the Local Spare Drive but you should not do the reverse.

5.3 Creating a Logical Drive

Step 1. To create a logical drive, press **ENT** for two seconds to enter the Main Menu. Use the up or down arrow keys to navigate through the menus. Choose "View and Edit Logical Drives," and then press **ENT**.

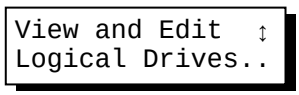

Step 2. Press the up or down arrow keys to select a logical drive entry, then press **ENT** for two seconds to proceed. "LD" is short for Logical Drive.

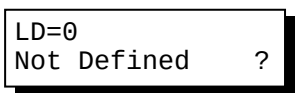

5.3.1 Choosing a RAID Level

Step 3. Use the up or down arrow keys to select the desired RAID level, then press **ENT** for two seconds. "**TDRV**" (Total Drives) refers to the number of all available disk drives.

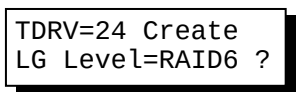

5.3.2 Choosing Member Drives

Step 4. Press **ENT** for two seconds; the message, "**RAID X selected To Select drives**", will prompt. Confirm your selection by pressing **ENT**.

```
RAID 6  Selected
To Select drives
```

Step 5. Press **ENT**, then use the up or down arrow keys to browse through the available drives.

```
Slot=* 239367MB
New DRV SEAGATE
```

Step 6. Press **ENT** again to select/deselect individual disk drives. An asterisk (*) mark will appear on the selected drive(s). To deselect a drive, press **ENT** again on the selected drive. The (*) mark will disappear.

```
SLOT=1 239367MB
*LG=0 SL SEAGATE
```

Step 7. After all the desired hard drives have been selected, press **ENT** for two seconds to continue.

5.3.3 Logical Drive Preferences

Step 8. You may enter the following screen to "**Change Logical Drive Parameter**" by pressing **ENT** before initializing the logical drive.

```
Change Logical
Drive Parameter?
```

5.3.4 Maximum Drive Capacity

Step 8-1-1. Choose "**Maximum Drive Capacity**," then press **ENT**. The maximum drive capacity refers to the maximum capacity that will be used in each member drive.

```
Maximum Drive
Capacity      ..
```

Step 8-1-2. If necessary, use the up and down arrow keys to change the maximum size that will be used on each drive.

```
MaxSiz= 239367MB
Set to 239367MB?
```

5.3.5 Spare Drive Assignments

The Local (Dedicated) Spare Drive can also be assigned here.

Step 8-2-1. Press the up or down arrow keys to choose "**Spare Drive Assignments**," then press **ENT**.

Step 8-2-2. Available disk drives will be listed. Use the up or down arrow keys to browse through the drive list, then press **ENT** to select the drive you wish to use as the Local (Dedicated) Spare Drive.

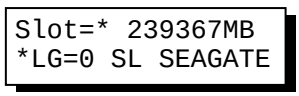

Step 8-2-3. Press **ENT** again for two seconds.

5.3.6 Disk Reserved Space

This menu allows you to see the size of disk reserved space. Default is 256MB. The reserved space is used for storing array configuration and other non-volatile information.

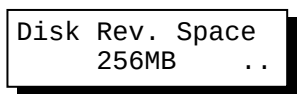

5.3.7 Logical Drive Assignment (Dual-Active Controllers)

If your subsystem is powered by dual-redundant RAID controllers, you may choose to assign a logical drive to a specific RAID controller. The assignment can take place during or after the initial configuration.

IMPORTANT!

Firmware releases before rev. **3.48**:

A logical drive can be assigned to **either** the Primary controller **or** the Secondary controller.

Firmware releases rev. **3.48** or later:

A logical drive can be assigned to **both** the Primary and Secondary controllers. Since the Primary/Secondary role can change during a controller failover, RAID controllers are identified as **Slot A** and **Slot B** controllers in firmware rev. 3.48.

Slot A refers to the upper controller slot, which is the default location for the dominating or master controller. **Slot B** refers to the lower controller slot.

Step 8-3-1. In a dual-controller configuration, the assignment menus should appear as listed on the right.

```
Change logical
Drv Parameter ..
```

Step 8-3-2. If you want to assign this logical drive to the Slot B controller, use your arrow keys and press **ENT** once to display the question mark "?". Press again for two seconds to confirm the change.

```
Logical Drive
Assignments   ..
```

```
Red Ctlr LG
Assign to Slot B?
```

5.3.8 Write Policy

This menu allows you to set the caching mode policy for this specific logical drive. "**Default**" is a neutral value that is coordinated with the subsystem's general caching mode setting. Other choices are "**Write-back**" and "**Write-through**."

```
Write Policy
Default      ..
```

```
Write Policy
Default      ?
```

Step 8-4-1. Press **ENT** once to change the status digits into a question mark "?".

Step 8-4-2. Use the arrow keys to select "**Default**," "**Write-back**," or "**Write-through**."

Step 8-4-3. Press **ENT** for two seconds to confirm your change.

IMPORTANT!

- The "**Write-back**" and "**Write-through**" parameters are permanent for specific logical drives. The "Default" selection, however, is more complicated and more likely equal to "not specified."
- If set to "**Default**," a logical drive's write policy is determined not only by the system's general caching mode setting, but also by the "**Event trigger**" mechanisms. The "Event Trigger" mechanisms automatically disable the write-back caching and adopt the conservative "Write-through" mode in the event of battery or component failures.

5.3.9 Initialization Mode

This menu allows you to determine if the logical drive is immediately accessible. If the Online method is used, data can be written onto it before the array's initialization is completed. You may continue with other array configuration processes, e.g., including this array in a logical volume.

```
Initialization
Mode   Online..
```

Array initialization can take a long time especially for those comprising a large capacity and parity data. Setting to "Online" means the array is immediately accessible and that the controller will complete the initialization in the background or I/Os become less intensive.

Step 8-5-1. Press **ENT** once to change the status digits into a question mark "?".

```
Initialization
Mode   Online?
```

Step 8-5-2. Use the arrow keys to select either the "**Online**" or the "**Off-line**" mode.

Step 8-5-3. Press **ENT** for two seconds to confirm your change.

5.3.10 Stripe Size

```
Stripe size
Default      ..
```

This menu allows you to change the array stripe size. Setting to an incongruous value can severely drag performance. This item should only be changed when you can test the combinations of different I/O sizes and array stripe sizes and can be sure of the performance gains it might bring you.

For example, if the I/O size is 256k, data blocks will be written to two of the member drives of a 4-drive array while the RAID firmware will read the remaining member(s) in order to generate the parity data.

* For simplicity reasons, we use RAID3 in the samples below.

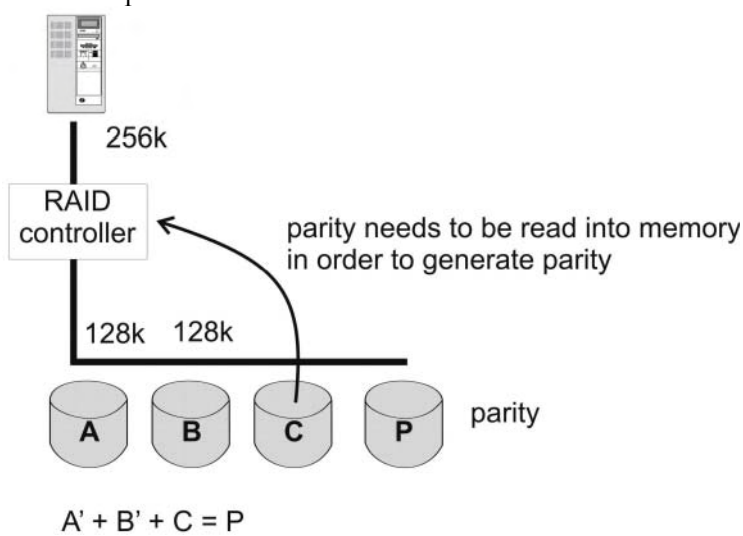

Figure 5-1 256k I/O size & an Array with a 128k Stripe Size

In an ideal situation, a 384k I/O size allows data to be written to 3 member drives and parity data to be simultaneously generated without the

effort to consult data from other members of an array.

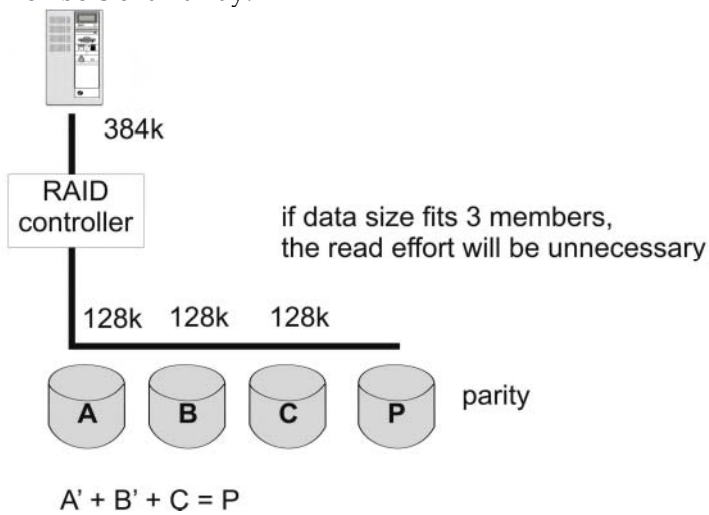

Figure 5-2 384k I/O size & an Array with a 128k Stripe Size

If the I/O size is larger than the combined stripe depths, the extra data blocks will be written to the member drives on the successive spins, and the read efforts will also be necessary for generating parity data.

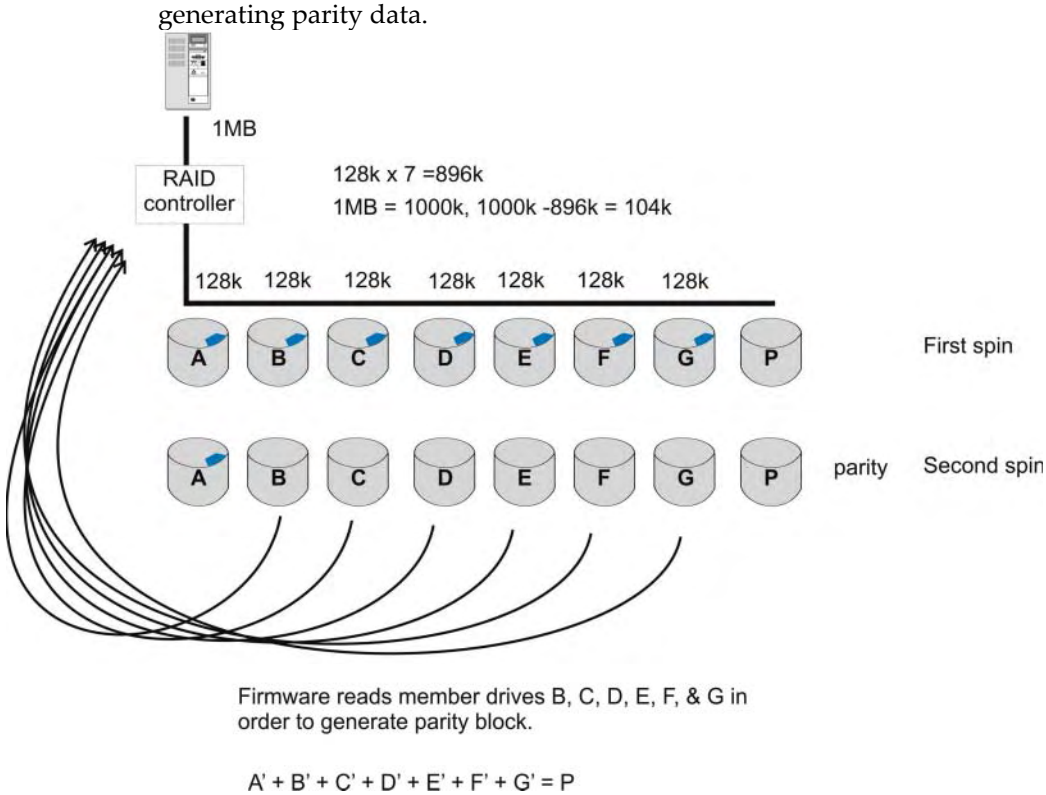

Figure 5-3 1MB I/O size & an Array with a 128k Stripe Size and 8 members

Although the real-world I/Os do not always perfectly fit the array stripe size, matching the array stripe size to your I/O characteristics can eliminate drags on performance (hard drive seek and rotation efforts) and will ensure the optimal performance.

Listed below are the default values for different RAID levels.

Table 2-1 RAID Levels and Stripe Sizes

RAID level	Stripe Size
RAID0	128KB
RAID1	128KB
RAID3	16KB
RAID5	128KB
RAID6	128KB
NRAID	128KB

Step 8-6-1. Press **ENT** once to change the status digits into a question mark "?".

```
Stripe size
128KB          ?
```

Step 8-6-2. Use the arrow keys to select a desired stripe size.

Step 8-6-3. Press **ENT** for two seconds to confirm your change.

5.3.11 Beginning Initialization

Step 9. Press **ESC** to return to the previous menu. Use the up or down arrow keys to select "**Create Logical Drive?**".

```
Create Logical
Drive          ?
```

Step 10. Press **ENT** for two seconds to start initializing the logical drive.

The Online Mode:

If the online initialization method is applied, the array will be immediately

```
LG=0 Creation
Completed!
```

available for use. The array initialization runs in the background and the array is immediately ready for I/Os. Engineers can continue configuring the RAID subsystem.

The Offline Mode:

The RAID controller will immediately start to initialize the array parity if the "offline" mode is applied. Note that if NRAID or RAID0 is selected, initialization time is short and completes almost within a second.

```
Initializing090%
Please Wait!
```

```
LG=0 Initializat
Ion Completed
```

The logical drive's information displays when the initialization process is completed. If the "online" mode is adopted, array information will be displayed immediately.

```
LG=0 RAID6 DRV=4
478734MB GD SB=0
```

NOTE:

Due to the operation complexity, the RAID Migration option is not available using the LCD keypad panel.

5.4 Creating a Logical Volume

NOTE:

Logical volume is not a must for RAID configuration.

Step 1. Press **ENT** for two seconds to enter the Main Menu. Press the up or down arrow keys to select "**View and Edit Logical Volume**," then press **ENT**.

```
View and Edit
Logical Volume
```

Step 2. Press the up or down arrow keys to select an undefined entry for logical volume, then press **ENT** for two

```
LV=0
Not Defined    ?
```

seconds to proceed. "**LV**" is short for Logical Volume.

Step 3. Proceed to select one or more logical drives as the members of a logical volume. Press **ENT** to proceed. "**LD**" is short for Logical Drive.

```
LV=0 Selected To
Select LD Drives?
```

Step 4. Use the up or down arrow keys to browse through the logical drives.

```
LG0 RAID5 DRV=3
2021MB GD SB=0
```

Step 5. Press **ENT** again to select/deselect the members. An asterisk (*) mark will appear in front of a selected logical drive.

Step 6. After all the desired logical drive(s) have been selected, press **ENT** for two seconds to continue.

Two sub-menus will appear.

5.4.1 Initialization Mode

Array initialization can take a long time especially for those comprised of a large capacity. Setting to "Online" means the array is immediately accessible and that the controller will complete the initialization in the background and when I/O demands become less intensive.

```
Initialization
Mode    Online..
```

Step 7-1-1. Press **ENT** once to change the status digits into a question mark "?".

```
Initialization
Mode    Online?
```

Step 7-1-2. Use the arrow keys to select either the "**Online**" or the "**Off-line**" mode.

Step 7-1-3. Press **ENT** for two seconds to confirm your change.

5.4.2 Write Policy

This menu allows you to set the caching mode policy for this specific logical volume. "**Default**" is a neutral value that is coordinated with the controller's general caching mode setting. Other choices are "**Write-back**" and "**Write-through**."

```
Write Policy
Default        ..
```

Step 7-2-1. Press **ENT** once to change the status digits into a question mark "?".

Step 7-2-2. Use the arrow keys to select "**Default**," "**Write-back**," or "**Write-through**."

Step 7-2-3. Press **ENT** for two seconds to confirm your change.

IMPORTANT!

- The "**Write-back**" and "**Write-through**" parameters are permanent for specific logical drives. The "Default" selection, however, is more complicated and more likely equal to "not specified."
- If set to "**Default**," a logical drive's write policy is controlled not only by the subsystem's general caching mode setting, but also by the "**Event trigger**" mechanisms. The "Event Trigger" mechanisms automatically disable the write-back caching and adopt the conservative "Write-through" mode in the event of a battery or component abnormalities.

Step 8. When you are finished setting the preferences, press **ENT** for two seconds to display the confirm box. Press **ENT** for two seconds to start initializing the logical volume.

```
Create
Logical Volume ?
```

```
Lv=0 Creation
Completed
```

Step 9. A message shows that the logical volume has been successfully created.

```
Create Logical
Volume Successed
```

Step 10. Press **ESC** to clear the message.

Logical volume information will be displayed next.

```
Lv=0 ID=07548332
0024488MB DRV=2
```

NOTE:

Once a logical drive is included in a logical volume, its "Controller Assignment" option will disappear. The controller assignment option displays under the logical volume sub-menu instead.

5.4.3 Logical Volume Assignment (Dual-active Controllers)

In a dual-controller configuration, you may choose to assign this logical volume to the Slot B controller (Default is Slot A, the default dominant/master controller). The assignment can take place during or after the initial configuration.

```
Change Logical
Volume Params  ?
```

```
Logical Volume
Assignments   ..
```

In a dual-controller configuration, the assignment menus should appear as listed on the right.

```
Red Ctlr LV
Assign to SlotB?
```

```
Logical Volume
Assignments   ..
```

```
Red Ctlr LV
Assign to SlotB?
```

Step 1. Press **ENT** on a configured logical volume. Use arrow keys to select "**Logical Volume Assignment..**", and press **ENT** to proceed. Press **ENT** for two seconds to confirm.

Step 2. Press **ESC**, and the LCD will display the logical volume's information when initialization is completed.

```
LV=0 ID=685AE502
2021MB   DRV=1
```

5.5 Partitioning a Logical Drive/Logical Volume

IMPORTANT!

Partitioning is **NOT** a requirement for building a RAID subsystem. The configuration processes for partitioning a logical volume are identical to those for partitioning a logical drive.

WARNING!

Partitioning a configured volume will destroy all data stored on it. It is best to partition a logical array during the initial system configuration.

Step 1. Press **ENT** for two seconds to enter the Main Menu. Press the up or down arrow keys to select "**View and Edit Logical Volume**," then press **ENT**.

```
View and Edit
Logical Volume
```

Step 2. Use the up or down arrow keys to select a logical volume, then press **ENT**.

```
LV=0 ID=685AE502
2021MB DRV=1
```

Step 3. Use the up or down arrow keys to select "**Partition Logical Volume**," then press **ENT**.

```
Partition
Logical Volume..
```

Step 4. The total capacity of the logical volume will be displayed as one partition. Press **ENT** for two seconds to change the size of the first partition.

```
LV=0   Prt=0:
    2021MB     ?
```

Step 5. Use the up or down arrow keys to change the number of the flashing digit, (see the arrow mark) then press **ENT** to move to the next digit.

```
LV=0   Part=0:
   2021MB
```

Step 6. After changing all the digits, press **ENT** for two seconds to confirm the capacity of this partition. You may then use arrow keys to move to the

```
LV=0 Part=0:
  700MB       ?
```

next partition to configure more partitions.

```
LV=0 Partition=1
  1321MB       ?
```

The rest of the drive space will be automatically allocated to the next partition. You may repeat the process to create up to 16 partitions using the same method described above.

Step 7. Press **ESC** several times to return to the Main Menu.

IMPORTANT!

If operating with a Unix-based system, reset the subsystem for the configuration to take effect if any changes have been made to partition sizes and partition arrangement.

5.6 Mapping a Logical Volume/Logical Drive to Host LUN

The process of mapping (associating with) a logical drive is identical to that of mapping a logical volume. The process of mapping a logical volume is used here as a configuration sample. The idea of host LUN mapping is diagrammed as follows:

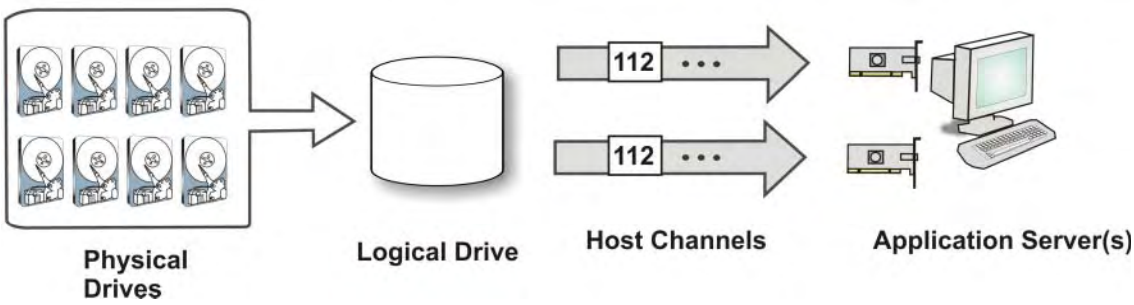

Figure 5-4 Host LUN Mapping

Step 1. Press **ENT** for two seconds to enter the Main Menu. Use the up or down arrow keys to select "**View and Edit Host Luns,**" then press **ENT**.

```
View and Edit  ↕
Host Luns      ..
```

IMPORTANT!

Note the following details before proceeding:

- **Slot A or Slot B IDs**: If using firmware before rev.3.48, logical group of drives (logical drive/logical volume) previously assigned to the Slot A controller cannot be re-assigned to a Slot B controller. Neither can those assigned to the Slot B controller be mapped to a Slot A controller ID.
- Your subsystem comes with Slot A IDs only. You need to manually create Slot B IDs in a dual-controller configuration. Please enter "**View and Edit Channels**" menu to create or remove a host ID.

CAUTION!

Firmware revisions 3.48 and later also support the **cross-controller** ID mapping. The cross-controller mapping allows you to associate a logical drive with BOTH controller A and controller B IDs.

However, mapping to both controllers' IDs is only beneficial when it is difficult making the fault-tolerant host links between RAID controllers and host HBAs, e.g., using SAS-to-SAS RAID systems. Currently, there is no external SAS switch available on the market. For Fibre-host systems, fault-tolerant links can easily be made with the help of external bypass such as Fibre Channel switches.

For details of fault-tolerant link connections, please refer to your system Hardware Manual.

Step 2. The first available ID on the first host channel appears (usually channel0).

```
CH=0 AID=112
             ..
```

Step 3. Press the up or down arrow keys to select a configured host ID, and then press **ENT** for two seconds to confirm.

Step 4. Press the up or down arrow keys to select the type of logical configuration to be associated with a host ID/LUN. Available choices are: "**Map to Logical Drive**" or "**Map to Logical Volume**".

```
Map to
Logical Drive ?
```

```
Map to
Logical Volume ?
```

Step 5. Confirm your choice by pressing **ENT** for two seconds.

Step 6. Press the up or down arrow keys to select a LUN number, then press **ENT** to proceed.

```
CH0 ID112 LUN0
Not Mapped
```

Step 7. Press **ENT** for two seconds to confirm the selected LUN mapping.

```
Map Host LUN   ?
```

Step 8. Press the up or down arrow keys to select a logical drive, a logical volume, or a partition within.

```
LV=0 ID=685AE502
478734MB  DRV=1
```

Step 9. Press **ENT** for two seconds to map the selected partition to this LUN. If the logical configuration has not

```
LV=0 PART=0
   700MB      ?
```

been partitioned, you can map the whole capacity (Partition 0) to a host ID/LUN.

Step 10. Press **ENT** for two seconds when prompted by "**Map Host LUN**" to proceed.

```
Map Host LUN
            ..
```

Step 11. Mapping information will be displayed on the subsequent screen. Press **ENT** for two seconds to confirm the LUN mapping.

```
CH0 ID112 LUN0
MAP to LV0 PRT0?
```

Step 12. The mapping information will appear for the second time. Press **ENT** or **ESC** to confirm, and the host ID/LUN screen will appear.

```
CH0 ID112 LUN0
MAPto LV0 PRT0
```

Step 13. Use the arrow keys to select another ID or LUN number to continue mapping other logical configurations or press **ESC** for several times to leave the configuration menu.

```
CH=0 ID=112
            ..
```

When any of the host ID/LUNs is successfully associated with a logical array, the "**No Host LUN**" message in the initial screen will change to "**Ready**."

5.7 Assigning Spare Drive and Rebuild Settings

5.7.1 Adding a Local Spare Drive

Step 1. Press **ENT** for two seconds to enter the Main Menu. Press the up or down arrow keys to select "**View and Edit Drives**," then press **ENT**.

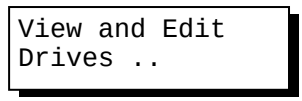

Step 2. Disk drive information will be displayed on the LCD. Press the up or down arrow keys to select a drive that is stated as "**NEW DRV**" or "**USED DRV**" that has not been included in any logical drive, nor specified as a "**FAILED**" drive, then press **ENT** to select it.

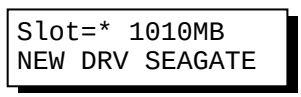

Step 3. Press the up or down arrow keys to select "**Add Local Spare Drive,**" then press **ENT**.

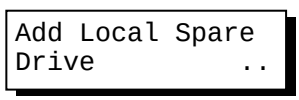

Step 4. Press the up or down arrow keys to select the logical drive where the Local Spare Drive will be assigned, then press **ENT** for two seconds to confirm.

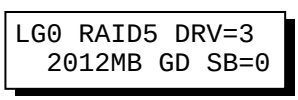

Step 5. The message "**Add Local Spare Drive Successful**" will be displayed on the LCD.

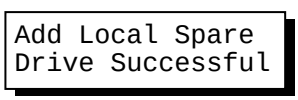

5.7.2 Adding a Global Spare Drive

Step 1. Press **ENT** for two seconds to enter the Main Menu. Press the up or down arrow keys to select "**View and Edit Drives**," then press **ENT**.

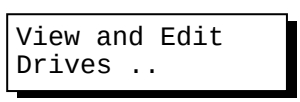

Step 2. Disk drive information will be displayed on the LCD. Press the up or down arrow keys to select a disk drive that has not been assigned to any logical drive, then press **ENT**.

Step 3. Press the up or down arrow keys to select "**Add Global Spare Drive**," then press **ENT**.

```
Add Global Spare
Drive          ..
```

Step 4. Press **ENT** again for two seconds to add the spare drive. The message, "**Add Global Spare Drive Successful**," will be displayed on the screen.

```
Add Global Spare
Drive Successful
```

WARNING!

Assigning a hot-spare to an array composed of drives of a different interface type should be avoided. For example, a SATA Global spare may accidentally participate in the rebuild of an array using SAS members. It is better to prevent mixing SAS and SATA drives in a logical drive configuration.

5.7.3 Adding an Enclosure Spare Drive

In environments where RAID volumes might span across several enclosures, e.g., using JBODs, this option can designate a spare drive to rebuild a failed drive within the same enclosure.

Step 1. To create an Enclosure Spare Drive, press **ENT** for two seconds to enter the Main Menu. Press the up or down arrow keys to select "**View and Edit Drives**," then press **ENT**.

Step 2. Disk drive information will be displayed on the LCD. Press the up or down arrow keys to

select a disk drive that has not been assigned to any logical drive, then press **ENT**.

Step 3. Press the up or down arrow keys to select "**Add Enclosure Spare Drive**," then press **ENT**.

```
Add Enclosure
Spare Drive  ..
```

Step 4. When the last digit changes to a question mark "**?**", press **ENT** again for two seconds to create the enclosure spare. The message, "**Add Spare Drive Successful**," will be displayed on the screen.

```
Add Enclosure
Spare Drive  ?
```

```
Add Spare Drive
Successful !!
```

Step 5. Press **ESC** and the drive status displays as shown on the right.

```
Slot=0 239112MB
ENCLOS Maxtor  7
```

5.8 Deleting Existing RAID Arrays

5.8.1 Deleting a Logical Drive

WARNING!

Deleting a logical drive erases all data stored in it.

Step 1. Press **ENT** for two seconds to enter the Main Menu. Press the up or down arrow keys to select "**View and Edit Logical Drives**," then press **ENT**.

```
View and Edit
Logical Drives
```

Step 2. Press the up or down arrow keys to select a logical drive, then press **ENT**.

```
LG0 RAID5 DRV=3
  2012MB GD SB=1
```

Step 3. Use the up or down arrow keys to select "**Delete Logical Drive**," then press **ENT**.

```
Delete Logical
Drive         ..
```

Step 4. Press **ENT** for two seconds to confirm.

```
LG=0
Not Defined   ?
```

5.8.2 Deleting a Partition of a Logical Drive

WARNING!

Whenever there is a partition change, data will be erased. Prior to partition change, you have to remove its associated host LUN mappings. After the partition change, you also need to re-arrange the disk volumes from your host system OS.

Step 1. Press **ENT** for two seconds to enter the Main Menu. Press the up or down arrow keys to select "**View and Edit Logical Drives..**," then press **ENT**.

```
View and Edit
Logical Drives
```

Step 2. Press the up or down arrow keys to select a logical drive, then press **ENT**.

```
LG0 RAID5 DRV=3
  2012MB GD SB=1
```

Step 3. Press the up or down arrow keys to choose "**Partition Logical Drive**," then press **ENT**.

```
Partition
Logical Drive ..
```

Step 4. The first partition's information will be shown on the LCD. Press the up or down arrow keys to browse through the existing partitions in the logical drive. Select a partition by pressing **ENT** for two seconds.

```
LG=0 Prt=0
   200MB       ?
```

```
LG=0 Prt=1
   300MB       ?
```

Step 5. Use the up or down arrow keys to change the number of the flashing digit to "**0**," then press **ENT** to move to the next digit. After changing all the digits, press **ENT** for two seconds.

```
LG=0 Prt=2
   600MB       ?
```

The disk space of the deleted partition will be automatically allocated to the previous partition as diagrammed below. For example, if partition 1 is deleted, its disk space will be added to partition 0.

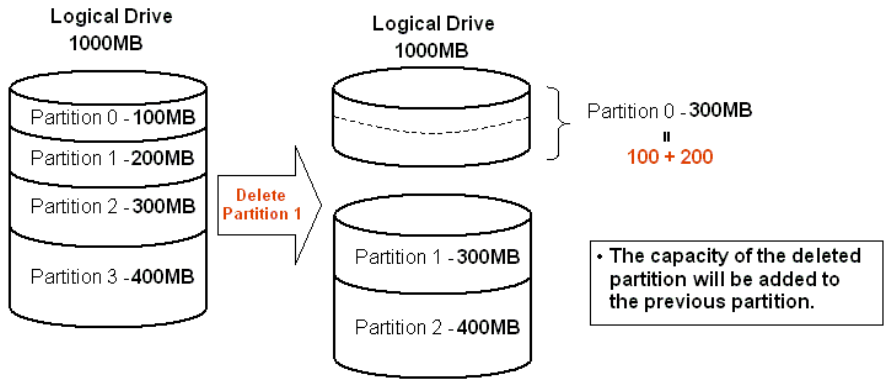

Figure 5-5 Drive Space Allocated to the Previous Partition

5.9 Naming a Logical Drive

Step 1. Press **ENT** for two seconds to enter the Main Menu. Press the up or down arrow keys to select "**View and Edit Logical Drives..**," then press **ENT**.

```
View and Edit
Logical Drives
```

Step 2. Press the up or down arrow keys to select a logical drive, then press **ENT**.

```
LG0 RAID5 DRV=3
  2012MB GD SB=1
```

Step 3. Press the up or down arrow keys to select "**Logical Drive Name**," then press **ENT**.

```
Logical Drive
Name           ..
```

Step 4. Press the up or down arrow keys to change the character of the flashing cursor. Press **ENT** to move the cursor to the next space. The maximum number of characters for a logical drive name is 32.

```
Enter LD Name:
_
```

5.10 Deleting Host LUNs

Step 1. Press **ENT** for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Host Luns", then press **ENT**.

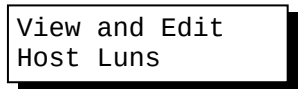

Step 2. Press the up or down arrow keys to select a host ID, then press **ENT** to proceed.

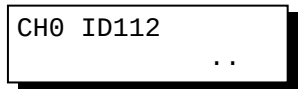

Step 3. Use the up or down arrow keys to browse through the LUN number and its LUN mapping information.

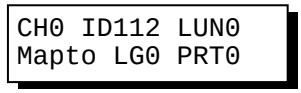

Step 4. Press **ENT** on the LUN you wish to delete.

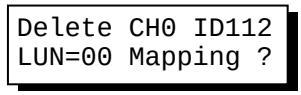

Step 5. Press **ENT** for two seconds to confirm deletion. The deleted LUN has now been unmapped.

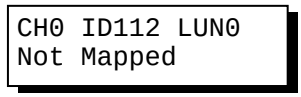

5.11 Deleting Spare Drives

Deleting Spare Drive (Global / Local/Enclosure Spare Drive)

Step 1. Press **ENT** for two seconds to enter the Main Menu. Press the up or down arrow keys to select "View and Edit Drives," then press **ENT**.

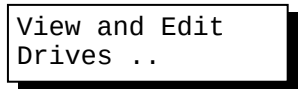

Step 2. Drive information will be displayed on the LCD. Press the up or down arrow keys to select the spare drive you wish to delete, then press **ENT.**

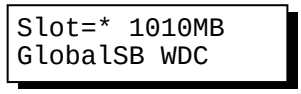

Step 3. Press the up or down arrow keys to select "Delete Spare Drive," then press **ENT** to continue.

```
Delete Spare
Drive          ..
```

Step 4. Press **ENT** for two seconds to delete the spare drive.

```
Delete Spare
Drive Successful
```

5.12 Restore Firmware Default

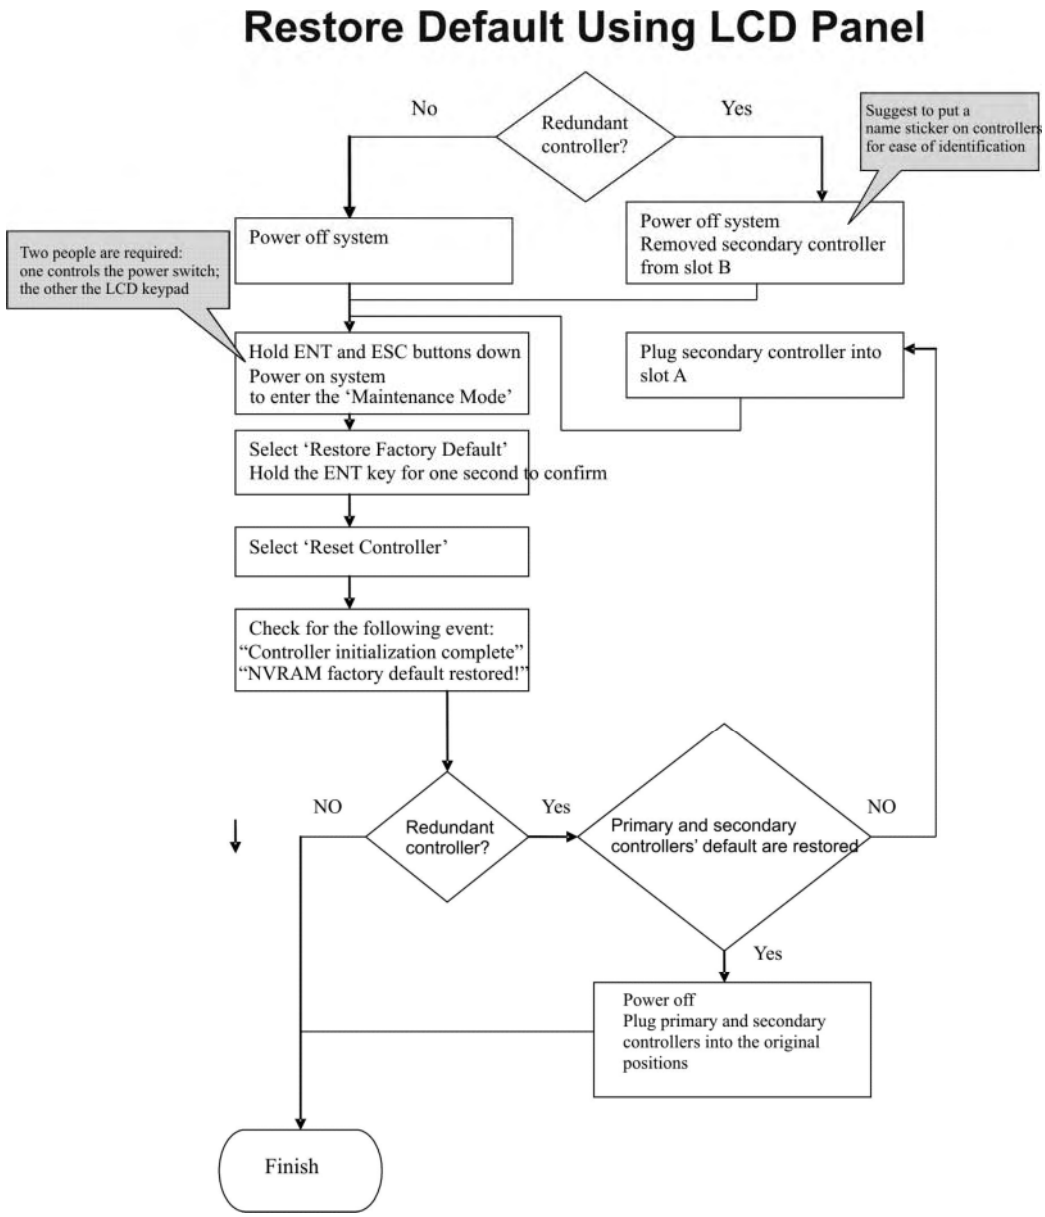

Chapter

6

Creating RAID Arrays & Host LUN Mapping

(Using Terminal Access)

6.1 Initial Screen

Hardware installation should be completed before powering on your RAID enclosure. The subsystem and disk drives must be properly configured and initialized before the host computer can access the storage capacity. The text-based configuration and administration utility resides in the controller's firmware.

Open the initial terminal screen: use the arrow keys to move the cursor bar through the menu items, then press **[ENTER]** to choose the terminal emulation mode, and **[ESC]** to dismiss current selection and/or to return to the previous menu/screen.

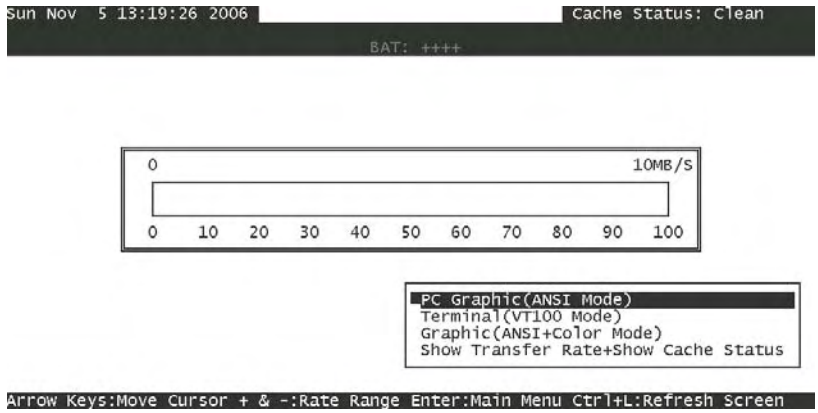

6.2 Individual Disk Menu:

Prior to configuring individual disk drives into a logical drive, it is necessary to understand the status of all physical drives in your enclosure.

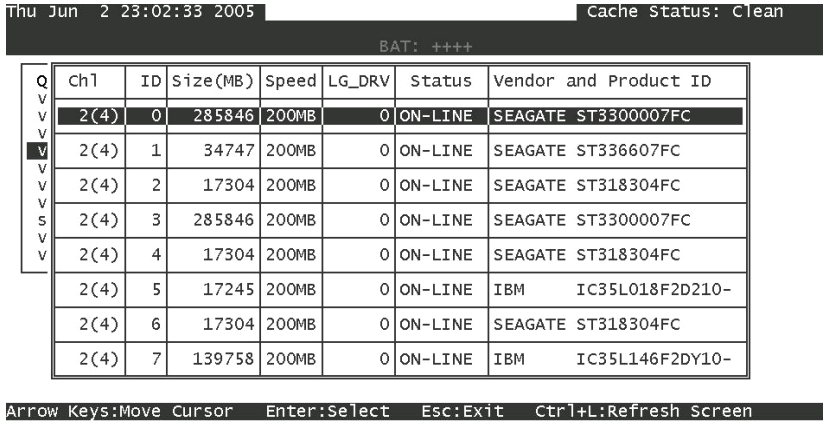

Chl	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
2(4)	0	285846	200MB	0	ON-LINE	SEAGATE ST3300007FC
2(4)	1	34747	200MB	0	ON-LINE	SEAGATE ST336607FC
2(4)	2	17304	200MB	0	ON-LINE	SEAGATE ST318304FC
2(4)	3	285846	200MB	0	ON-LINE	SEAGATE ST3300007FC
2(4)	4	17304	200MB	0	ON-LINE	SEAGATE ST318304FC
2(4)	5	17245	200MB	0	ON-LINE	IBM IC35L018F2D210-
2(4)	6	17304	200MB	0	ON-LINE	SEAGATE ST318304FC
2(4)	7	139758	200MB	0	ON-LINE	IBM IC35L146F2DY10-

Use the arrow keys to scroll down to "**View and Edit Drives**" to display information on all the physical drives installed.

Physical hard drives are listed in the "View and Edit Drives" table. Use the arrow keys to scroll the table. First examine whether there is any drive installed but not listed here. If a disk drive is installed but not listed, the drive may be faulty or not installed correctly. Reinstall the hard drives and contact your supplier for replacement drives.

IMPORTANT!

- Drives of the same brand/model/capacity may not have the same block number.
- The basic read/write unit of a hard drive is block. If members of a logical drive have different block numbers (capacity), the smallest block number will be taken as the maximum capacity to be used in every drive. Therefore, use drives of the same capacity.
- You may assign a Spare Drive to a logical drive whose members have a block number equal or smaller than the Local/Global Spare Drive, but you should not do the reverse.

6.2.1 Drive Read/Write Test

From the "**View and Edit Drives**" menu, select a "**new**" drive that the utility is to be performed on, then press [**ENTER**].

TIPS:

To perform Read/Write test on a "**FRMT DRV** (formatted)" drive, you can manually remove the 256MB reserved space on that drive to change its status into the "**NEW DRV**."

You cannot perform Read/Write test on members of a logical drive.

Select "**Drive Utilities,**" then press [**ENTER**]. Choose "**Read/Write Test**" and press [**ENTER**]. You can choose to enable/disable the following parameters:

- **Ready-Write Test Type:**

  Two test types are available: **Ready-Only** and **Read + Write**

- **Actions to be Taken on Detection of Errors:**
  - **Continue Test**
  - **Abort on Detection of Any Errors** (including hardware, media errors, etc.)
  - **Abort on Detection of Hard Errors only** (including hardware errors defined by SATA 8-bit encoding)
- **Recovery Process:**
  - **No Recovery**
  - **Mark Block Bad** (Mark the affected blocks as bad)
  - **Reassignment** (try to reassign data residing on the affected sectors to other drive sectors)
  - **First Try Reassignment** (Operate by attempting to reassign and mark bad if the reassignment attempt fails)
- **Execute Drive Testing**

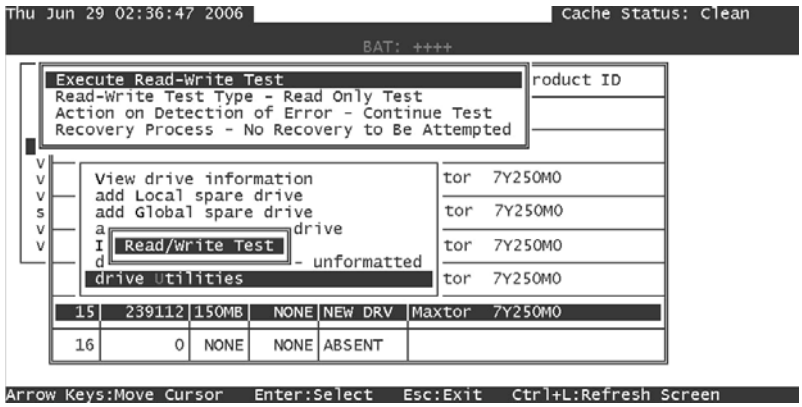

When finished with the configuration, select "**Execute Drive Testing**" and press **[ENTER]** to proceed.

The Read/Write test progress will be indicated by a status bar.

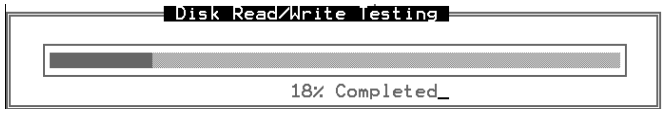

You may press **[ESC]** and select "**Read/Write Test**" later and choose to "**View Read/Write Testing Progress**." If you want to stop testing the drive, select "**Abort Drive Testing"** and press **[ENTER]** to proceed.

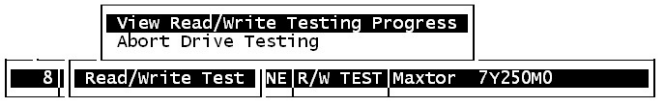

IMPORTANT!

- The disk drive on which a read/write test will be performed cannot be a spare drive (local or global) nor a member of a logical drive.
- The "**Read/Write Test**" option will not appear if the drive is not indicated as a "**New Drive**". Also, a drive formatted with a 256MB reserved space is also excluded from selection.
- The "**Read/Write Test**" function only applies to Fibre Channel, SAS, or SCSI disk drives.

6.3 Creating a Logical Drive

Step 1. Browse through the Main Menu and select "**View and Edit Logical Drive**."

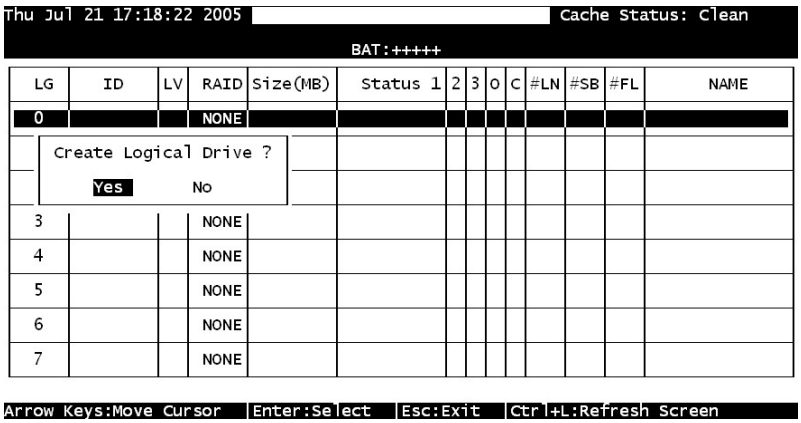

Step 2. For the first logical drive on the RAID subsystem, simply choose the first logical drive entry, "**LG 0**," and press **[ENTER]** to proceed. You may create as many as 16 logical drives or more using drives in a RAID subsystem or in a cascaded enclosure.

Step 3.

When prompted to "**Create Logical Drive?**," select **Yes** and press **[ENTER]** to proceed.

6.3.1 Choosing a RAID Level

Step 4. A pull-down list of supported RAID levels will appear. Choose a RAID level for this logical drive. In this chapter, RAID 6 will be used to demonstrate the configuration process.

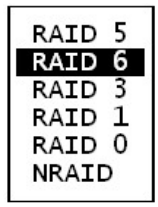

6.3.2 Choosing Member Drives

Step 5. Choose your member drive(s) from the list of available physical drives. Tag the drives for inclusion by positioning the cursor bar on the drive and then pressing **[ENTER]**. An asterisk "*" will appear in front of the selected physical drive(s). To deselect the drive, press **[ENTER]** again on the selected drive and the asterisk "*" will disappear. Use the same method to select more member drives.

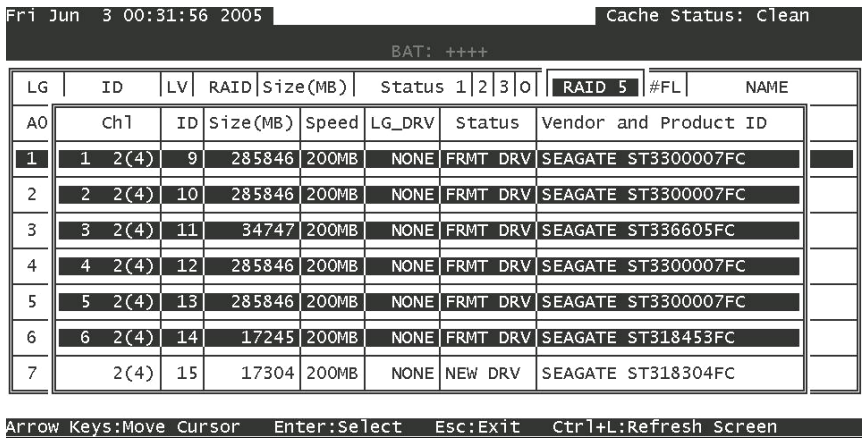

6.3.3 Logical Drive Preferences

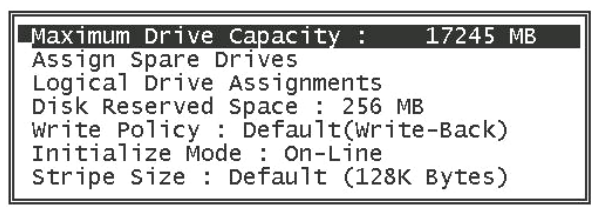

Step 6. After you selected the directed members, press **[ESC]** to proceed. A **Logical Drive Preference** menu will prompt.

6.3.3.1 Maximum Drive Capacity

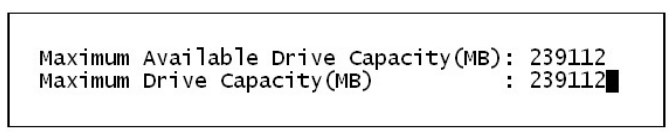

Step 6-1. As a rule, a logical drive should be composed of drives of the same capacity. A logical drive can only

use the capacity of each drive up to the maximum capacity of the smallest member selected for the array. The capacity of the smallest member will be listed here as the maximum drive capacity.

6.3.3.2 Assign Spare Drives

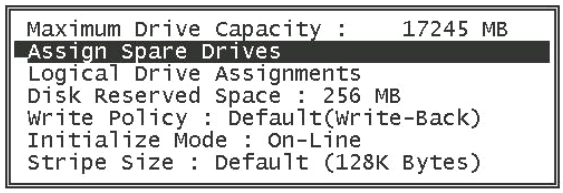

Chl	ID	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
2(4)	15	17304	200MB	NONE	NEW DRV	SEAGATE ST318304FC

Step 6-2. You can assign a "**Local Spare**" drive to the logical drive from a list of unused disk drives. The spare chosen here is a spare exclusively assigned and will automatically replace a failed drive within the logical drive. The controller will then rebuild data onto the replacement drive in the event of disk drive failure.

NOTE:

- A logical drive composed in a non-redundancy RAID level (NRAID or RAID0) has no fault-tolerance and does not support spare drive rebuild.

6.3.3.3 Disk Reserved Space

Step 6-3. The reserved space is a small section of disk space formatted for storing array configuration, RAIDWatch program, and other non-volatile data. This item is for display only - you cannot change the size of the reserved space.

6.3.3.4 Logical Drive Assignments (Dual-Active Controllers)

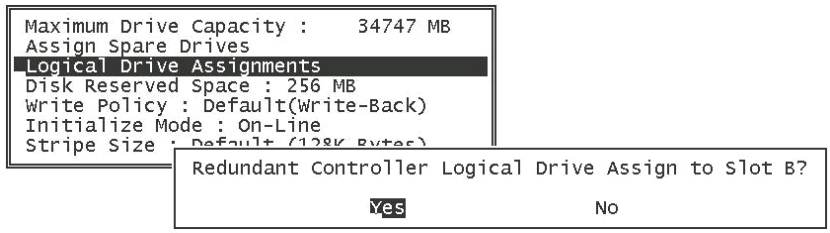

Step 6-4. If your subsystem is powered by dual-redundant RAID controllers, you may manually assign a logical drive to a specific RAID controller. The assignment can take place during or after the initial configuration. The default is the "Slot A" controller; press **[ESC]** if change is not preferred.

IMPORTANT!

Firmware releases before rev. **3.48**:

A logical drive can be assigned to **either** the **Primary** controller **or** the **Secondary** controller.

Firmware releases rev. **3.48** or later:

A logical drive can be assigned to **both** the **Slot A** and **Slot B** controllers. Since the Master/Secondary role can change during a controller failover, RAID controllers are identified as **Slot A** and **Slot B** controllers in firmware rev. 3.48.

Slot A refers to the upper controller slot, which is the default location for the dominating or master controller. **Slot B** refers to the lower controller slot.

6.3.3.5 Write Policy

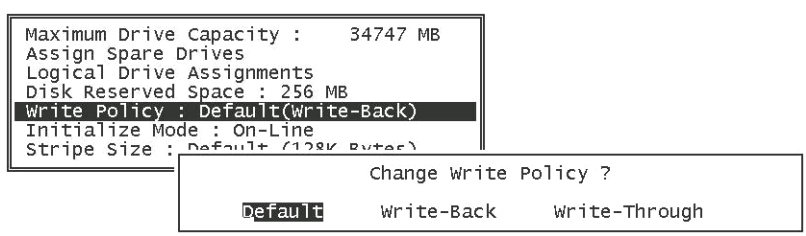

Step 6-5. This sub-menu allows you to select the caching mode for this specific logical drive. "**Default**" is a neutral value that is coordinated with the subsystem's general caching mode setting bracketed in the Write Policy status.

IMPORTANT!

- The "**Write-back**" and "**Write-through**" parameters are permanent for specific logical drives. The "**Default**" selection, however, is more complicated and more likely equal to "not specified."
- If set to "**Default**," a logical drive's write policy is determined not only by the system's general caching mode setting, but also by the "**Event trigger**" mechanisms. The "Event Trigger" mechanisms automatically disable the write-back caching and adopt the conservative "**Write-through**" mode in the event of component failures or elevated temperature.

6.3.3.6 Initialization Mode

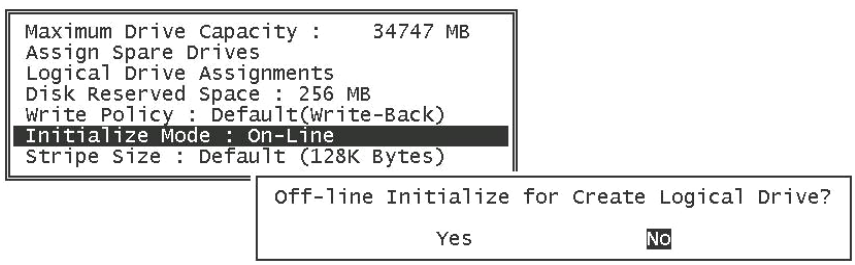

Step 6-6. This sub-menu allows you to configure if the logical drive is immediately available. If the online (default) mode is used, logical drive is immediately ready for I/Os and you may continue with array configuration, e.g., including the array into a logical volume, before the array's initialization process is completed.

6.3.3.7 Stripe Size

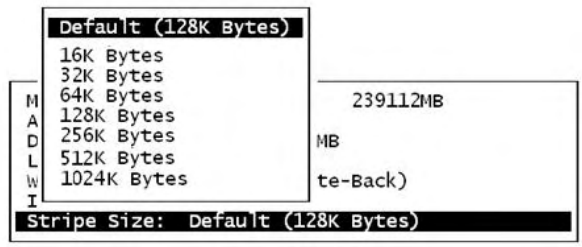

Step 6-7. This option should only be changed by experienced technicians. Setting to an incongruous value can severely drag performance; therefore, this option should only be changed when you can be sure of the performance gains it might bring you.

For example, if your array is often stressed by large and sequential I/Os, a small stripe size will force hard disks to spin many more times in order to conduct data in different data blocks and hence reduce the efficiency brought by parallel executions.

Diagrammed below are conditions featuring host I/Os in 512KB transfer size and a RAID3 array using 128KB and 32KB stripe sizes. The first condition shows a perfect fit where each host I/O is efficiently satisfied by writing to 4 disks simultaneously.

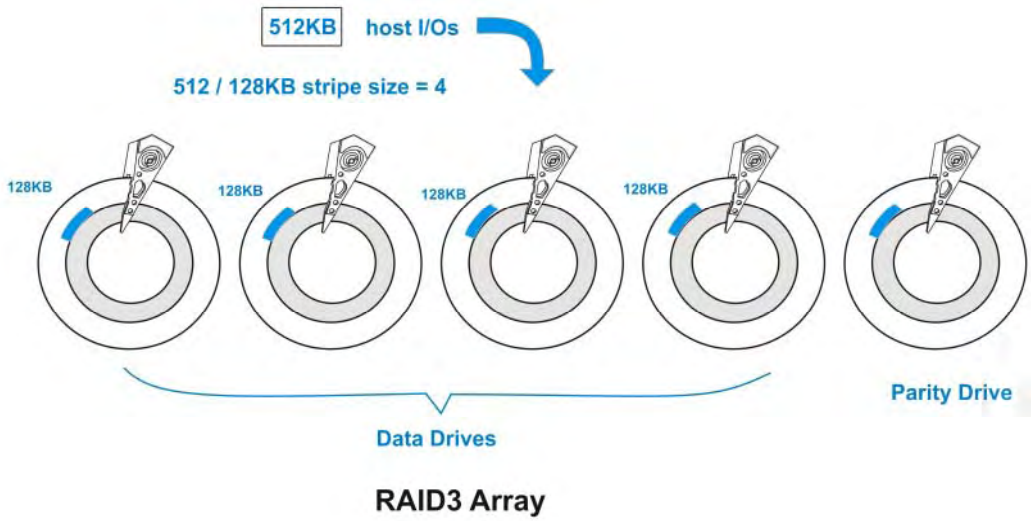

Figure 6-1: 512KB Transfer Size and 128KB Stripe Size

As the contrast, an inadequately small, 32KB stripe size will force the hard disks to write four times and controller firmware to generate 4 parity blocks.

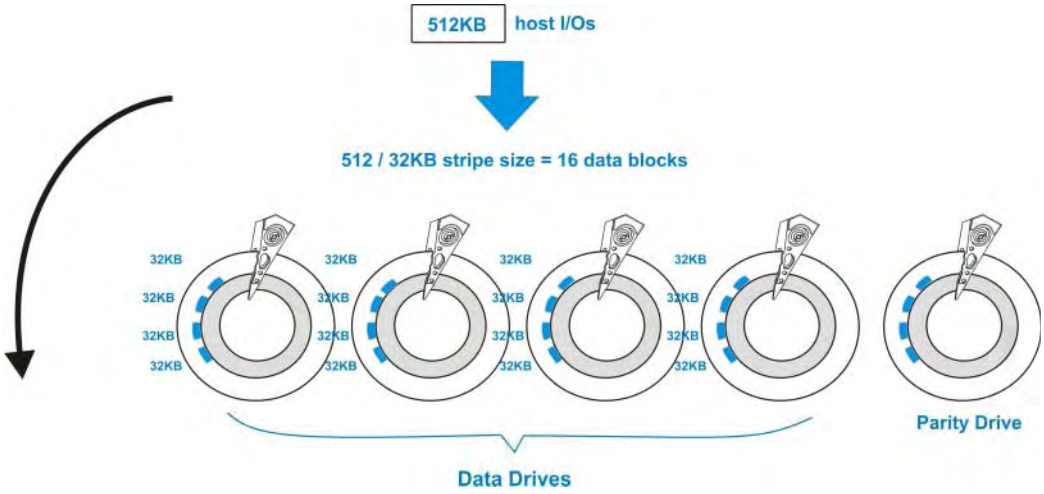

Figure 6-2: 512KB Transfer Size and 32KB Stripe Size

The "**Default**" value is determined by the combined factors of the controller **Optimization Mode** setting and the **RAID level** selected for the specific logical drive. See the table below for default values:

Table 6-1 RAID Levels and Stripe Sizes

RAID Level	Stripe Size
RAID0	128KB
RAID1	128KB
RAID3	16KB
RAID5	128KB
RAID6	128KB
NRAID	128KB

Press **[ESC]** to continue when all the preferences have been set.

NOTE:

The Stripe size here refers to the "**Inner Stripe Size**" specifying the chunk size allocated on each individual data disk for parallel access instead of the "**Outer Stripe Size**" that is the sum of chunks on all data drives.

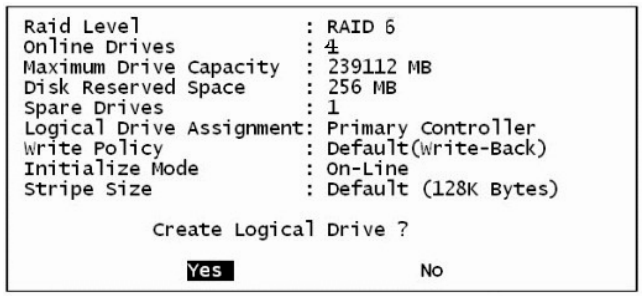

Step 7. A confirmation box will appear on the screen. Verify all information in the box before choosing **Yes** to confirm and proceed.

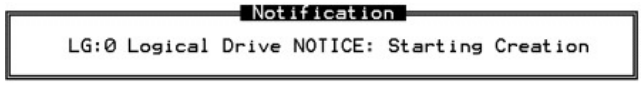

Step 8. If the online initialization mode is applied, the logical drive will first be created and the controller will initialize the array in the background or when the array is less stressed by I/Os.

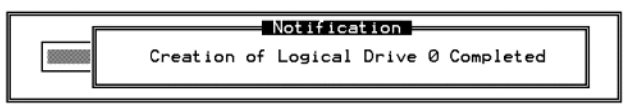

Step 9. The completion of array creation is indicated by the message prompt above.

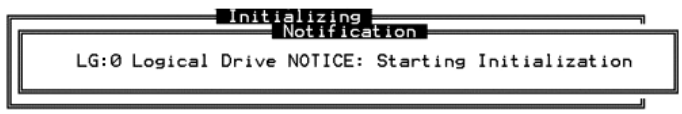

Step 10. A controller event will prompt to indicate that the logical drive initialization has begun. Press **ESC** to cancel the "Notification" prompt, and a progress indicator will display on the screen as a percentage bar.

Step 11. While the array initialization runs in the background, you can continue configuring your RAID subsystem, e.g., with host LUN mapping.

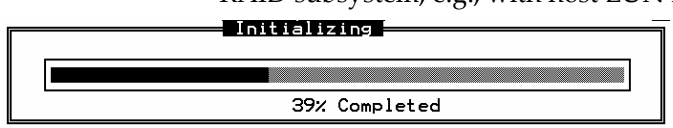

When a fault-tolerant RAID level (RAID 1, 3, 5 or 6) is selected, the subsystem will start initializing parity.

Step 12. Use the **ESC** key to view the status of the created logical drive.

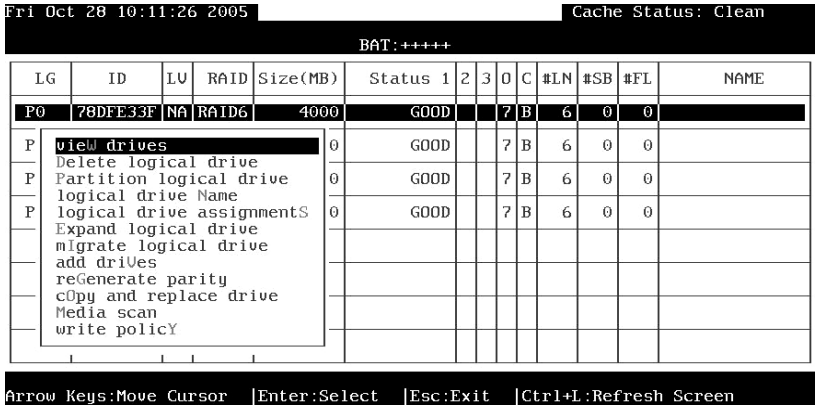

IMPORTANT!

Only logical drives with RAID levels 1, 3, 5, or 6 will take the time to initialize the logical drive. Logical drives with RAID level 0 and NRAID do not perform logical drive initialization. With RAID0 or NRAID, the drive initialization process finishes almost immediately.

6.4 RAID Migration

Currently the RAID migration function supports the migration between RAID5 and RAID6.

Before proceeding with RAID migration, make sure you have sufficient free capacity or unused disk drives in your RAID array. RAID6 arrays require at least four (4) member drives and use additional capacity for the distribution of secondary parity. For example, if you want to migrate a RAID5 array consisting of three (3) drives to RAID6, one additional disk drive should be available.

Different features of RAID5 and RAID6 arrays are summarized as follows:

Table 6-2 RAID5 and RAID6 Comparison

	RAID5	RAID6
Min. No. of Member Drives	3	4
Usable Capacity	N-1 (1 drive's capacity used for storing parity data)	N-2 (2 drives' capacity used for storing parity data); N>=4
	If individual disk capacity = 100G, Capacity of a 4-drive RAID5 = (4 -1) x 100G = 300G Capacity of a 4-drive RAID6 = (4 -2) x 100G = 200G	
Redundancy	Single disk drive failure	2 disk drives to fail at the same time

6.4.1 Requirements for Migrating a RAID5 Array

The precondition for migrating a RAID5 array to RAID6 is:

- The "**usable capacity**" of the RAID6 array should be equal or larger than the "usable capacity" of the original RAID5 array.

To obtain a larger capacity for migrating to RAID6, you can:

1). **Add Drive(s)**: Include one or more disk drives into the array.

2). **Copy and Replace**: Use larger disk drives in the array to replace the original members of the RAID5 array.

6.4.2 Migration Methods

The conditions for migrating a RAID5 array to RAID6 are diagrammed as follows:

Fault condition:

The usable capacity of the RAID6 array is smaller than the usable capacity of the original RAID5 array.

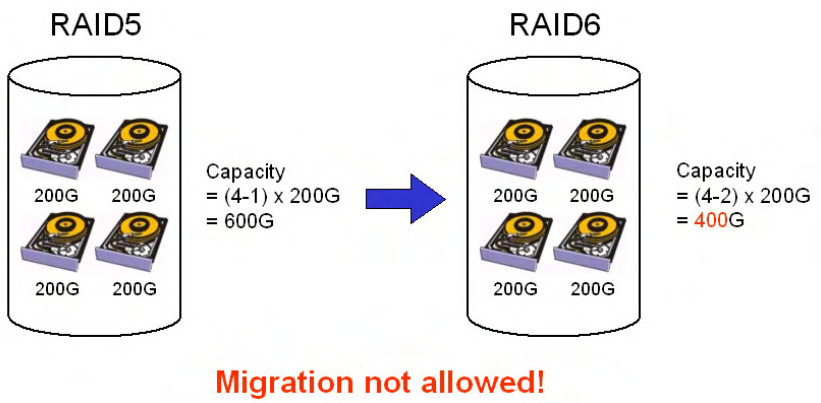

Figure 6-3: Migration Not Allowed for Insufficient Capacity

Migration by Adding Drive(s):

The additional capacity for composing a RAID6 array is acquired by adding a new member drive.

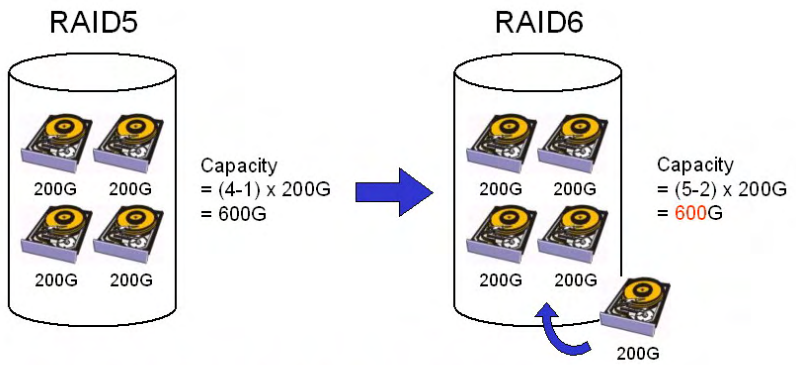

Figure 6-4: Migration Condition Met by Adding Drive(s)

Migration by Copy and Replace:

The additional capacity for composing a RAID6 array is acquired by using larger drives as the members of the array.

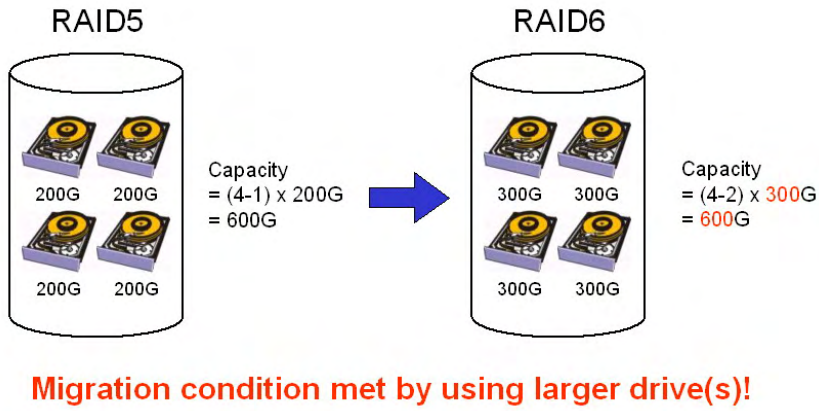

Figure 6-5: Migration Condition Met by Using Larger Drives

6.4.3 Migration: Exemplary Procedure

To migrate a RAID5 array to RAID6, follow the steps below:

Step 1. From the "**View and Edit Logical Drives**" status screen, select an existing RAID5 array and press **[ENTER]**. A sub-menu will appear. Move your cursor bar to select **Migrate Logical Drive.**

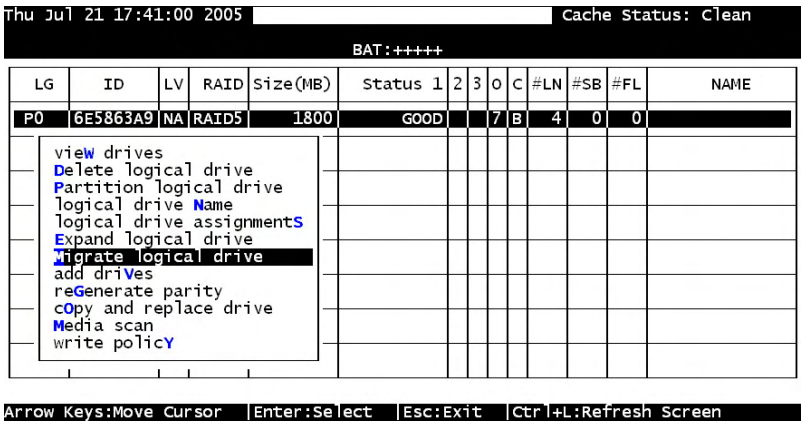

Step 2. A confirmation box should appear. Select **Yes** to proceed.

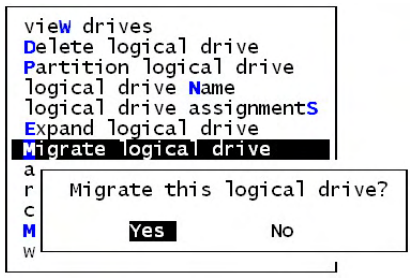

Step 3. A selection box should prompt allowing you to choose a RAID level to migrate to. Press **[ENTER]** on RAID6.

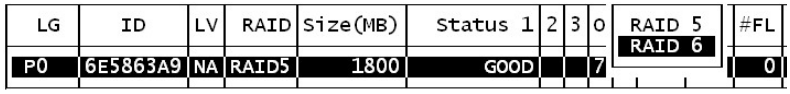

Step 4. A list of member drives and unused disk drives (new or used drives) should prompt. In the case of migrating a 3-drive RAID5 to 4-drive RAID6, you can select the original members of the RAID5 array and select one more disk drive to meet the minimum requirements of RAID6. You may also select different disk drives in your RAID enclosure for composing the new RAID6 array.

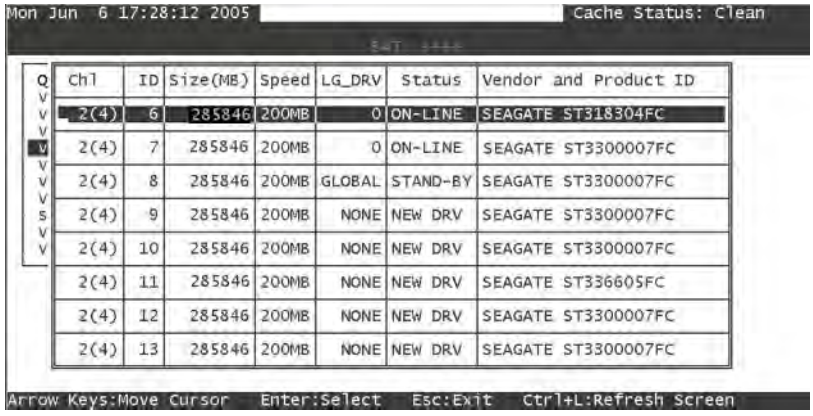

Step 5. Press **[ESC]** to proceed to the next configuration screen. A sub-menu should prompt.

Step 6. You may either change the maximum capacity to be included in the new RAID6 array or change the array stripe size.

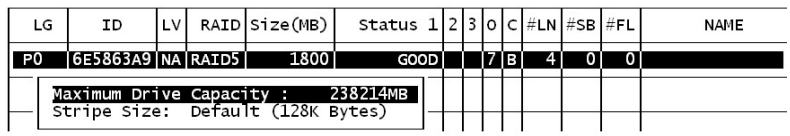

Step 7. A confirmation box should prompt. Check the configuration details and select **Yes** to start the migration process.

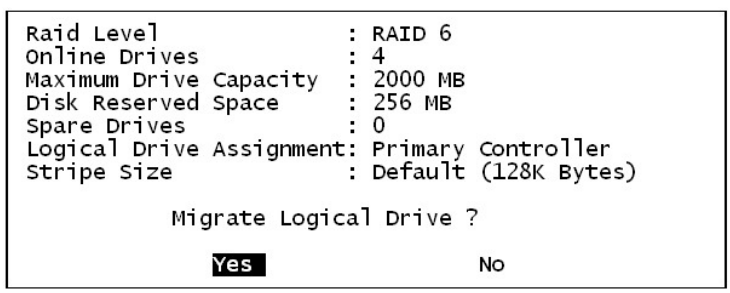

Step 8. A message should prompt indicating the migration process has started.

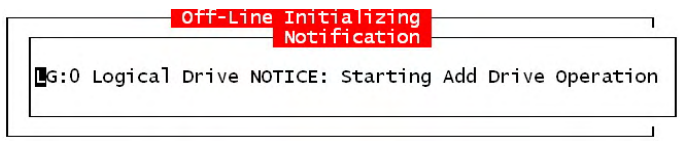

Step 9. Press **[ESC]** to clear the message. The initialization progress is shown below.

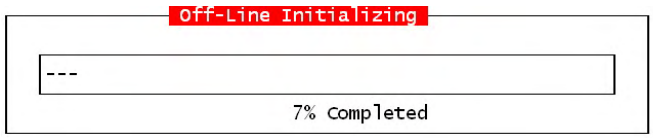

Step 10. Since the migration process includes adding a new member drive, the completion of RAID migration is indicated as follows:

Once the migration is completed, associate the RAID6 array with the ID/LUN number originally associated with the previous RAID5 array.

Please refer to ***Section 6.7 Mapping an Array Volume to Host LUNs*** for more information.

6.5 Creating a Logical Volume

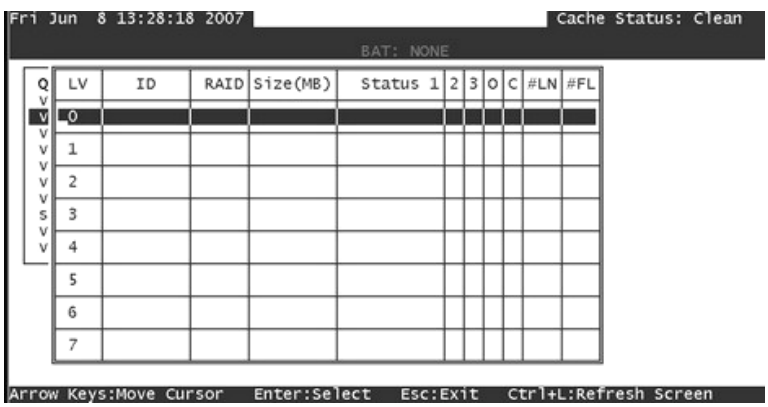

NOTE:

- A logical volume consists of one or several logical drives. These logical drives are striped together.
- For example, A RAID50 volume is one that consists of several RAID5 arrays using the block-level striping method known as RAID0. A RAID50 logical volume is actually a RAID0 volume striped across RAID5 arrays.

Step 1. Select "**View and Edit Logical Volumes**" in the Main Menu to display the current logical volume configuration and status on the screen. Select a logical volume index number (0 to 7) that has not yet been defined, and then press **[ENTER]** to proceed.

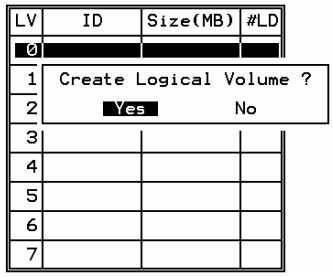

Step 2. A prompt "**Create Logical Volume?**" will appear. Select **Yes** and press [**ENTER**].

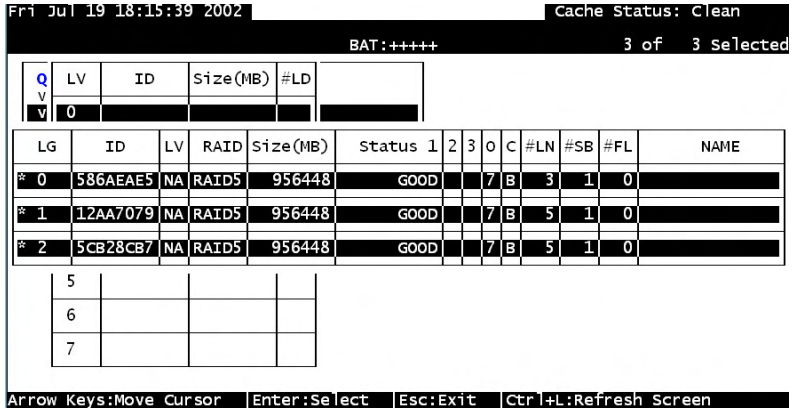

Step 3. Select one or more logical drive(s) available on the list. The same as creating a logical drive, the logical drive(s) can be tagged for inclusion by positioning the cursor bar on the desired disk drive and pressing **[ENTER]** to select. An asterisk (*) will appear on the selected logical drive. Pressing **[ENTER]** again will deselect a logical drive.

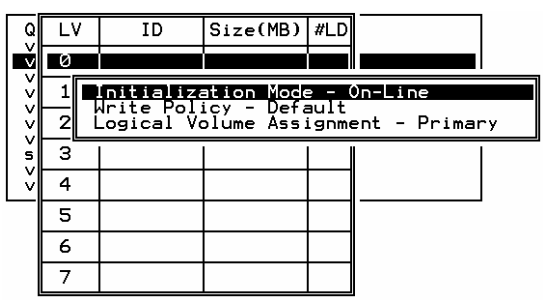

Step 4. Use the arrow keys to select a sub-menu and change the initialization mode, write policy, or the managing controller.

Step 5. Logical volumes can be assigned to different controllers (**primary** or **secondary**; **Slot A** or **Slot B** controllers). The default is the primary or Slot A controller.

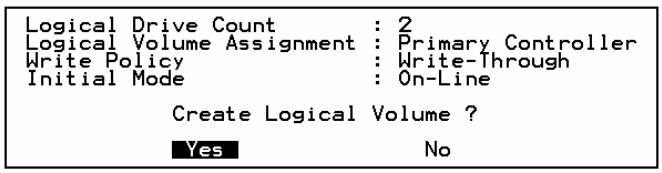

Note that if a logical volume is manually assigned to a specific controller, all its members' assignments will also be shifted to that controller.

Step 6. When all the member logical drives have been selected, press **[ESC]** to continue. The confirm box displays. Choose **Yes** to create the logical volume.

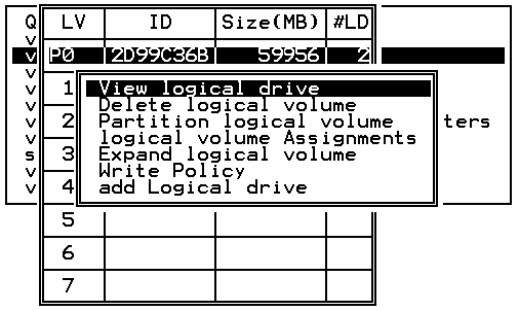

Step 7. Press **[ENTER]** on a configured volume, and the information of the created logical volume displays.

LV: Logical Volume ID

ID: Unique ID for the logical volume, randomly generated by the RAID controller firmware

Size: Capacity of this volume

#LD: Number of the included members

6.6 Partitioning RAID Arrays

WARNING!

- Whenever a partition is changed, it is necessary to reconfigure all host LUN mappings. All data in it will be lost and all host LUN mappings will be removed when there is any change to the partition capacity.
- If operating in a Unix-based system, reset the subsystem for the configuration changes to take effect if any changes were made to partition sizes and partition arrangement.

The process of partitioning a logical drive is the same as that of partitioning a logical volume. The partitioning of a logical volume is used as an example in the proceeding discussion.

NOTE:

- Partitioning can be useful when dealing with complex storage; e.g., dividing large logical configurations. However, partitioning a logical drive or logical volume is not a requirement for RAID configuration.

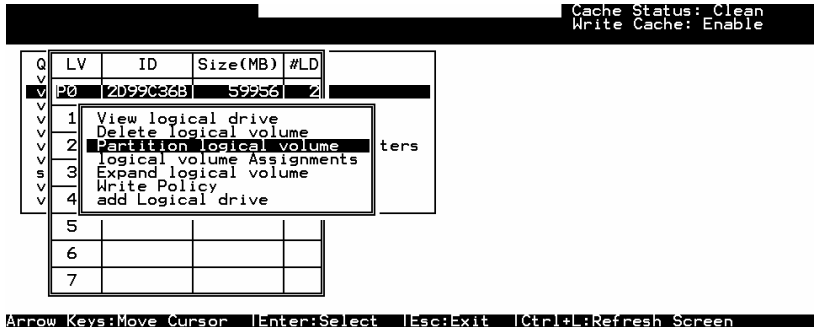

Step 1. Select the logical volume you wish to partition, then press **[ENTER]**. Choose **Partition logical volume**, then press **[ENTER]**. Select from the list of undefined partitions and press **[ENTER]**.

Step 2. A partition list displays. If the logical volume has not yet been partitioned, all volume capacity will list as "**partition 0**."

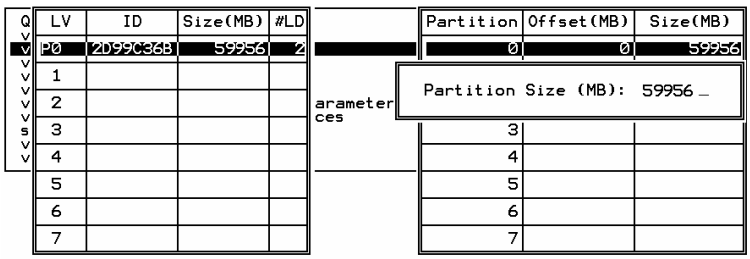

Step 3. Press [**ENTER**] and key in the desired capacity for the selected partition, and then press [**ENTER**] to proceed. The remaining capacity will be automatically allotted to the next partition.

Step 4. When prompted by the "**Partition Logical Volume?**" message, choose **Yes** to confirm then press **[ENTER]**. Follow the same procedure to partition the remaining capacity of your logical volume.

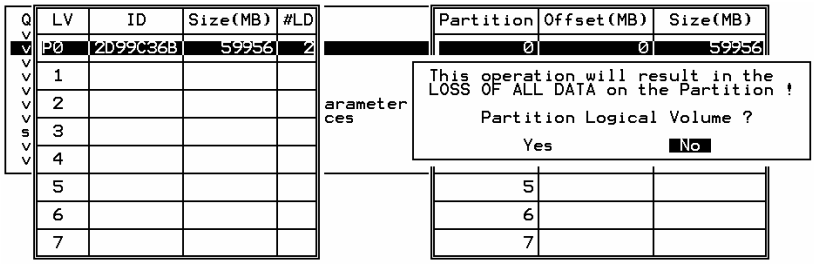

When a partition of a logical drive/logical volume is deleted, the capacity of the deleted partition will be added to the previous partition.

6.7 Mapping an Array Volume to Host LUNs

Step 1. Select "**View and Edit Host luns**" in the Main Menu, then press **[ENTER]**.

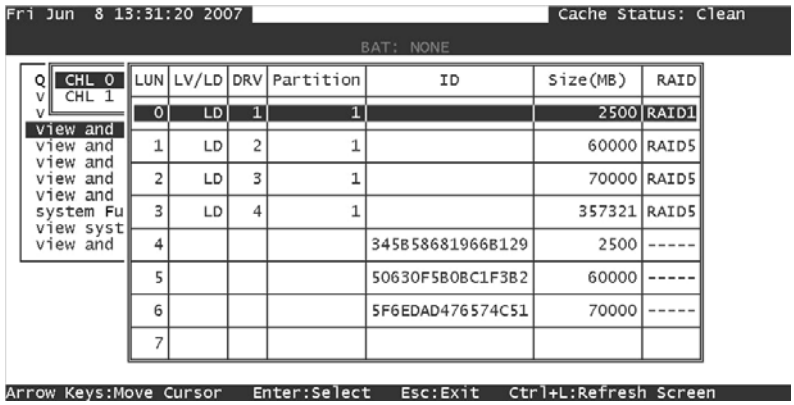

IMPORTANT!

Note the following details before proceeding:

- **Slot A or Slot B IDs**: If using firmware before rev.3.48, logical group of drives (logical drive/logical volume) previously assigned to the Slot A controller cannot be re-assigned to a Slot B controller. Neither can those assigned to the Slot B controller be mapped to a Slot A controller ID.

- Your subsystem comes with Slot A IDs only. You need to manually create Slot B IDs in a dual-controller configuration. Please enter "**View and Edit Channels**" menu to create or remove a host ID.

CAUTION!

Firmware revisions 3.48 and later also support the **cross-controller** ID mapping. The cross-controller mapping allows you to associate a logical drive with BOTH controller A and controller B IDs. However, mapping to both controllers' IDs is only beneficial when it is difficult making the fault-tolerant host links between RAID controllers and host HBAs, e.g., using SAS-to-SAS RAID systems. Currently, there is no external SAS switch available on the market. For Fibre-host systems, fault-tolerant links can easily be made with the help of external bypass such as Fibre Channel switches.

For details of fault-tolerant link connections, please refer to your system Hardware Manual.

Step 2. A list of host channel/ID combinations appears on the screen. The diagram above shows two host channels and each is designated with at least a default ID. More can be manually added on each channel.

Multiple IDs on host channels are necessary for creating access to RAID arrays through fault-tolerant data links. Details on creating multiple IDs and changing channel modes have been shown in the previous chapter. Select a host ID by pressing **[ENTER]**. The idea of host LUN mapping is diagrammed below:

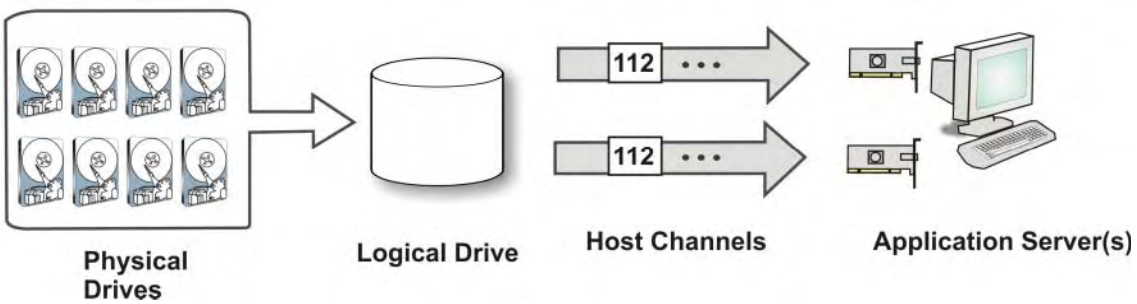

Figure 6-6: Host LUN Mapping

Step 3. Select the channel-ID combination you wish to map, then press **[ENTER]** to proceed. An index of LUN numbers will display. Select an LUN number under the ID. Press **[ENTER]** on an LUN number to proceed and press **[ENTER]** again on **"Map Host LUN"** to proceed.

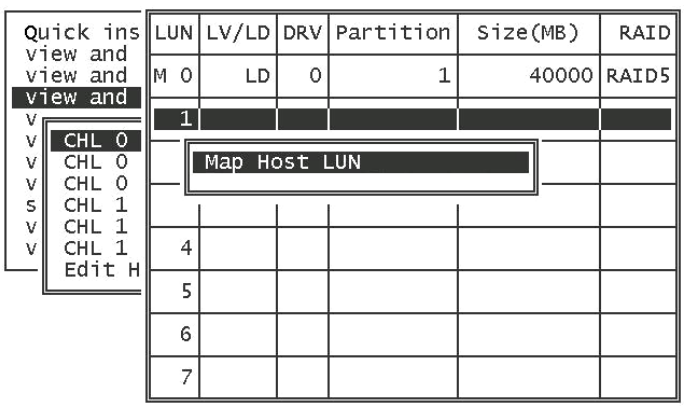

NOTE:

- If your host adapter cards do not support multiple LUN numbers under a channel ID, select **LUN0**. You should refer to the documentation that came with your host adapters to see whether multiple LUNs are an option.

Step 4. Choose mapping either a **Logical Drive** or a **Logical Volume** on the drop box.

Step 5. Existing logical arrays will be listed. Select the array you wish to associate with the target ID by pressing **[ENTER]**.

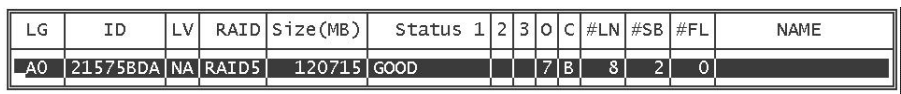

LG	ID	LV	RAID	Size(MB)	Status 1	2	3	O	C	#LN	#SB	#FL	NAME
A0	21575BDA	NA	RAID5	120715	GOOD			7	B	8	2	0	

Step 6. A partition list will prompt. If your array has been divided into many partitions, select one of them to proceed. If not, the whole capacity will be displayed as one logical partition. Press **[ENTER]** on a partition.

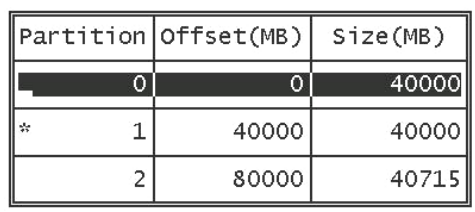

Partition	Offset(MB)	Size(MB)
0	0	40000
* 1	40000	40000
2	80000	40715

Step 7. When prompted by the confirmation message, check the mapping details and select **Yes** to complete the process.

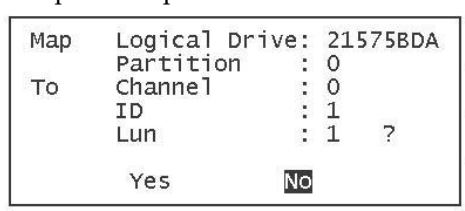

The details in the confirmation box read: partition 0 of logical drive "xxxxxxBDA" will map to (be associated with) LUN 1 of ID 1 on host channel 0.

You may now continue to map other arrays (partitions) to host LUNs.

NOTE:

Once any host ID/LUN is successfully associated with a logical capacity, the "**No Host LUN**" message in the LCD screen will change to "**Ready**."

6.8 Assigning Spare Drive and Rebuild Settings

6.8.1 Adding a Local Spare Drive

A spare drive is a standby drive that automatically participates in the rebuild of logical arrays. A spare drive must have an equal or larger capacity than the array members. A **Local Spare** is one that participate in the rebuild of a logical drive it is assigned to. A **Global Spare** participates in the rebuild of all configured logical drives, and it should have a capacity equal to or larger than all physical drives in a RAID subsystem.

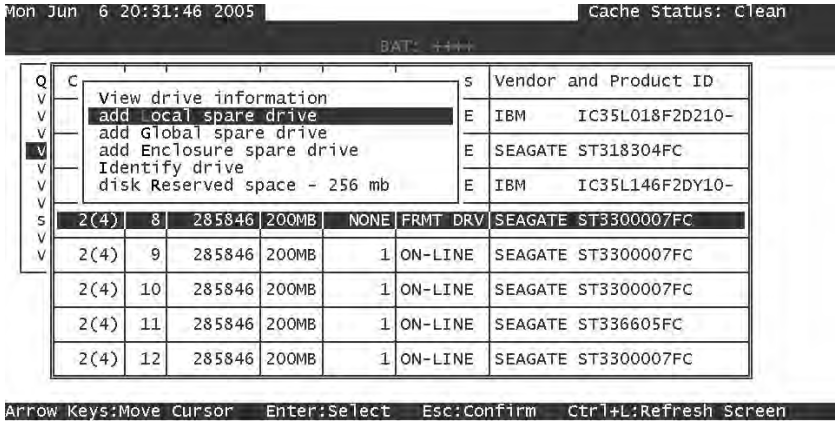

Step 1. Select "**View and Edit Drives**" on the Main Menu, then press **[ENTER]**. Move the cursor bar to a drive that is not included in a logical drive or as a spare drive (usually indicated as a "**New Drive**"), and then press **[ENTER]**.

Step 2. Select "**Add Local Spare Drive**" and press **[ENTER]**. A list of existing logical drives displays.

Step 3. Move the cursor bar to a logical drive, then press **[ENTER]**. The unassigned disk drive will be associated with this logical drive as a Local Spare.

Step 4. When prompted to "**Add Local Spare Drive**?," choose **Yes** to confirm.

6.8.2 Adding a Global Spare Drive

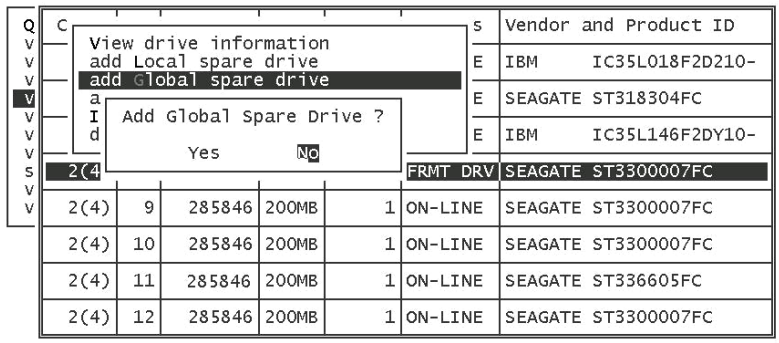

Step 1. Move the cursor bar to the drive that is not a member drive or a spare (usually indicated as a "**New Drive**"), and then press **[ENTER]**.

Step 2. Select "**Add Global Spare Drive**." When prompted by "Add Global Spare Drive?," select ***Yes*** and press **[ENTER]** to complete the configuration process.

6.8.3 Adding an Enclosure Spare Drive

An **Enclosure Spare** only participates in the rebuild of a failed drive located within the same enclosure.

NOTE:

An Enclosure Spare is one that is used to rebuild a failed drive that resides in the same enclosure.

In configurations that span across multiple enclosures, a Global Spare may participate in the rebuild of a failed drive in a different enclosure. Using Enclosure Spare can avoid disorderly locations of member drives in a multi-enclosure configuration.

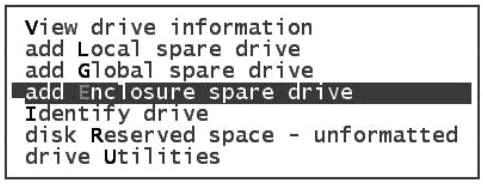

Step 1. Move the cursor bar to a disk drive that is not a member drive or a configured spare (usually indicated as a "New Drive"), and then press **[ENTER]**.

Step 2. Select "**Add Enclosure Spare Drive**." When prompted to "Add Enclosure Spare Drive?," select **Yes** to complete the process.

6.9 Logical Drive and Drive Member Related Functions

Select "**View and Edit Logical Drives**" in the Main Menu to display the array status. Refer to the previous chapter for more details on the legends used in the Logical Drive's status. To see the drive member information, choose the logical drive by pressing **[ENTER]**.

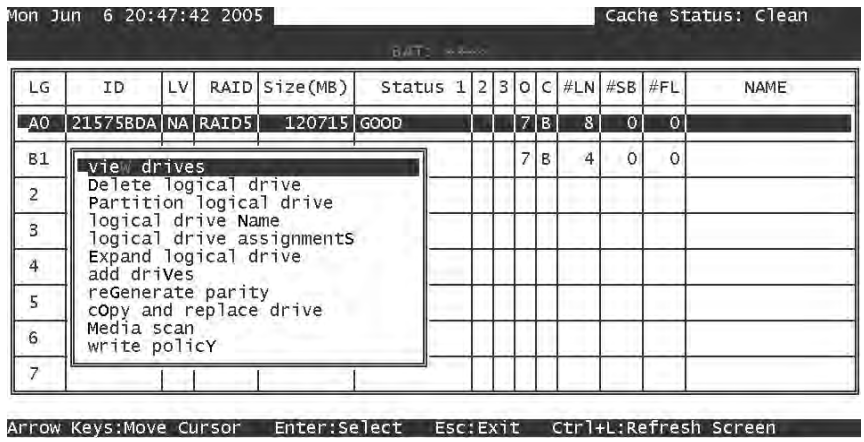

The logical drive-related functions include:

	Function	Description
1	**View Drive**	Displays member drive information
2	**Delete Logical Drive**	Deletes a logical drive
3	**Partition Logical Drive**	Creates or removes one or more partition within a logical drive
4	**Logical Drive Name**	Assigns a name to a logical drive
5	**Logical Drive Assignments**	Assigns a logical drive to the Primary or Secondary RAID controller
6	**Rebuild Logical Drive**	Manually rebuilds a logical drive when a failed drive is replaced
7	**Expand Logical Drive**	Expands the unused capacity of a logical drive
8	**Migrate Logical Drive**	Migrates a logical drive to a different RAID level
9	**Add Drives**	Adds physical drive(s) to a logical drive
10	**Regenerate Parity**	Regenerates a logical drive's parity
11	**Copy and Replace Drive**	Copies or replaces members of a logical drive

12	**Media Scan**	Configures Media Scan priority, iteration count, and task schedules
13	**Write Policy**	Changes the write policy associated with the logical drive

NOTE:

- The **Rebuild** and **Regenerate Parity** processes will be discussed in **Array integrity** chapter.
- The **Logical Drive Assignments** process is complex and will be discussed in the **Redundant Controller** chapter where you can find all details about configuring a dual-active RAID controller configuration.
- The **Add Drive** and **Copy + Replace** processes will be discussed in the **Array Expansion** chapter.

6.9.1 Deleting a Logical Drive

WANRING!

Deleting a logical drive destroys all data stored on it.

NOTE:

Unmap the logical drive from all configured ID/LUNs before deleting it.

Step 1. Select the logical drive you wish to delete, then press **[ENTER]**.

Step 2. Select "**Delete logical drive**," then choose **Yes** when prompted to confirm.

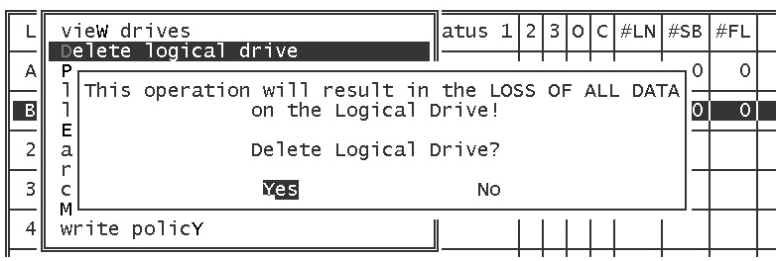

6.9.2 Deleting a Partition of a Logical Drive

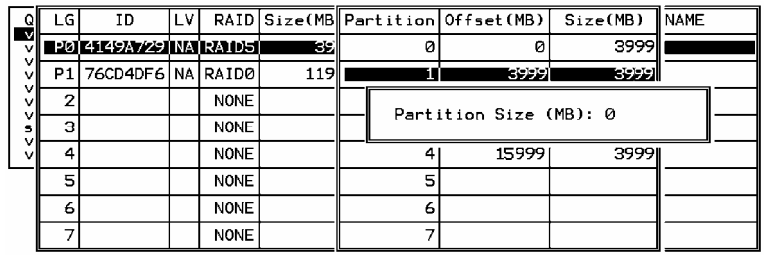

Step 1. Choose the logical drive which has a partition you wish to delete, then press **[ENTER]**.

Step 2. Choose "**Partition logical drive**." Partitions of the logical drive will be displayed in tabulated form.

Step 3. Move the cursor bar to the partition you wish to delete, then press **[ENTER]**. Enter "0" on the partition size to delete the partition.

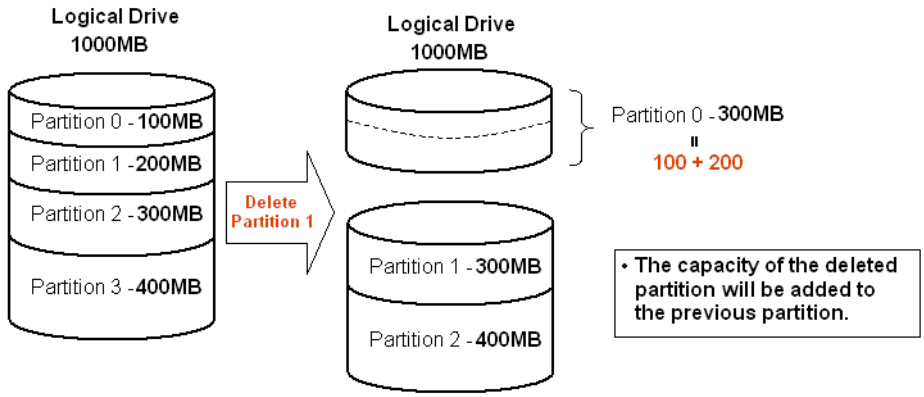

Figure 6-7: Drive Space Allocated to the Previous Partition

As illustrated above, the capacity of the deleted partition will be added to the previous partition. The number tags of partitions following the deleted partition will also be changed, e.g., partition 2 becomes partition 1.

WARNING!

Whenever a partition is changed, it is necessary to reconfigure the associated host LUN mappings. All data kept in the related partitions and the host LUN mappings will be invalidated with partition change.

6.9.3 Naming a Logical Drive

Naming can help identify different arrays in a multi-array configuration.

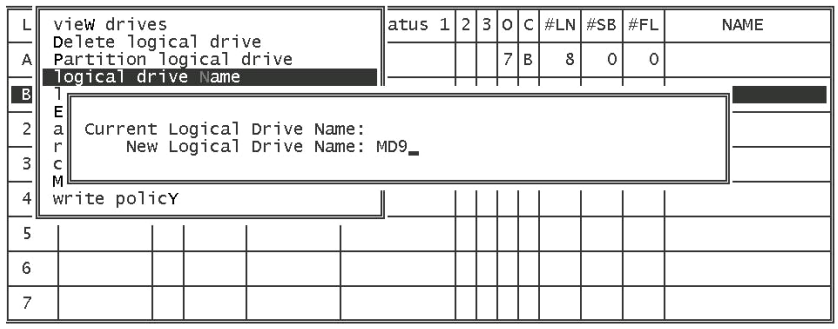

TIPS:

This function is especially helpful in situations such as the following:

One or more logical drives have been deleted, the array indexing is changed after system reboot, e.g., LD0 deleted and the succeeding LD1 becomes LD0. The designating numbers of logical drives following a deleted configuration will all be affected.

Step 1. Choose the logical drive for which you wish to assign a name, then press **[ENTER]**.

Step 2. Choose "**logical drive name**," then press **[ENTER]** again. The current name will be displayed. You may now enter a new name in this field.

Step 3. Enter a name, then press **[ENTER]** to save the configuration. The maximum number of characters for a logical drive name is 32.

6.9.4 Expand Logical Drive

WARNING!

- If the logical drive to be expanded has already been partitioned to the maximum number allowed, e.g., 64 partitions, then the expansion capacity will be added to the last partition. The partition change will then invalidate the data previously stored in the array.

If there is an amount of unused capacity in a logical drive, the logical drive can be expanded. If there is no unused capacity, a logical drive can still be expanded using the "**Add Drive**" or "**Copy & Replace**" technologies.

To expand a logical drive, first select the "**Expand Logical Drive"** function. Press **[Enter]** to activate the configuration window.

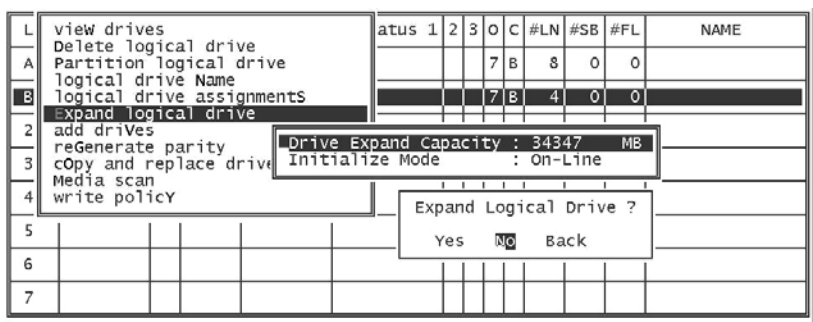

IMPORTANT!

1. The **Drive Expand Capacity** here refers to the unused capacity on each member drive. If a RAID5 array has 4 members and each member drive features a 2GB unused capacity, then the total unused capacity will be 4 - 1 (parity drive) x 2G = 6GB.

2. The capacity brought by the array expansion process will be available as a "new" partition.

3. Chances are if you have partitioned you array into the maximum number of logical partitions and you commenced the "Expand" command, the last partition, e.g. no. 63, will be expanded and will destroy the data in it.

6.10 Deleting Host LUNs

Step 1. Choose the host channel and host ID combination you wish to view or delete.

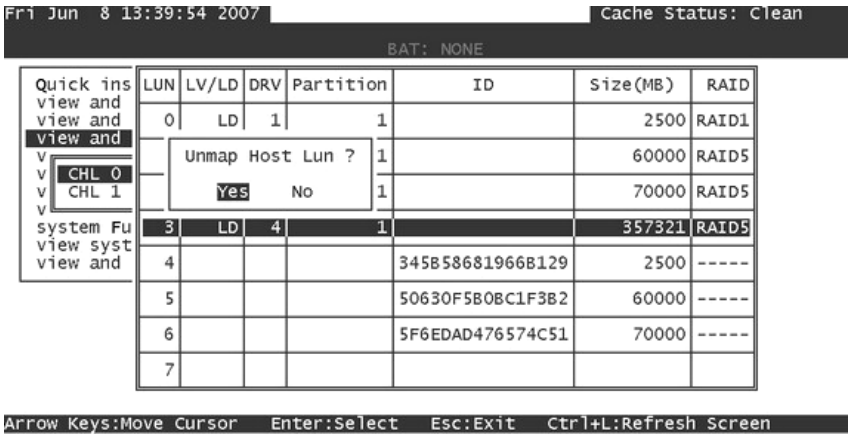

Step 2. A list of current LUN mappings will be displayed on the screen. Move the cursor bar to the LUN mapping you wish to delete, then press **[ENTER].**

Step 3. Select **Yes** to delete the LUN mapping, or **No** to cancel.

6.11 Deleting Spare Drive (Global/Local/Enclosure Spare Drive)

Move the cursor bar to a "Local", "Global," or "Enclosure" Spare Drive in the "**View and Edit Drives**" menu, then press **[ENTER].** Choose "**Delete Global/Local/Enclosure Spare Drive**," then press **[ENTER]** again. Choose **Yes** to confirm.

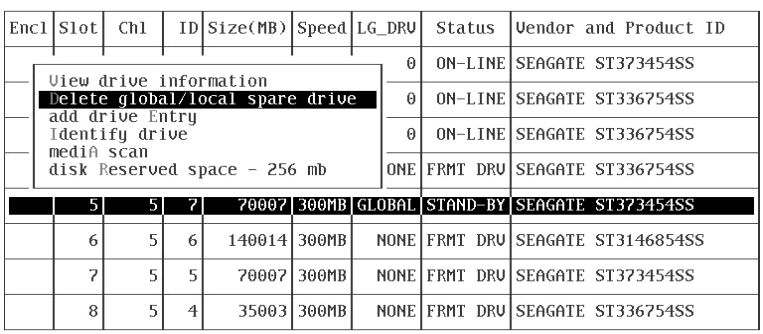

NOTE:

The spare drive you deleted (disassociated or reassigned as a normal disk drive) or any drive you replaced from a logical unit will be indicated as a "used drive."

Chapter

7

Fibre Channel Options

7.1 View and Edit Channels

Most EonStor subsystems come with preset data paths and there is no need to modify channel configurations, e.g., channel mode.

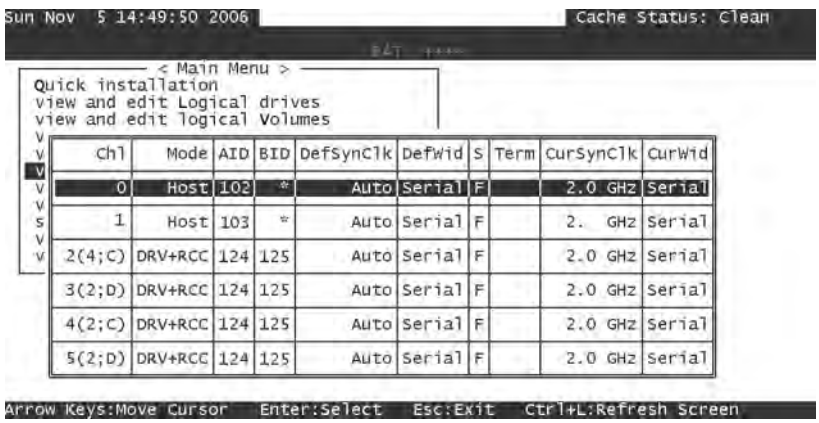

Chl	Mode	AID	BID	DefSyncClk	DefWid	S	Term	CurSyncClk	CurWid
0	Host	102	*	Auto	Serial	F		2.0 GHz	Serial
1	Host	103	*	Auto	Serial	F		2. GHz	Serial
2(4;C)	DRV+RCC	124	125	Auto	Serial	F		2.0 GHz	Serial
3(2;D)	DRV+RCC	124	125	Auto	Serial	F		2.0 GHz	Serial
4(2;C)	DRV+RCC	124	125	Auto	Serial	F		2.0 GHz	Serial
5(2;D)	DRV+RCC	124	125	Auto	Serial	F		2.0 GHz	Serial

Choose "**View and Edit Channels**" in the Main Menu to display information of your I/O channels.

NOTE:

The EonStor ASIC400 models come with dedicated I/O channels that are strung betweens partner RAID controllers. These channels have no external interfaces and cannot be used for I/Os. Information about these dedicated RCC channels can be found on the status menu but no configuration options will be available for them.

7.1.1 Channel IDs - Host Channel

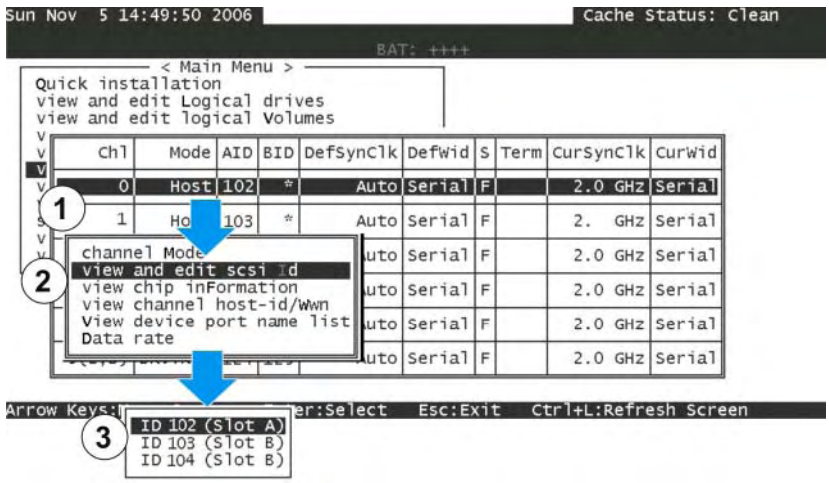

Step 1. Choose a host channel, then press **[ENTER]**.

Step 2. Choose "**View and Edit ID**." A list of existing ID(s) will be displayed on the screen.

Step 3. Select one of the existing IDs and press **[ENTER]**. You may then create a new ID or delete an existing ID.

NOTE:

If a host channel connection is configured in an arbitrated FC loop, FC-AL mode, the maximum number of host IDs will be limited to "16."

7.1.2 Adding an ID (Slot A / Slot B Controller ID)

In a single-controller mode, the Slot B controller ID is unavailable. In a dual-controller configuration, you should manually create one or more Slot B controller IDs on your host channels.

The co-existing **Slot A** and **Slot B** IDs enable the connection of fault-tolerant links to your subsystem. You may refer to ***Chapter 14*** of this manual for the configuration samples.

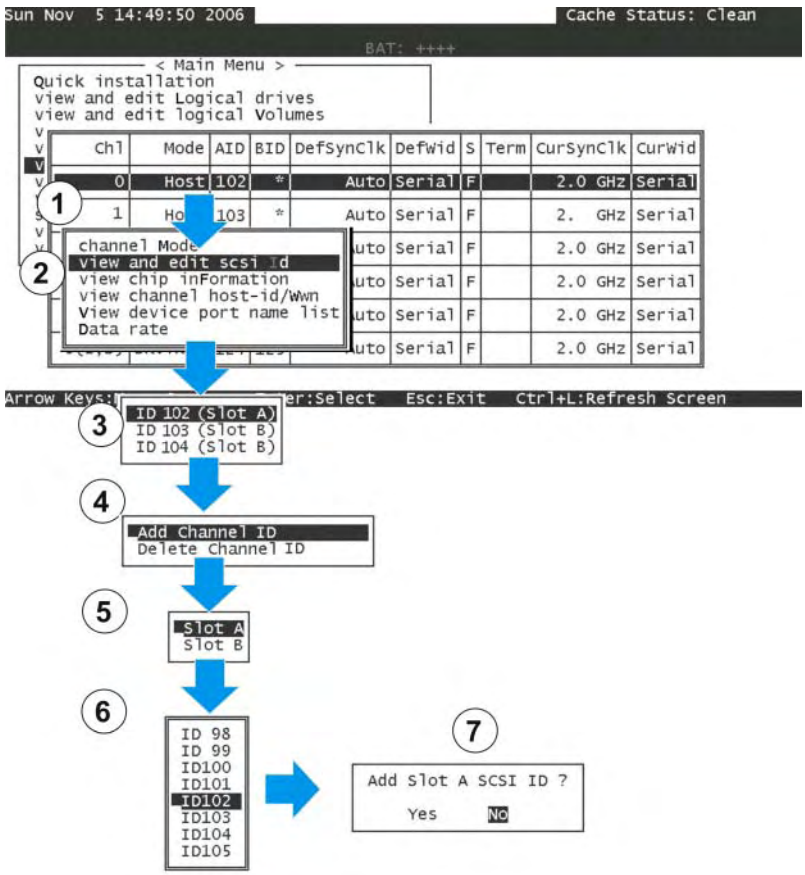

Once Slot B controller IDs are available, you can associate logical arrays with both Slot A and Slot B IDs so that system workload can be shared between partner RAID controllers.

Step 1. Press **[ENTER]** on one of the existing IDs. Choose "**Add Channel ID**," then choose to designate an ID as either the "**Slot A**" ID or "**Slot B**" ID.

Step 2. A list of channel IDs will appear. Select an ID from the pull-down list. In a dual-controller configuration, logical drives associated with a Slot A ID will be managed by the Slot A controller (usually one that resides in the upper controller slot), and a logical drive associated with a Slot B ID managed by the Slot B controller (one in the lower controller slot).

Step 3. A confirm box will prompt reminding you that the configuration change will only take effect after the controller resets. Select **Yes** to confirm.

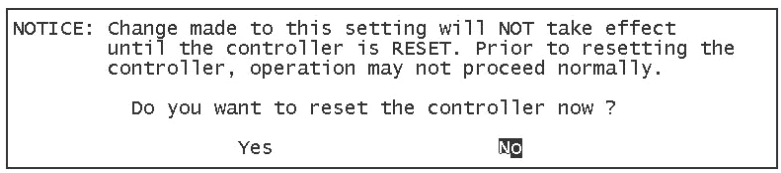

7.1.3 Deleting an ID

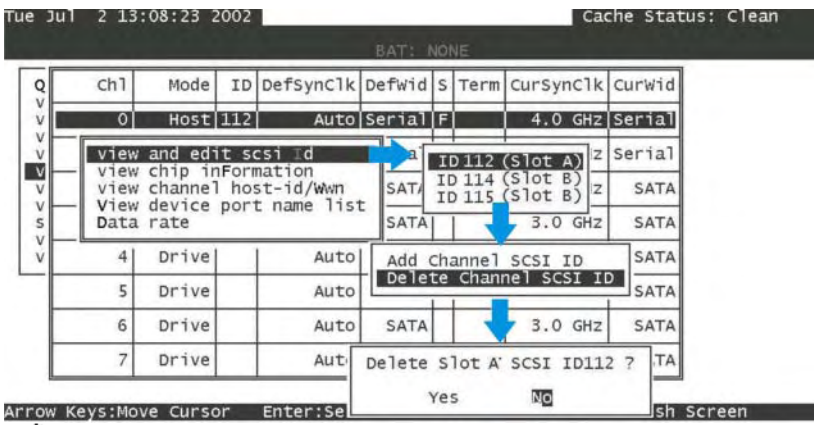

Step 1. Choose the host bus ID you wish to delete. Select "**Delete Channel ID**."

Step 2. The dialog box "**Delete Slot X ID #?**" will appear. Select **Yes**, then press **[ENTER]** to confirm.

Step 3. A confirm box will prompt reminding you that the configuration change will only take effect after the controller resets. Select **Yes** to confirm.

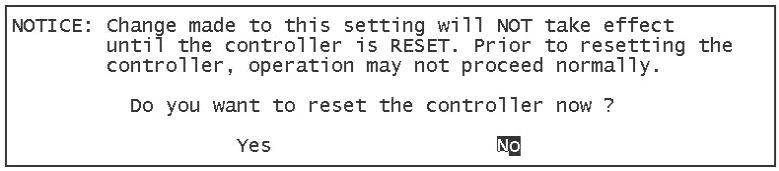

IMPORTANT!

- Every time you add/delete a channel ID, you must reset the subsystem/controller for the changes to take effect.
- Multiple target IDs can co-exist on a host channel while every drive channels in a dual-controller subsystem has two preset IDs.
- At least one ID should be present on each channel bus.

For details on the relationship between host IDs and physical configurations in a dual-controller configuration, please refer to **Chapter 14 Redundant Controller**.

7.1.4 Data Rate (Channel Bus)

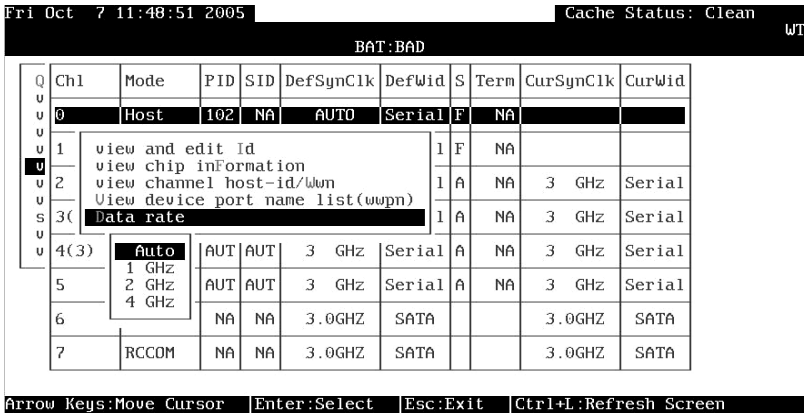

This option is available in the configuration menu of Fibre host channel and of the drive channel configuration menus in Fibre-, SAS-, or SATA-based subsystems. Default is "AUTO" and should work fine with most disk drives. Changing this setting is not recommended unless some particular bus signal issues occur.

7.1.5 View Chip Information

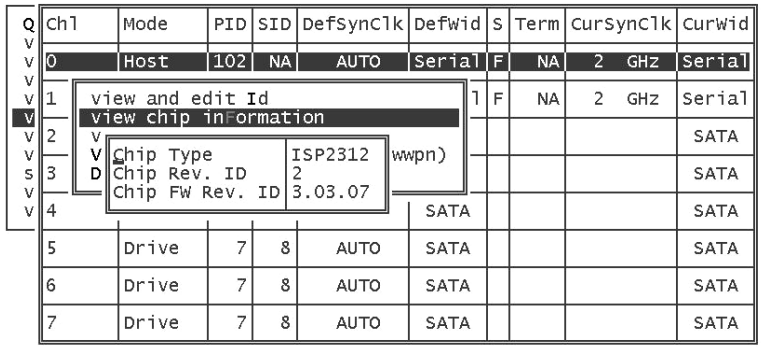

This is a view only option showing basic information about each of the host/drive chip processor.

7.1.6 View Channel Host ID/WWN

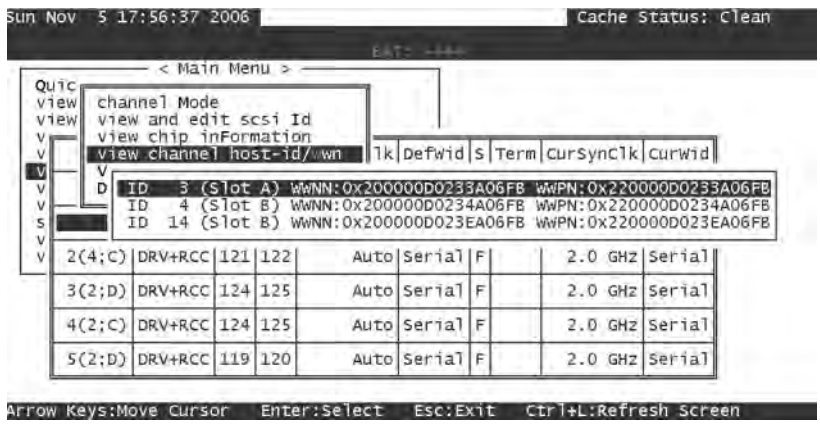

WWNN Node name and **WWPN Port name** are unique eight-byte addresses that appear on a Fibre Channel host port. Every host channel ID appears as a Fibre Channel device and carries both a node name and a port name. A RAID configuration associated with a host ID will also be associated with both a unique node name and a port name.

Corresponding to the dual-ported connectivity defined in Fibre Channel specifications, some of the SAN management software on the market may identify a RAID storage by checking its specific node name and port names.

If a RAID array needs to appear through fault-tolerant links, it needs to be associated with multiple host IDs on separate host ports (channels).

Two identical host IDs (e.g., ID0 on CH0 and ID0 on CH1) on two different host channels carry an identical node name. If an array is associated with these IDs, the array will appear with one node name and two different port names. Some management software will then be able to identify these port names as alternate data paths to a storage device.

The Host ID/WWN option allows users to inspect the node names and port names assigned to specific host IDs. Some management software running on host computers needs these names to properly address a storage subsystem.

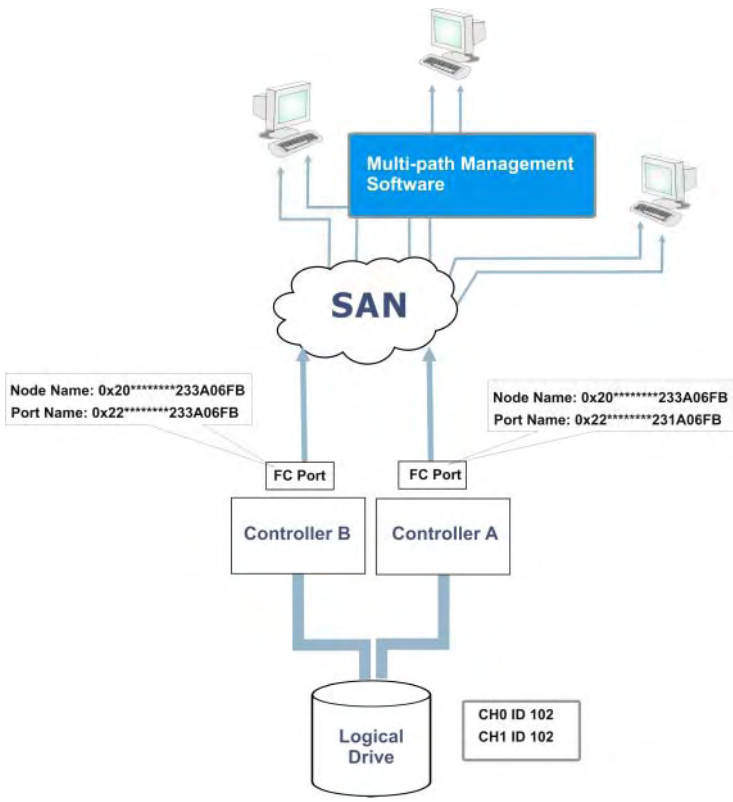

Figure 7-1 Logical Drive Access via Fault-tolerant Links

7.1.7 View Device Port Name List (WWPN)

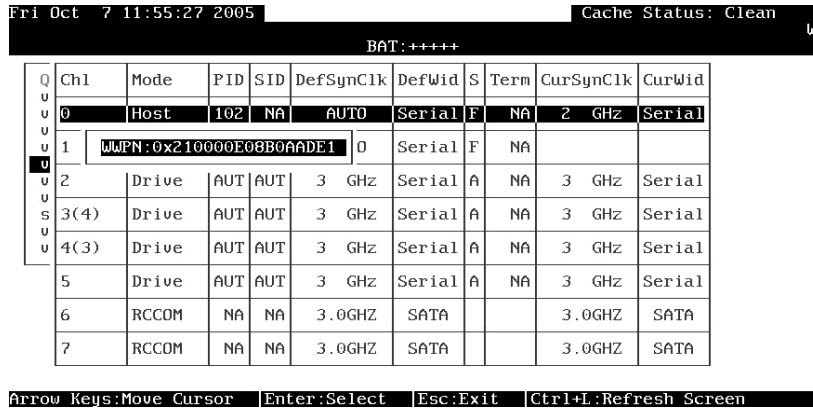

This function displays the device port names (host adapter WWN) of the adapters that appear on a host loop or through a switched fabric connection.

The HBA port names detected here can be manually added to the "**Host-ID WWN name list**" in "**View and Edit Host LUN**" menu. Adding port names to the list can facilitate the LUN mapping.

Each port name can then be assigned a nickname for ease of identification, e.g., financial pool.

7.1.8 Add Host – ID/WWN Label Declaration

A nickname can be appended to any host adapter WWN for ease of identification in SAN environments where multiple servers reside in a storage network.

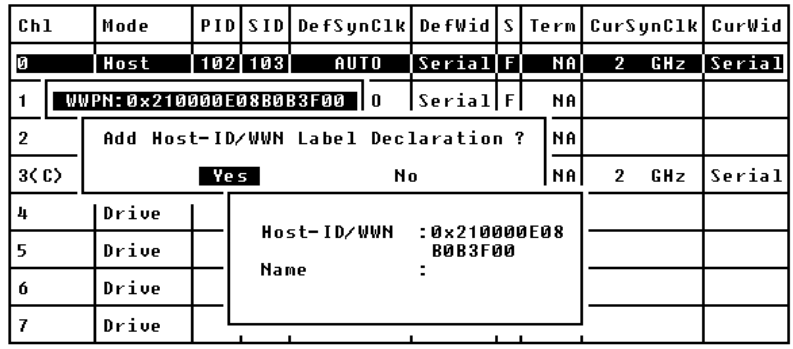

Press **[ENTER]** on a displayed adapter **WWPN** and you will be prompted with a confirm box as shown above. Choose **Yes** and enter a name for the host adapter port.

7.2 Fibre-related Host-side Parameters

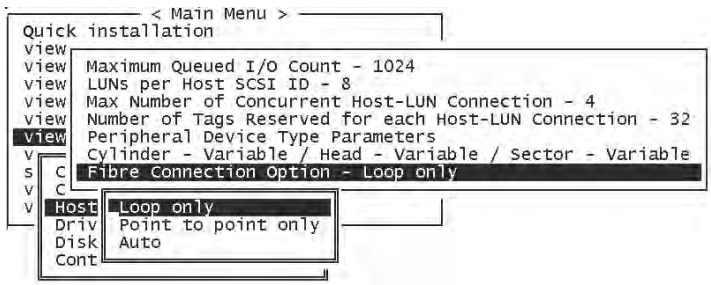

The "Auto" negotiation mode may not be supported by all RAID models and is dependent on the Fibre Channel networking devices.

7.2.1 Fibre Channel Connection Option:

Use the ↑ ↓ keys to scroll down to "**View and Edit Configuration Parameters**," "**Host-side Parameters**," and then "**Fibre Connection Option**." A prompt will display all options. Select an option that is appropriate for your Fibre Channel topology.

With the "**Point-to-Point**" mode, one target host ID will be available on each host port.

With the "**Loop-only**" mode, multiple host IDs will be available on each host port.

7.2.2 Controller Unique Identifier

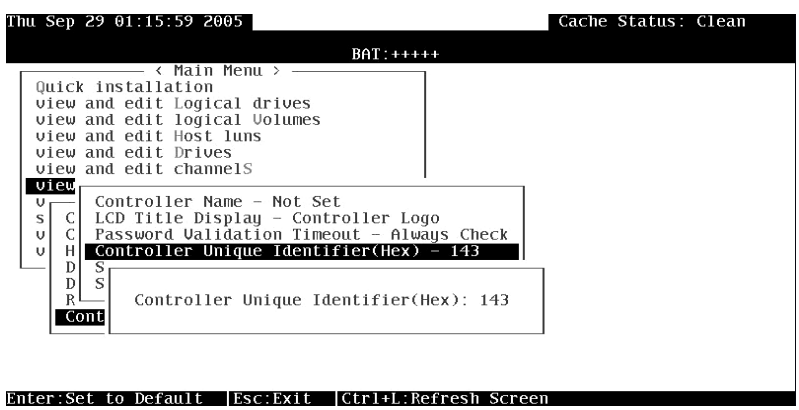

A Controller Unique Identifier is **required** for operation with the **Redundant Controller Configuration**. All EonStor subsystems come with a preset identifier.

The unique identifier will be used to generate a Fibre Channel "**node name**" (**WWNN**). The node name is device-unique and comprised of information such as the IEEE company ID and this user-configurable identifier in the last two bytes.

In redundant mode when a controller fails and a replacement is combined, the node name will be passed down to the replacement, making the host unaware of controller replacement so that the controller failover and failback process can complete in a host-transparent manner.

All EonStor subsystems come with a default identifier. This identifier guarantees your FC ports' port names and node names are unique over a Fibre Channel network. Making changes to the default value is only necessary if the port name conflicts should occur.

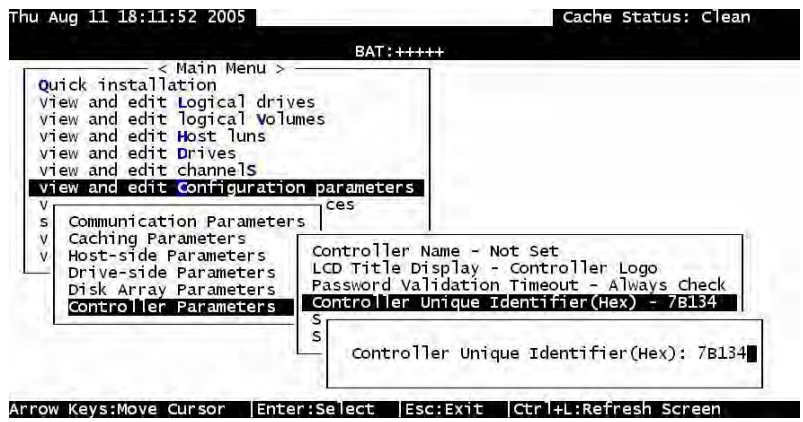

The unique identifier can also be accessed from "**View and Edit Configuration Parameters**" -> "**Controller Parameters**" -> "**Controller Unique Identifier (Hex)**."

Chapter

8

Host-side and Drive-side Parameters

This chapter discusses the advanced options for tuning various firmware parameters. Each function is given a brief explanation as well as a configuration sample. Terminal screens are used in the configuration samples. Some of the operations require basic knowledge of RAID technology and are only recommended for an experienced user.

NOTE:

- All figures in this chapter are showing examples using the management hyper terminal screen.
- Some of the host- or drive-side configurable options are included in different chapters. For example, details about the disk drive S.M.A.R.T. support is provided in ***Chapter 12***.

8.1 Host-side Parameters

The controller supports the following Host-side configurations:

- **Maximum Queued I/O Count**
- **LUNs per Host ID**
- **Num of Host-LUN Connect**
- **Tag per Host-LUN Connect**
- **Peripheral Dev Type Parameters**

- **Cyl/Head/Sector Mapping Config**

8.1.1 Maximum Concurrent Host LUN Connection ("Nexus" in SCSI)

This configuration option adjusts the internal resources for use with a number of concurrent host nexus. If there are four host computers (A, B, C, and D) accessing the array through four host ID/LUN combinations (ID 0, 1, 2, and 3), host A through ID 0 (one nexus), host B through ID 1 (one nexus), host C through ID 2 (one nexus) and host D through ID 3 (one nexus) - all queued in the cache - that is called 4 nexus.

If there are I/Os in the cache through four different nexus, and another host I/O comes down with a nexus different than the four in the cache memory (for example, host A accesses ID 3), the controller will return "busy." Note that it is "concurrent" nexus; if the cache is cleared up, it will accept four different nexus again. Many I/Os can be accessed via the same nexus.

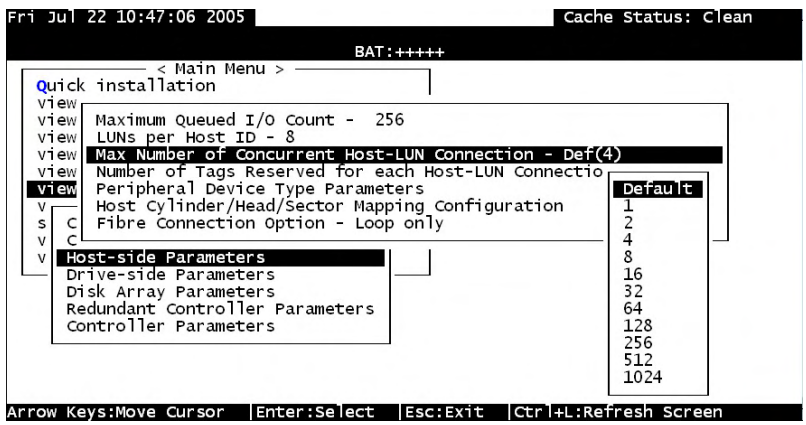

From the Main Menu, select "**View and Edit Configuration Parameters**," "**Host-side Parameters**," then press **[ENTER]**. Choose "**Max Number of Concurrent Host-LUN Connection**," then press **[ENTER]**. A list of available selections will appear. Move cursor bar to an item, then press **[ENTER]**. Choose **Yes** in the dialog box that follows to confirm your setting. The default is "4."

8.1.2 Number of Tags Reserved for Each Host-LUN Connection

Each nexus has 32 (the default setting) tags reserved. When the host computer sends 8 I/O tags to the controller, and the controller is too busy to process them all, the host might start to send less than 8 tags during every certain period of time since then. This setting ensures that the controller will accept at least 32 tags per nexus.

The controller will be able to accept more than that as long as the controller internal resources allow - if the controller does not have enough resources, at least 32 tags can be accepted per nexus.

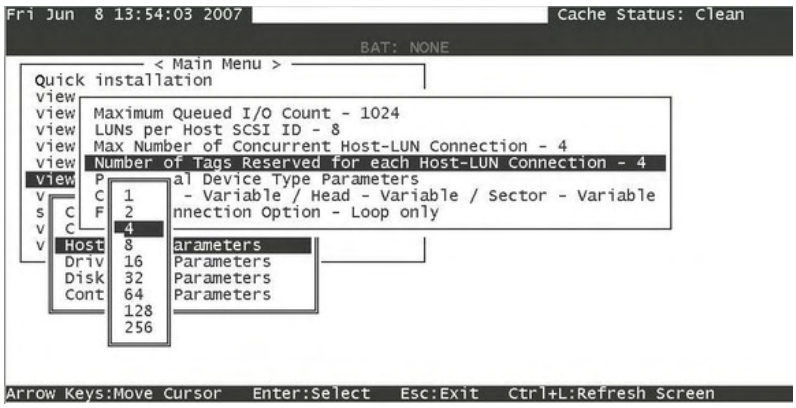

Choose "**Host-side Parameters**," then press **[ENTER]**. Choose "**Number of Tags Reserved for each Host-LUN Connection**," then press **[ENTER]**. A list of available selections will appear. Move the cursor bar to an item, then press **[ENTER]**. Choose **Yes** in the dialog box that follows to confirm the setting.

8.1.3 Maximum Queued I/O Count

This function allows you to configure the maximum queued I/O count the controller can receive from the host computer.

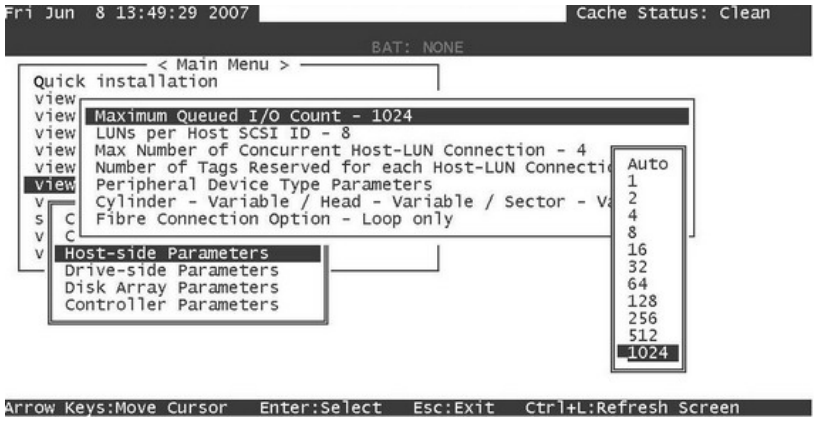

Choose "**Host-side Parameters**," then press **[ENTER]**. Choose "**Maximum Queued I/O Count**," then press **[ENTER]**. A list of available selections will appear. Move the cursor bar to an item, then press **[ENTER]**. Choose **Yes** in the dialog box that follows to confirm the setting.

8.1.4 LUNs per Host ID

The highly scalable Fibre Channel technology can address up to 126 devices per loop, and theoretically more than a million using FC switches. Each configured RAID volume is associated with host IDs and appears to the host as a contiguous volume.

If you file a document into a cabinet, you must put the document into one of the drawers. As defined by storage interface architecture, a Fibre channel ID is like a cabinet, and the drawers are the LUNs (Logical Unit Numbers). Each Fibre channel ID encapsulates up to 32 LUNs and up to 1024 LUNs are configurable through all host ports. A RAID volume can be associated with any of the LUNs under the Fibre channel IDs. Most Fibre host adapters treats a LUN like another Fibre device.

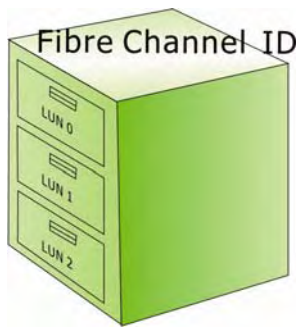

Figure 8-1: ID and LUNs as a Cabinet and Its Drawers

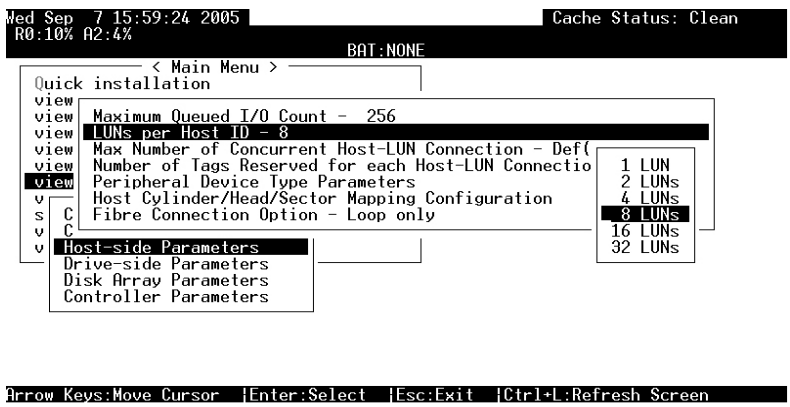

Choose "**LUNs per Host ID**," then press **[ENTER]**. A list of selections will appear. Move the cursor bar to an item, then press **[ENTER]**. Choose **Yes** in the dialog box that follows to confirm the setting.

8.1.5 LUN Applicability

Why Setting LUN Applicability?

The LUN Applicability settings apply in environments where system administrators use in-band methodology for management access to a RAID subsystem.

If no logical drive has been created and mapped to a host LUN, and the RAID controller is the only device connected to the host computer, usually the operating system will not load the driver for the host adapter. If the driver is not loaded, the host computer will not be able to use the in-band utility to communicate with the RAID controller. This is often the case when users want to start configuring a brand new subsystem using RAIDWatch manager software.

By then it will be necessary to configure the "**Peripheral Device Type**" setting for the host to communicate with the RAID controller.

- **LUN-0 only**

  If the "**LUN-0's only**" is selected, only LUN-0 of the host ID will appear as a device with the user-defined peripheral device type.

- **All Undefined LUNs**

  If "**all undefined LUNs**" is selected, every LUNs under that host ID will appear as a device with the user-defined peripheral device type.

Please refer to ***8.1.8 Peripheral Device Type Parameters for Various Operating Systems***
for details about the compatibility with various operating systems.

8.1.6 Peripheral Device Type

For management access to a new subsystem without pre-configured logical units and Ethernet connection, the in-band SCSI protocol can be used in order for a management computer to “see” the RAID subsystem. Please refer to the reference table below. You will need to make adjustments in the following submenus: **Peripheral Device Type**, **Peripheral Device Qualifier**, **Device Support for Removable Media**, and **LUN Application**.

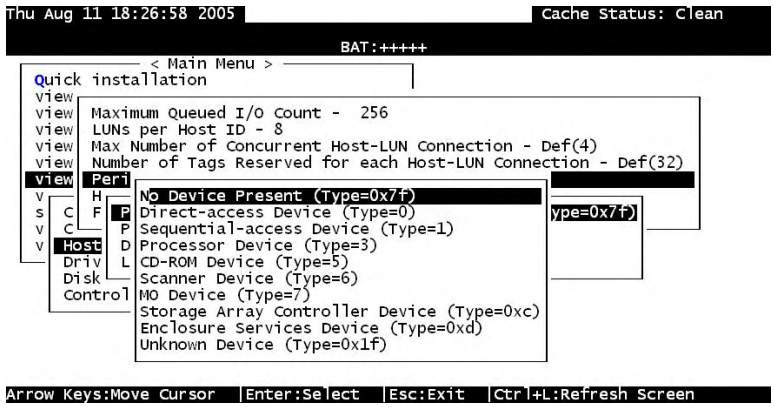

8.1.7 In-band Management Access

External devices (including a RAID subsystem; from the view of an application server or management PC) require communication links with a management computer for device monitoring and administration. In addition to the regular RS-232C or Ethernet connection, in-band SCSI can serve as an alternative means of management communications. In-band SCSI translates the original configuration commands into standard SCSI commands. These SCSI commands are then sent to and received by the controller over the existing host links, either SCSI or Fibre.

8.1.8 Peripheral Device Type Parameters for Various Operating Systems

IMPORTANT!

There is no need to configure the Peripheral Device setting if you are trying to manage a RAID subsystem from a RAIDWatch/SANWatch station through an Ethernet connection (to the EonStor subsystem's Ethernet port). An Ethernet connection to RAID uses TCP/IP as the communication protocol.

With an in-band connection, a host computer cannot "see" a RAID controller **UNLESS** one of the following conditions is met:

Condition 1:

(a) A logical unit has been created.

(b) At least one logical unit is mapped to a host ID or LUN using the RS-232/LCD keypad interface.

Condition 2:

(a) Host bus connection is established.

(b) The RAID subsystem/controller is configured to appear as a peripheral device through the host links.

With a brand new array, there is no association between disk drive configurations and the logical ID/LUN presentations on the host bus. If users want to start configuring a RAID system before any RAID configuration is made, the host will not be able to "see" the RAID subsystem. In order for a host to "see" the subsystem, it will be necessary to define the controller/subsystem as a peripheral device.

Different host operating systems require different adjustments. See the tables below to find appropriate settings for your host operating system. References to "**Peripheral Device Qualifier**" and "**Device Support for Removable Media**" are also included.

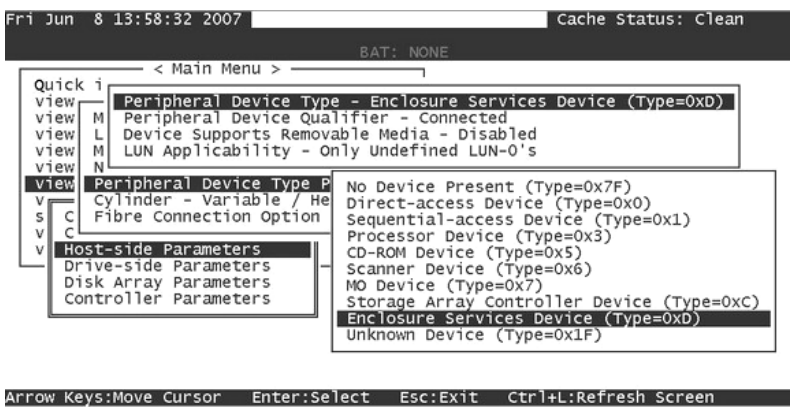

Table 8-1: Peripheral Device Type Parameters

Operating System	Peripheral Device Type	Peripheral Device Qualifier	Device Support for Removable Media	LUN Applicability
Windows 2000/2003	0xd	Connected	Either is okay	LUN-0's
Solaris™ 8/9 (x86 and SPARC)	0xd	Connected	Either is okay	LUN-0's
Linux RedHat 8/9; SuSE 8/9	0xd	Connected	Either is okay	LUN-0's

Table 8-2: Peripheral Device Type Settings

Device Type	Setting
Enclosure Service Device	0xd
No Device Present	0x7f
Direct-access Device	0
Sequential-access Device	1
Processor Type	3
CD-ROM Device	5
Scanner Device	6
MO Device	7
Storage Array Controller Device	0xC
Unknown Device	0x1f

8.1.9 Cylinder/Head/Sector Mapping

Drive capacity is decided by the number of blocks. For some operating systems (Sun Solaris, for example) the capacity of a drive is determined by the cylinder/head/sector count. For earlier Sun Solaris systems, the cylinder cannot exceed 65535; choose "**cylinder<65535**," then the controller will automatically adjust the head/sector count for your OS to read the correct drive capacity. Please refer to the related documents provided with your operating system for more information.

Cylinder, Head, and Sector counts are selectable from the configuration menus shown below. To avoid any difficulties with a Sun Solaris configuration, the values listed below can be applied.

Table 8-3: Cylinder/Head/Sector Mapping under Sun Solaris

Capacity	Cylinder	Head	Sector
< 64 GB	variable	64	32
64 - 128 GB	variable	64	64
128 – 256 GB	variable	127	64
256 – 512 GB	variable	127	127
512 GB - 1 TB	variable	255	127

Older Solaris versions do not support drive capacities larger than 1 terabyte. ***Solaris 10*** now supports array capacity larger than 1TB. Set the values to the values listed in the table below:

Table 8-4 Cylinder/Head/Sector Mapping under Sun Solaris 10

Capacity	Cylinder	Head	Sector
>1TB	<65536	255	variable
		variable	255

Configuring Sector Ranges/Head Ranges/Cylinder Ranges:

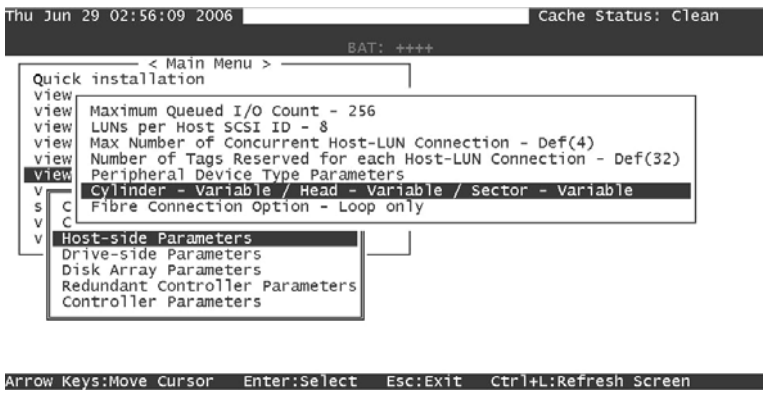

The sector, head, and cylinder variables are presented as preset combinations. Please refer to the documentation that came with your operating system and select one value set that is most appropriate for your OS file system.

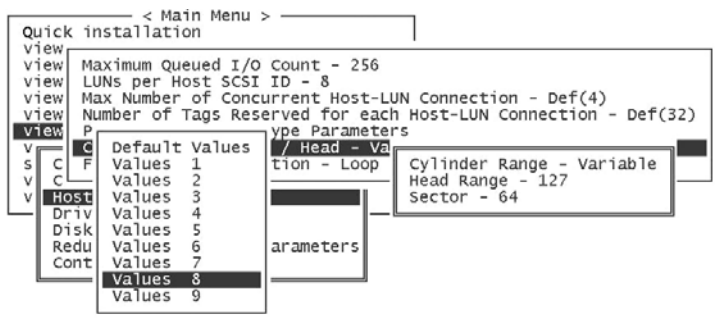

8.2 Drive-side Parameters:

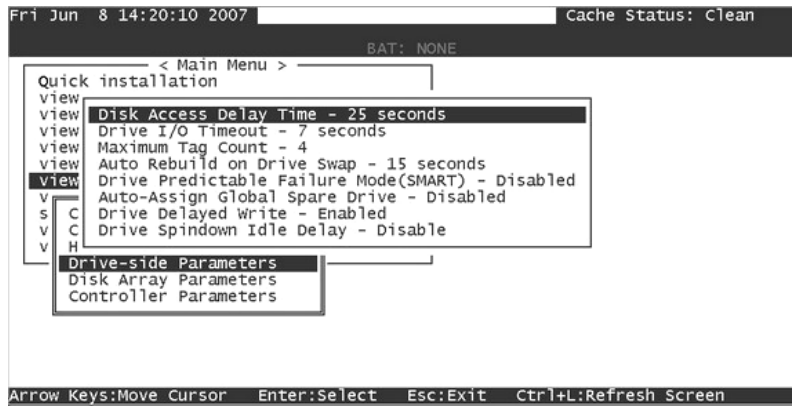

Select "**Drive-side Parameters**" from the Main Menu, then press **[ENTER]**. The Drive-side Parameters menu displays.

8.2.1 Disk Access Delay Time

This feature sets the delay time before the subsystem tries to access the hard drives after power-on. Default may vary from 15 seconds to 30 seconds, and is determined by the type of drive interface. This parameter can be adjusted to fit the spin-up speed of different disk drive models.

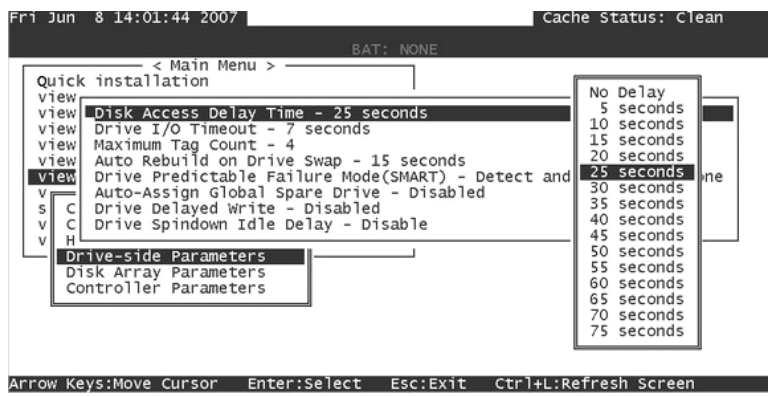

Select "**Disk Access Delay Time**," then press **[ENTER]**. A list of selections displays. Move the cursor bar to a selection, then press **[ENTER]**. Choose **Yes** in the dialog box that follows to confirm the setting.

8.2.3 Drive I/O Timeout

The "Drive I/O Timeout" is the time interval for the controller to wait for a drive to respond. If the controller attempts to read data from or write data to a drive but the drive does not respond within the Drive I/O Timeout value, the drive will be considered as a failed drive.

When the drive itself detects a media error while reading from the drive platter, it usually retries the previous reading or re-calibrates the read/write head. When a disk drive encounters a bad block on the media, it will attempt to reassign the bad block to a spare block. However, it takes time to perform the above operations. The time to perform these operations can vary between among disk drives by different vendors.

During channel bus arbitration, a device with higher priority can utilize the bus first. A device with lower priority will sometimes receive an I/O timeout when devices of higher priority keep utilizing the bus.

The default setting for "Drive I/O Timeout" is 7 seconds. It is highly recommended not to change this setting. Setting the timeout to a lower value will cause the controller to judge a drive as failed while a drive is still retrying, or while a drive is unable to arbitrate the drive bus. Setting the timeout to a greater value will cause the controller to keep waiting for a drive, and it may sometimes cause a host timeout.

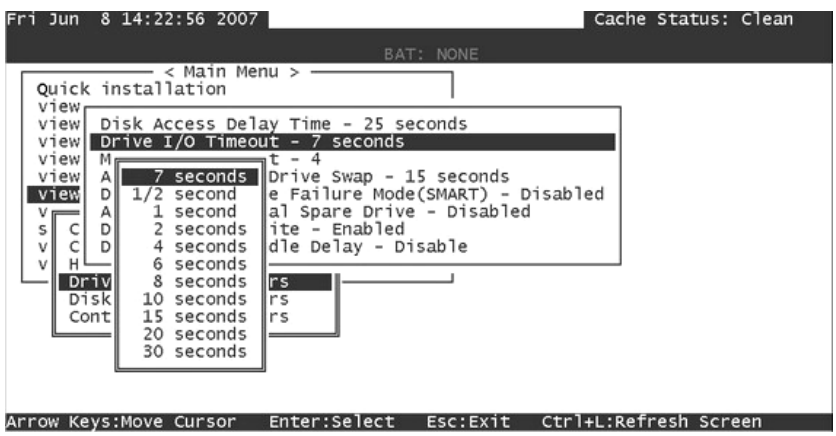

Choose "Drive I/O Timeout –Default (7 seconds)," then press **[ENTER]**. A list of selections will appear. Move the cursor bar to a selection, then press **[ENTER]**. Choose **Yes** in the dialog box that follows to confirm the setting.

8.2.4 Maximum Tag Count: Tag Command Queuing (TCQ) and Native Command Queuing (NCQ) Support

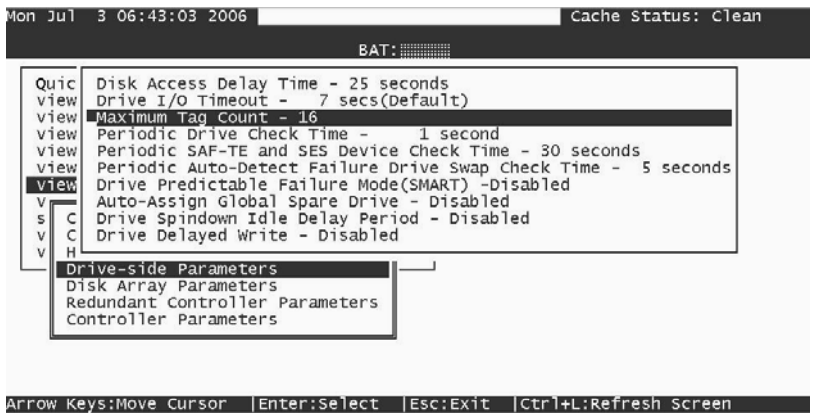

This sub-menu facilitates the support for both **Tagged Command Queuing** (TCQ) and **Native Command Queuing** (NCQ). TCQ is a traditional feature on SCSI, SAS, or Fibre Channel disk drives, while NCQ is recently implemented with SATA disk drives. The queuing feature requires the support of both host adapters and hard disk drives. Command queuing can intelligently reorder host requests to streamline random accesses for IOPS/multi-user applications.

Infortrend's subsystems support Tag Command Queuing with an adjustable maximum tag count from **1** to **128**. The default setting is "Enabled" with a maximum tag count of 32 (SCSI), 8 (for Fibre drives), or 4 (default for SAS/SATA drives). Choose "Maximum Tag Count", then press **[ENTER]**. A list of available tag count numbers displays. Move the cursor bar to a number, then press **[ENTER]**. Choose **Yes** in the dialog box that follows to confirm the setting.

IMPORTANT!

- Every time you change this setting, you must reset the controller/subsystem for the changes to take effect.
- Disabling Tag Command Queuing will disable the hard drives' built-in buffer.

NOTE:

The following options are categorized as related to array maintenance and data integrity:

- Auto Rebuild on Drive Swap

- Auto-Assign Global Spare Drive

Another option is associated with disk drive S.M.A.R.T. support and details can be found in ***Chapter 12:***

- Drive Predictable Failure Mode (SMART)

8.2.5 Drive Delayed Write

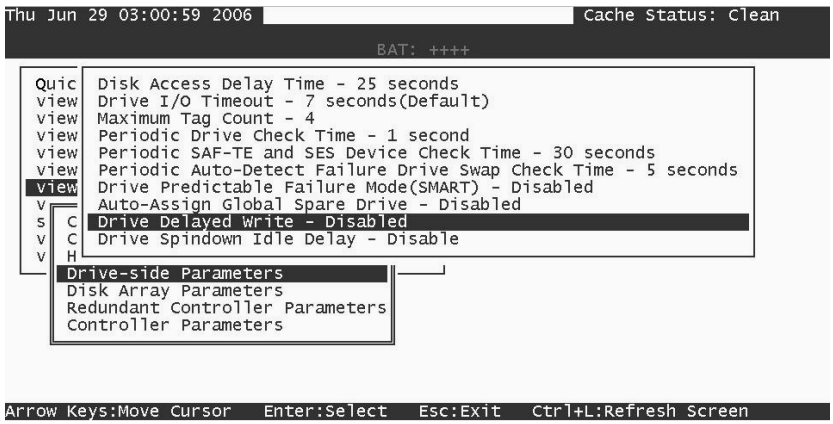

This option applies to disk drives which come with embedded buffers. When enabled, write performance can improve. However, this option should be disabled for mission-critical applications. In the event of power outage or drive failures, data cached in drive buffers may be lost, and data inconsistency will occur.

This page is intentionally left blank.

Chapter

9

Enclosure Management

This chapter discusses the configuration options related to enclosure monitoring. Each function is given a brief explanation as well as a configuration sample. Terminal screens will be used in the configuration samples. Some of the operations require basic knowledge of RAID technology and are only recommended for an experienced user.

NOTE:

All figures in this chapter are showing examples from a hyper terminal console.

9.1 Enclosure Device Statuses (Peripheral Device Status)

RAID Enclosure Devices

To check the operational statuses of enclosure devices, follow the steps below:

Step 1. Select "**View and edit Peripheral Devices**" on the Main Menu and press **[ENTER].**

Step 2. Choose "**View Peripheral Device Status**," then press **[ENTER]** again. The device list displays.

Step 3. Press **[ENTER]** on the "**SES Device**" or "**I2C Peripheral Device**" to display a list of peripheral devices (enclosure modules). Monitoring of device

status depends on enclosure implementation and is accessed through different interfaces, e.g., S.E.S., SAS wide links, or I2C serial bus. Enclosure devices usually include the following:

1. Drive failure output definition
2. Cooling FAN
3. Power supply
4. Temperature Sensors
5. Device set description

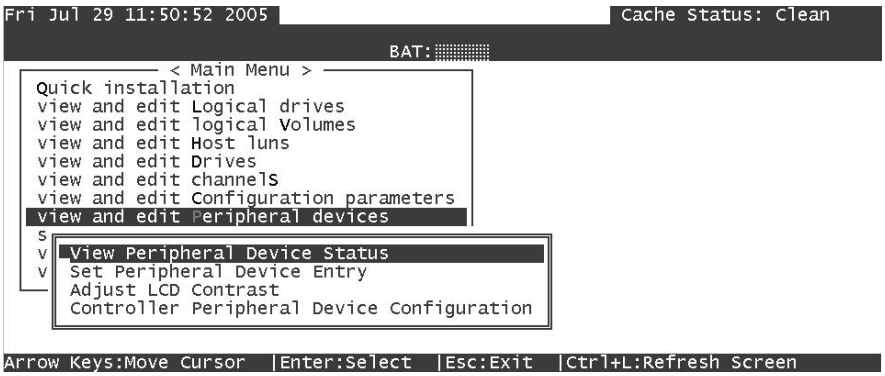

Below is a screen showing the enclosure devices interfaced through an I2C serial bus:

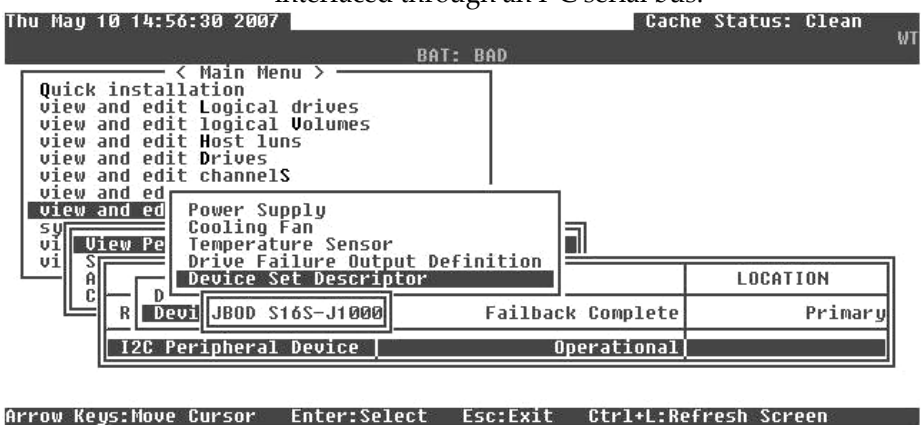

NOTE:

A SAS expansion enclosure connected through SAS links is also considered as an I2C Peripheral Device, which is defined as the **Device Set 1** (JBOD enclosure device) next to the **Device Set 0** (RAID enclosure device).

Step 4. Press **[ENTER]** on a component type to examine its operating status.

Following is a screen listing all cooling fans in a 3U enclosure, including those embedded in power supply modules.

Cooling fan0	High
Cooling fan1	High
Cooling fan2	High
Cooling fan3	High
PSU 0-FAN 0	8600RPM
PSU 0-FAN 1	8200RPM
PSU 1-FAN 0	Unknown
PSU 1-FAN 1	Unknown
PSU 2-FAN 0	8200RPM
PSU 2-FAN 1	8600RPM

Devices within the Expansion Enclosure

Devices in SAS expansion enclosures are monitored through a proprietary in-band methodology through a monitor chipset on JBOD controllers. Below is the device shown on the **View and Edit Drives** screen.

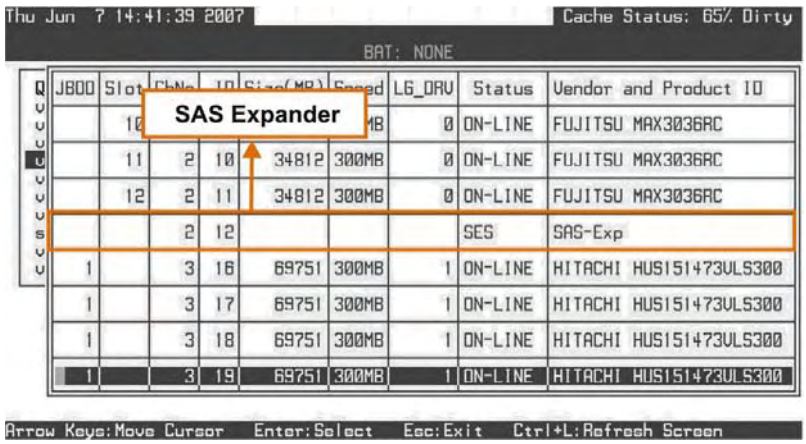

Information about the SAS expander handling SAS expansion links is shown as the last device in the RAID enclosure.

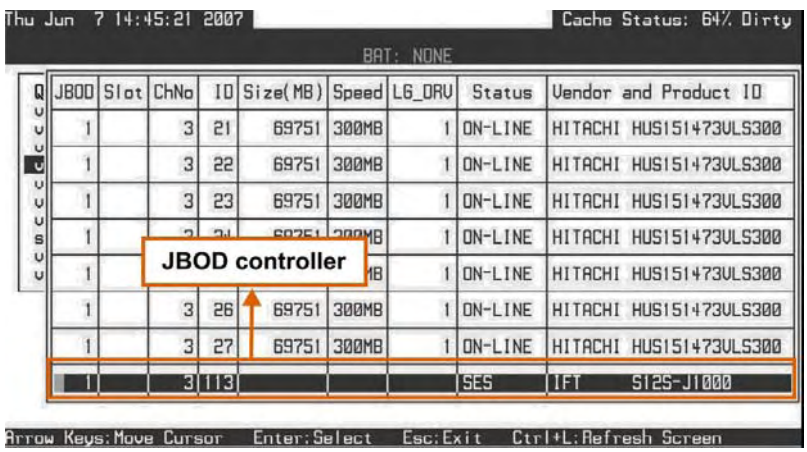

The JBOD controller within the expansion enclosure is shown as the last device in the expansion enclosure. You may press **[ENTER]** on

the device to check the revision number of the firmware running on SAS channel devices.

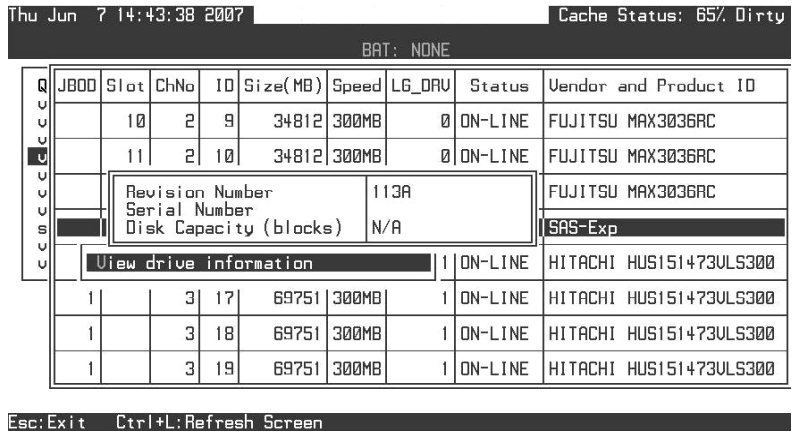

The operating statuses of individual enclosure devices within the expansion enclosures can be found in Main Menu -> "**View and Edit Peripheral Device**" -> "**View Peripheral Device**"- > "**I2C Peripheral Device**" and a list of devices will be displayed. Press **[ENTER]** on any to enter its status screen.

NOTE:

The JBOD enclosure devices will only display when firmware detects expansion enclosures across its expansion links.

Verifying Disk Drive Failure in a Multi-enclosure Application

You can verify disk drive locations by checking their channel number, slot number, and device IDs in "**Drive Failure Output Definition**." Note that the SAS channel number is a logically defined congregation of multiple physical links (PHYs) through the SAS expanders.

This information is important for locating and replacing a failed drive.

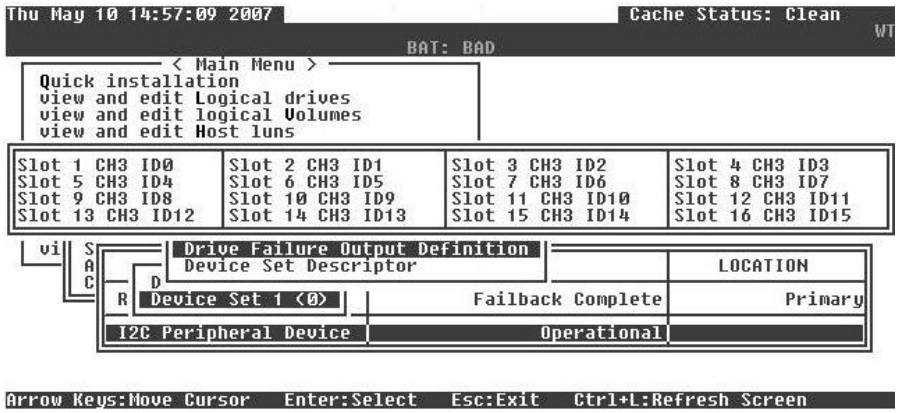

Another key factor in identify drive location is the JBOD/SBOD identifier that can be found under the Main Menu -> "**View and Edit Drives**" sub-menu. The JBOD identifier equals the enclosure ID you configure using the front panel rotary switch or the rear panel DIP switches.

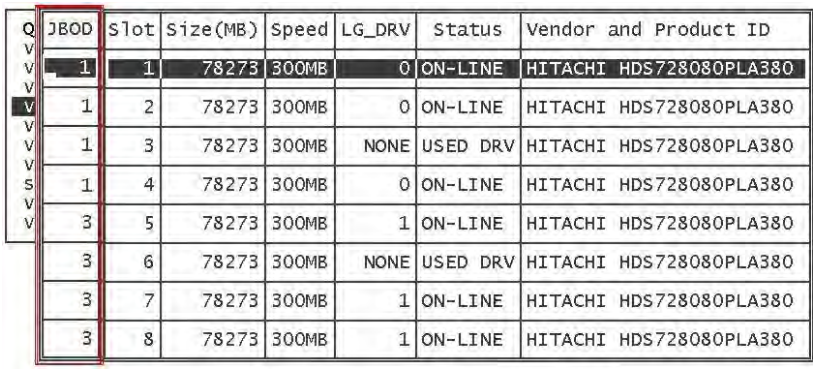

JBOD	Slot	Size(MB)	Speed	LG_DRV	Status	Vendor and Product ID
1	1	78273	300MB	0	ON-LINE	HITACHI HDS728080PLA380
1	2	78273	300MB	0	ON-LINE	HITACHI HDS728080PLA380
1	3	78273	300MB	NONE	USED DRV	HITACHI HDS728080PLA380
1	4	78273	300MB	0	ON-LINE	HITACHI HDS728080PLA380
3	5	78273	300MB	1	ON-LINE	HITACHI HDS728080PLA380
3	6	78273	300MB	NONE	USED DRV	HITACHI HDS728080PLA380
3	7	78273	300MB	1	ON-LINE	HITACHI HDS728080PLA380
3	8	78273	300MB	1	ON-LINE	HITACHI HDS728080PLA380

9.2 Enclosure Management Options

9.2.1 Enclosure Devices

9.2.1.1 Event Triggered Operations

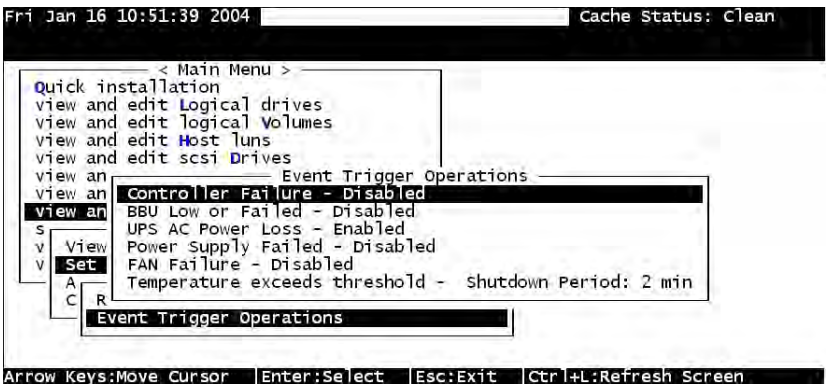

Step 1. Use arrow keys to move your cursor bar to select "**View and Edit Peripheral Devices**" on the Main Menu and press **[ENTER].**

Step 2. Choose "**Set Peripheral Device Entry**", press **[ENTER],** then select "**Event Trigger Operations**" by pressing **[ENTER]**. The event trigger menu displays.

Step 3. Select any of the monitoring elements by moving the cursor bar and pressing **[ENTER]** to enable or disable the association with related system events.

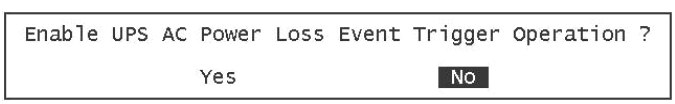

NOTE:

The last condition, the "**Temperature Threshold**," is associated with a configurable time buffer before an automatic shutdown. Please refer to the next section for details.

Operation Theory:

The Operation:

To reduce the chance of data loss due to hardware failure, the controller/subsystem automatically commences the following actions when a component failure is detected:

1). Switches its caching mode from "write-back" to the conservative "write-through."

2). Flushes all cached data.

3). Raises the rotation speed of cooling fans.

The Trigger:

The mode-switching and cache-flush operations can be triggered by the occurrences of the following conditions:

1. **Controller failure (Dual-controller Models)**

   If a controller fails in a dual-redundant controller configuration, the surviving controller no longer has the protection of synchronized cache by having the replica of unfinished writes in its partner.

2. **BBU low or failed:**

   If a battery fails or is under-charge, the unfinished writes cannot be supported if power outage occurs.

3. **UPS AC power loss:**

   Even with the buffer provided by the UPS, if power outage occurs, cached data should be immediately distributed to hard drives before the battery charge in UPS runs out.

4. **Power supply failure**
5. **Fan failure**
6. **Temperature exceeds threshold**

If one or more of the event triggers listed above are enabled, the occurrence of the above conditions forces the controller/subsystem to adopt the "**write-through**" caching mode. Once the faulty condition is corrected, the controller/subsystem automatically restores the previous caching mode.

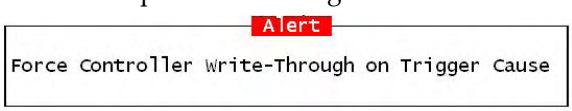

NOTE:

- The temperature thresholds refer to those set for both sensors on the RAID controller boards and those placed within the subsystem enclosure. In terms of the controller temperature, board 1 refers to the main circuit board and board 2 refers to the second-level I/O board or the daughter card. If any of the threshold values set for any sensor is exceeded, the reaction mode is automatically triggered.
- If a battery is not installed in your RAID subsystem, the "**BBU Low or Failed**" option should be disabled.

9.2.2 Auto Shutdown on Elevated Temperature

System components can be damaged if operated under elevated temperature. You can configure the time periods between the detection of exceeded thresholds and the controller's commencing an automatic shutdown.

The shutdown ***does not*** electrically disconnect the subsystem. When shutdown is commenced, the subsystem stops responding to I/O requests and flushes all cached writes in its memory. During that time, system administrators should have been notified of the condition and have begun restoring proper cooling of the subsystem. Extended operation under critical conditions like elevated temperature greatly reduces system efficiency and will eventually cause component failure.

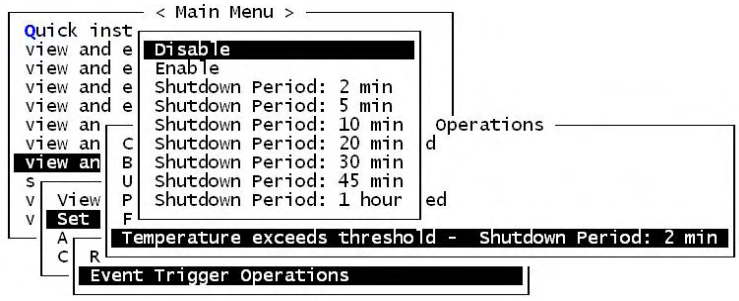

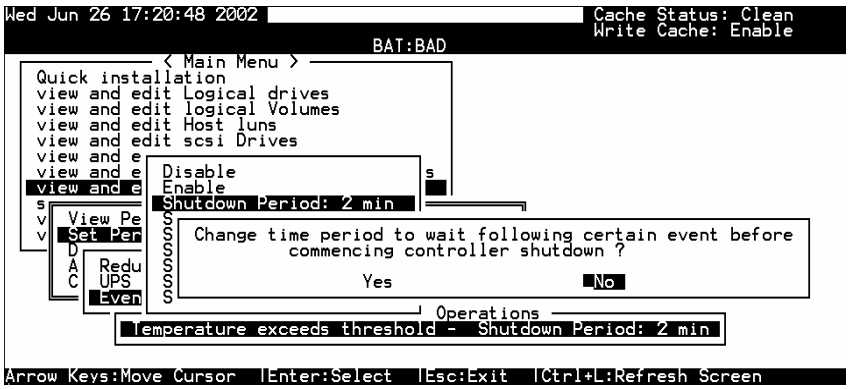

Step 1. Select "**View and Edit Peripheral Devices**" on the Main Menu and press **[ENTER].**

Step 2. Choose "**Set Peripheral Device Entry**" and "**Event Trigger Option**" by pressing **[ENTER]**. The auto-shutdown options display.

Step 3. Move your cursor bar to "**Temperature exceeds threshold**." Press **[ENTER]** and select a configurable time span between the detection of exceeded temperature and the controller's commencing an automatic shutdown.

9.2.3 Voltage and Temperature Self-monitoring

Follow the steps below to check voltage and temperature readings:

Step 1. Enter the Main Menu and select "**View and Edit Peripheral Devices**."

Step 2. Use the arrow keys to scroll down and select "**Controller Peripheral Device Configuration**," "**View Peripheral Device Status**," and then press **[ENTER]**.

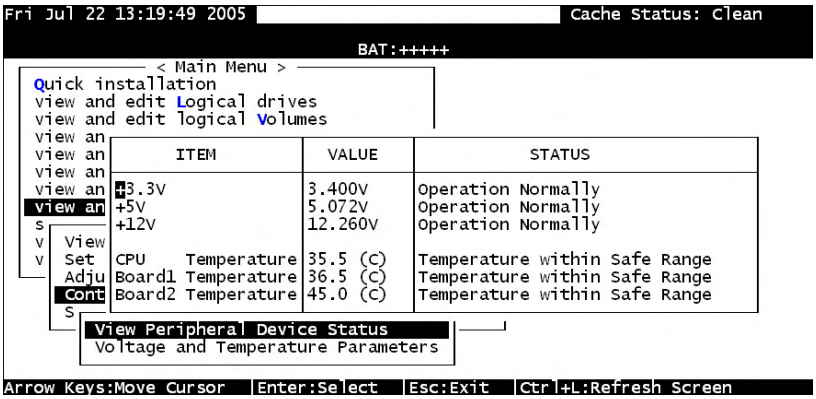

Step 3. The controller voltage and temperature readings display.

NOTE:

The enclosure fan status and the temperature readings from chassis sensors are shown in ***9.1 Enclosure Device Statuses (Peripheral Device Status).***

9.2.3.1 Changing Monitoring Thresholds

Step 1. Open your PC Terminal Emulation utility. Enter the Main Menu and select "**View and Edit Peripheral Devices**."

Step 2. Use the arrow keys to scroll down and select "**Controller Peripheral Device Configuration**,"

"**Voltage and Temperature Parameters**," and confirm by pressing **[ENTER]**.

CAUTION!

It is not recommended to change the threshold values unless you need to coordinate the RAID controller's values with that of your RAID enclosure. If a value exceeding the safety range is entered, an error message will prompt and the new parameter will be ignored.

For example, if the controller operates in a system enclosure where the upper limit on ambient temperature is relatively higher or lower, adjusting the default thresholds can coordinate the controller status monitoring with that of your specific enclosure specifications.

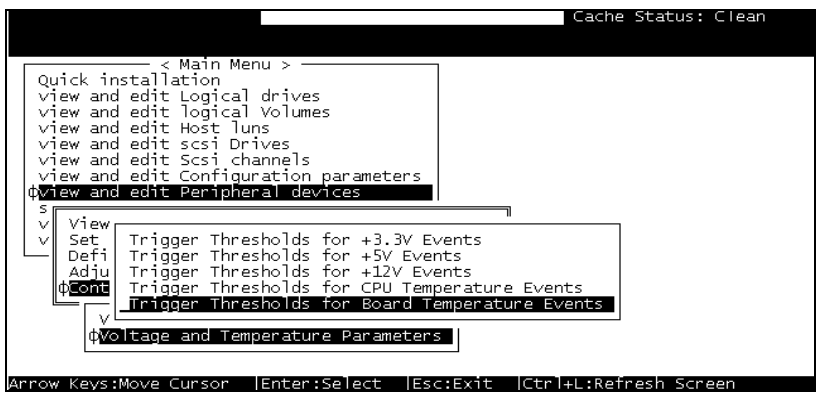

Step 3. Scroll down and select an item to configure.

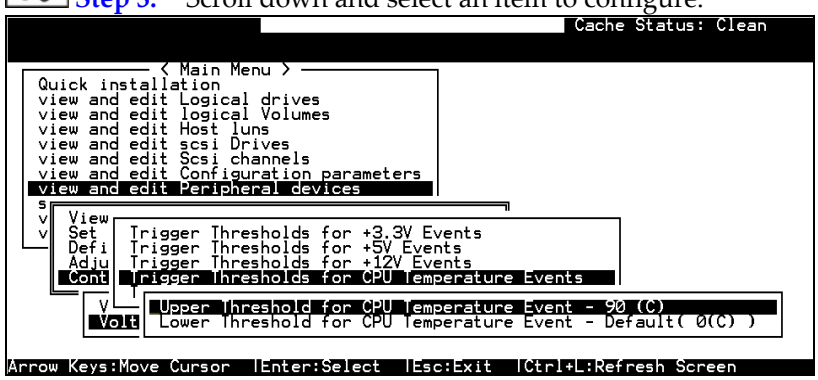

Step 4. Select an item, such as "**Trigger Thresholds for CPU Temperature Events**." Press **[ENTER]** and a list of selections will appear. You can change the upper or lower threshold values by keying a number. Press **[ENTER]** to confirm.

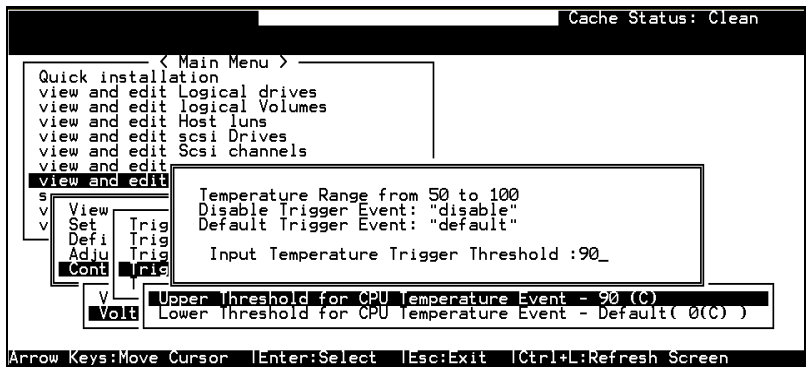

Step 5. A configuration window will prompt. Enter any value within the safety range. Values exceeding the safety range will be rejected by controller firmware.

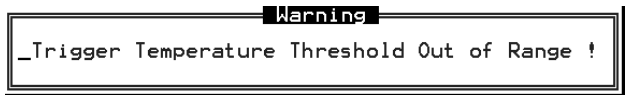

Step 6. Follow the same method to modify other threshold parameters.

9.3 UPS Support

UPS status monitoring is available from firmware release 3.34 onward. This feature enables the RAID subsystem to observe and respond to the status of the UPS (battery charge or AC power) via the COM2 serial port connection and dynamically switch the write policy.

Requirements for UPS Status Monitoring:

1. Connect the serial port on a UPS device to the controller/subsystem's COM2 serial port.
2. Set the same Baud Rate to the system COM2 and the UPS serial port. Please refer to ***Chapter One, RS-232C Serial Port Settings***.
3. Set the UPS option as described in ***9.2.1.1*** to "Enabled."

Condition Reports and Reactions:

1. When mains power is lost or when the UPS charge is low, an event is issued. The subsystem will commence an auto cache-flush and will be forced to adopt the conservative caching mode, the "write-through" mode.

2. When the serial port connection is lost or when the UPS is disconnected, an event is issued to notify system managers that the external UPS is absent.
3. When the UPS battery charge or mains power is restored to a safe functioning level, the subsystem automatically restores the original write policy.

This page is intentionally left blank.

Chapter

10

Data Integrity

This chapter discusses various firmware mechanisms that help to ensure data integrity.

No system is completely safe from hardware faults. For example, although the chance of occurrence is considerably low, the occurrences of bad blocks on two (RAID 5) or three (RAID 6) hard drives can fail a whole data set. When properly configured, the functions below help to minimize the chance of data loss:

1. Event Triggered Operations (please refer to the previous chapter)
2. Failed Drive Detection
3. Scheduled Maintenance
4. Regenerate Logical Drive Parity
5. Rebuild Priority
6. Verification on Writes

NOTE:

- Some of the configuration options may not be available to all sub-revisions of firmware.
- All figures in this chapter are showing examples of a management console over an RS-232 or telnet connection.

10.1 Failed Drive Detection

10.1.1 Auto Rebuild on Drive Swap Check Time

The "**Auto Rebuild on Drive Swap**" check time is the interval at which the controller checks to see if a failed drive has been swapped. When a member of a logical drive fails, the controller will continuously scan the drive bus (at the selected time interval). Once the failed drive has been swapped with a drive that has the adequate capacity to rebuild the logical drive, the rebuild will begin automatically.

The default setting is "15 seconds," meaning that the controller will automatically scan the drive busses if a failed drive has been replaced. To change the timeout, select a time interval.

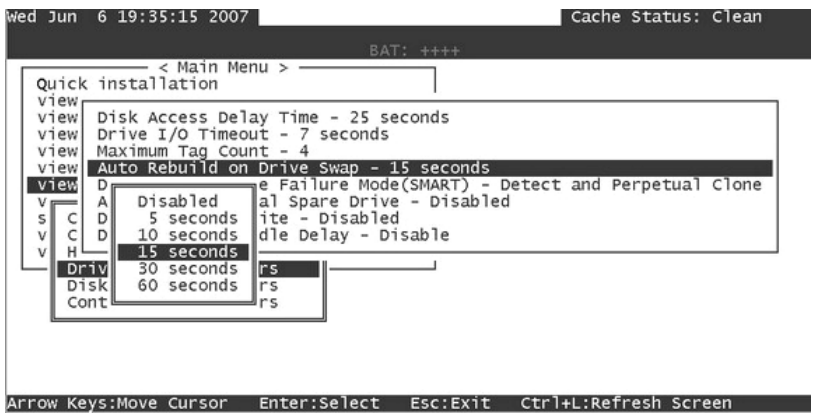

Step 1. Choose "**Auto Rebuild on Drive Swap**" from the "Drive-side Parameters" list, and press **[ENTER]**.

Step 2. Move your cursor bar to the desired interval; then press **[ENTER]**. Choose **Yes** in the dialog box that follows to confirm the setting.

IMPORTANT!

- The "**Auto Rebuild on Drive Swap**" timeout is enabled by choosing a time value. The RAID controller will poll all connected drives through the controller's drive channels at the assigned interval. Drive removal will be detected even if a host does not attempt to access data on that specific drive.

- If the "**Auto Rebuild on Drive Swap**" timeout is set to "Disabled" (the default setting is "Disabled"), the controller will not be able to detect any drive removal that occurs after the controller initialization process. The controller will only be able to detect drive removal when host access is directed to the drive side.

10.1.2 Auto-Assign Global Spare Drive

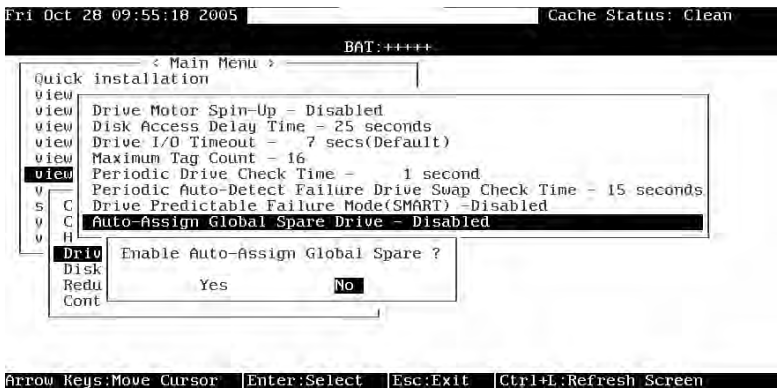

The "Auto-Assign" function automatically assigns any "new" drives that are not included in logical configurations as Global Spares.

Scenario:

Spare drives accelerate rebuild of a logical drive. In the example described below, multiple faults can occur at the same time making the array exposed to the risk of data loss:

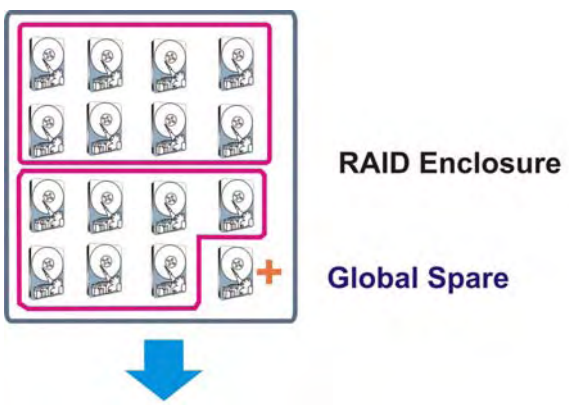

The **Auto-Assign Global Spare** feature is designed to reduce the chance of down time by operator's negligence. Shown on the left is a RAID enclosure with its drives configured into two arrays and a Global Spare. One logical drive consists of 8 members; the other consists of 7.

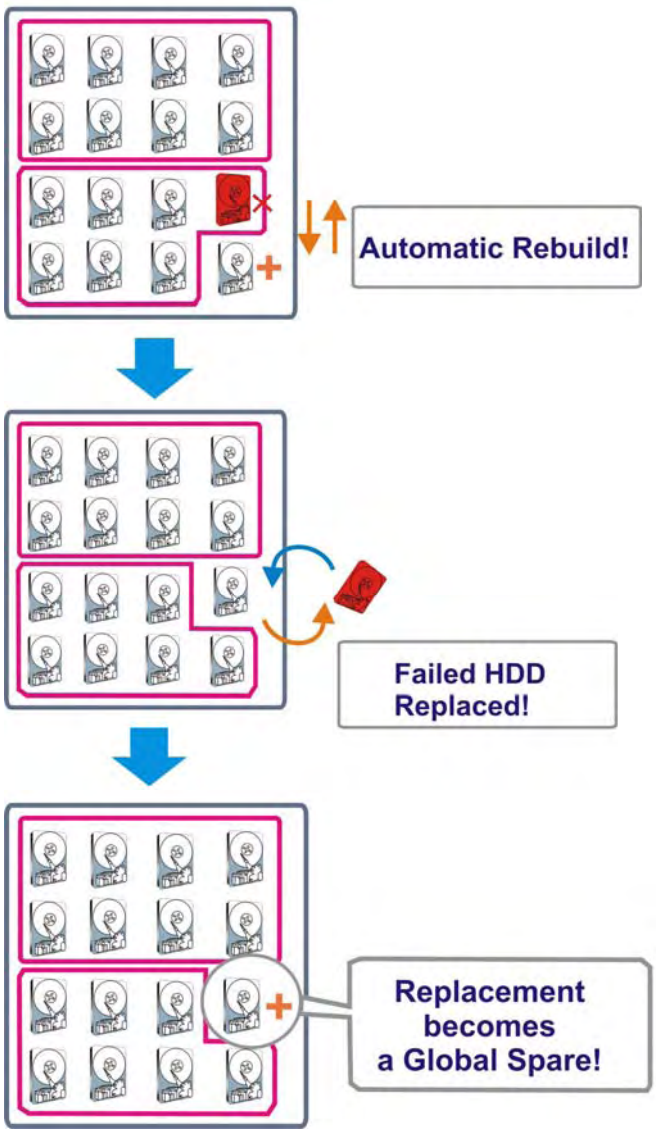

A member drive in one of the two logical drives fails. The Global Spare immediately participates in the rebuild.

The Failed drive is then replaced by a replacement drive.

The original Global Spare becomes a member of the 7-drive array.

With the Auto-Assign feature, firmware automatically configures the replacement drive as a Global Spare.

The Auto-Assign feature prevents the situation when a failed drive is replaced and the operator ***forgets*** to configure the replacement as another Global Spare leaving the array vulnerable to the occurrence of another drive failure.

Performance will also decrease if a failed drive is not immediately rebuilt. A considerable portion of system resources has to be conducted to generate data from the remaining members of the logical drive. If yet another member fails in a RAID 3 or RAID5 logical drive, data is lost.

NOTE:

- The **Auto-Assign Global Spare** applies to drive interfaces that support "auto detect," such as Fibre Channel, SATA, and SAS interfaces. Disk drives of these interfaces can be detected shortly after they are mated with the drive backplane.

Enabling the Function:

If a drive has a capacity smaller or apparently larger than the members of configured arrays, the controller may avoid using it as a global spare.

Enable the function and reset the controller for the configuration to take effect.

10.2 Scheduled Maintenance

Task Scheduler

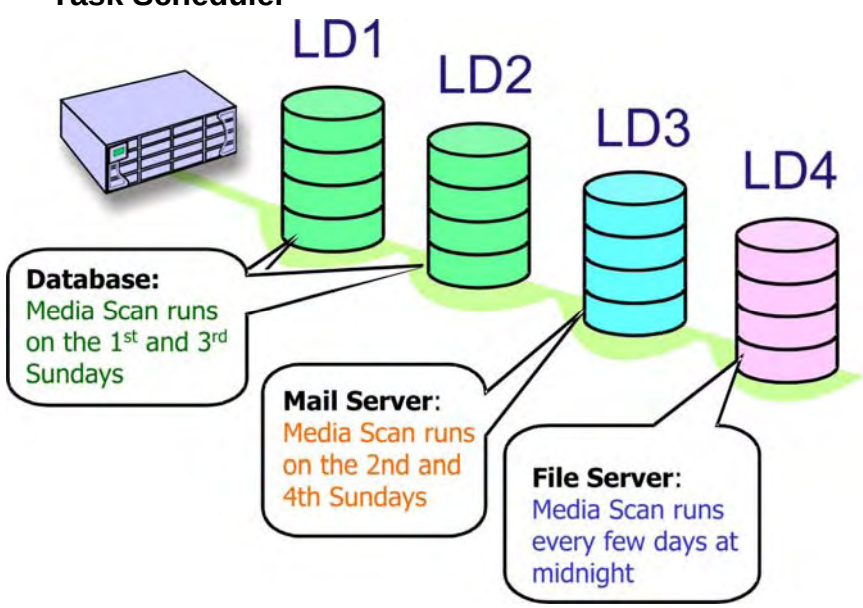

Figure 10-1 Scheduled Scans on Individual arrays

The Task Scheduler functionality allows Media Scans to be scheduled beginning at a specified start time and repeating at regular intervals defined by a configurable interval period. Each such schedule can be defined to operate on all drives of a certain class, all member drives of a specified logical drive, spare drives, or

all member drives of all logical drives. UIs supported are the text-based utility accessed through RS-232C serial connection/telnet and RAIDWatch GUI manager.

The Task Scheduler allows firmware to automatically perform media scans on specific RAID arrays saving you the efforts to manually initiate the processes. Scans take place at a preferred time when the subsystem is less stressed by daily service, e.g., Sundays or midnight.

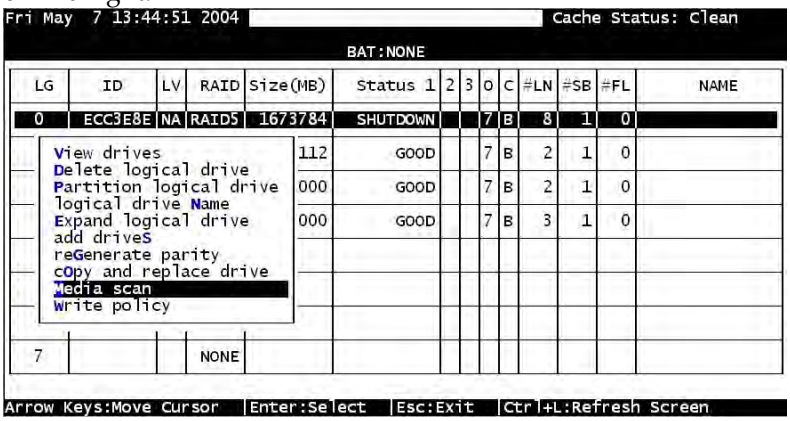

Step 1. Accessing the Task Scheduler Utility

Step 1-1. The scheduler utility can be accessed through the "**View and Edit Logical Drives**" menu by selecting a configured array and to display the list of array-related functions. Select "Media Scan" by pressing **[ENTER]**.

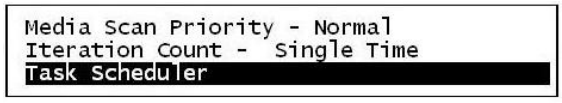

Step 1-2. Select "Task Scheduler" by pressing **[ENTER]**.

Step 2. Creating a New Schedule

Step 2-1. If there is no preset schedule, a confirm box will prompt.

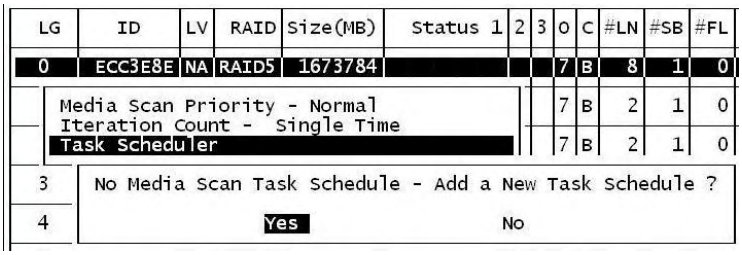

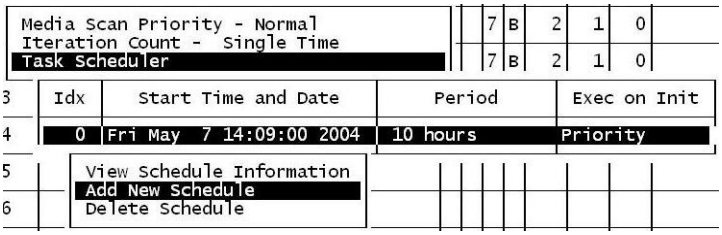

Step 2-2. Press **[ENTER]** on an existing schedule to display the configuration options. You may choose to check information of a task schedule, to create a new schedule, or to remove a configured schedule.

Step 3. Task Scheduler Options

To configure a task schedule, browse through the following options and make necessary changes:

Step 3-1. **Execute on Controller Initialization**

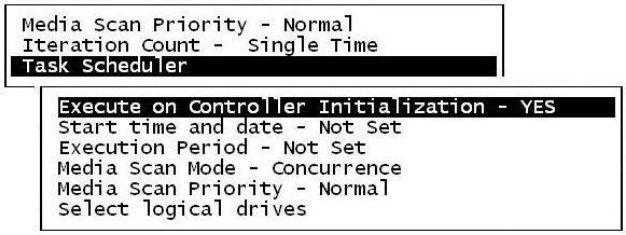

This option determines whether Media Scan is automatically conducted whenever the RAID system is reset or powered on.

Step 3-2. **Start Time and Date**

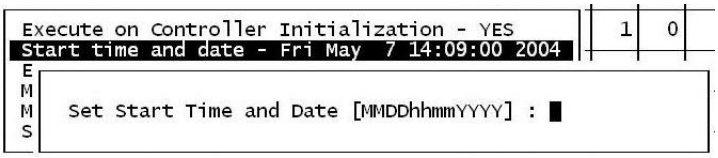

Enter time and date in its numeric representatives in the following order: month, day, hour, minute, and the year.

Step 3-3. Execution Period

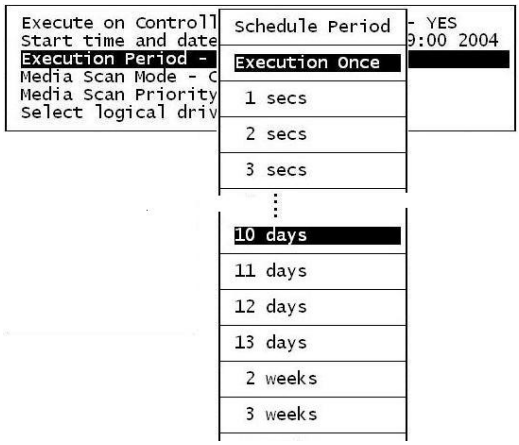

The scheduler memorizes the date and the time the actions are to be executed. Select one of the following:

- If the action is intended to be executed for one time only, select "Execution Once."

- In the case of a periodic action, the action is executed at the specified "start time," and then re-enacted at the time interval indicated in the execution period so as to be executed again later. The selectable interval ranges from one second to several weeks.

Step 3-4. Media Scan Mode

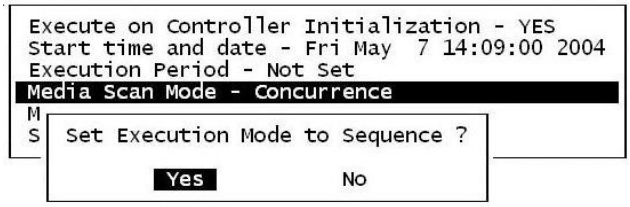

If the maintenance schedule includes more than one logical drive, the scan can be performed simultaneously on multiple logical drives together or separately on one logical drive at a time following a sequential order.

Step 3-5. Media Scan Priority

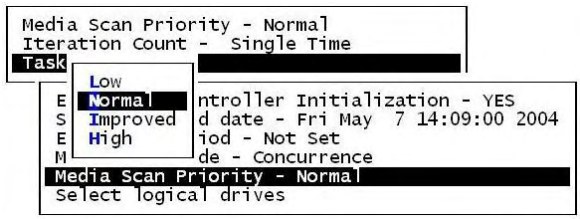

The scan priority determines how much of the system's resources will be consumed to perform the scheduled task. Select "Low" for better array performance and longer time to complete the media scan. Higher priority allows higher scan performance at the cost of reduced array performance.

Step 3-6. Select Logical Drives

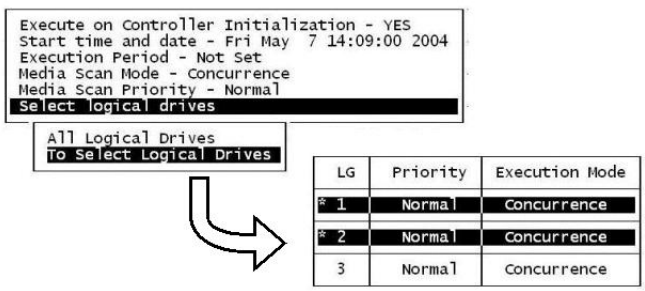

Press **[ENTER]** on "**Select Logical Drives**" to bring out a sub-menu. From there you may include all configured arrays or press **[ENTER]** on "**To Select Logical Drives**" to select one or more specific logical drive(s).

Logical drives can be tagged for inclusion by positioning the cursor bar on the logical drive and then pressing **[ENTER]**. An asterisk (*) mark will appear on the selected physical drive(s). To deselect the drive, press **[ENTER]** again on the selected drive. The "*" mark will disappear. Use the same method to select more logical drives if preferred.

When selection is done, press **[ESC]** to continue.

Step 4. Confirming the Creation of a Task Schedule

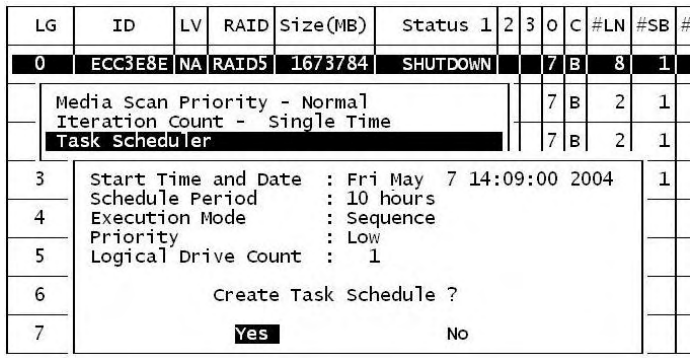

Step 4-1. When finished with setting the scheduler options, press **[ESC]** to display a confirm box.

Step 4-2. Verify all information in the box before choosing "**Yes"** to confirm and to complete the configuration process.

10.3 Manual Rebuild

If you want the controller to auto-detect a replacement drive, make sure you have a check time value set for the following option:

1. Auto Rebuild on Drive Swap check time

These two configurable options can be found under "**View and Edit Configuration Parameters**" -> "**Drive-side Parameters**."

IMPORTANT!

- A manual rebuild occurs in a subsystem that has no hot-spare.
- In a subsystem configured with hot-spares, the rebuild should take place automatically.
- The rebuild function will only appear if a logical drive (in RAID level 1, 3, 5, or 6) has a failed member.
- Carefully verify the location of a failed drive before replacement takes place. Removing the wrong drive will fatally fail a logical drive and the data loss will occur.

Step 1. Before physically replacing a failed drive, you should verify the messages as shown below:

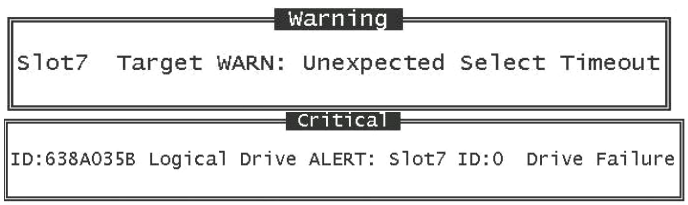

You should also check the logical drive member list in "**View and Edit Logical Drives**" -> "**View drives**." The failed drive's status should be indicated as "**BAD**."

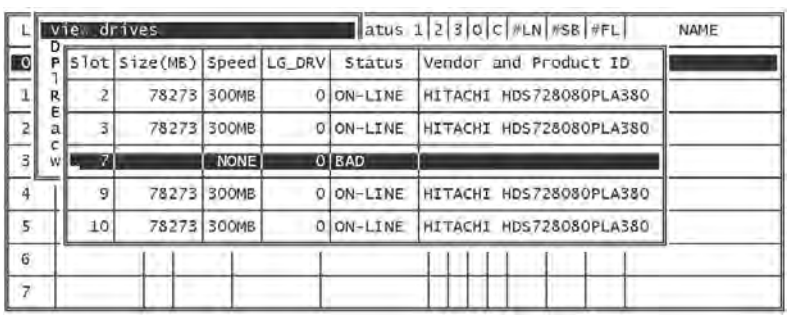

Step 2. Make sure you correctly identify the location of the failed drive and replace it with a new drive.

Step 3. Return to the "**View and Edit Logical Drives**" menu. Press [**ENTER**] on it and you should find the "**Rebuild**" option.

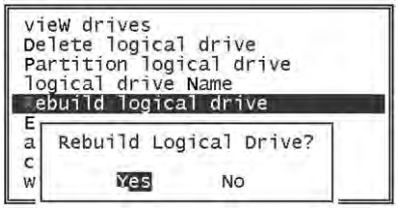

Step 4. The rebuild should start. Press **ESC** to skip the message.

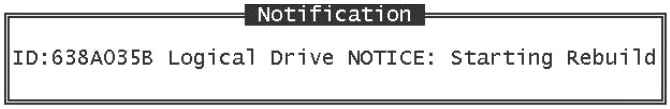

Step 5. The rebuild progress will be indicated by a status bar.

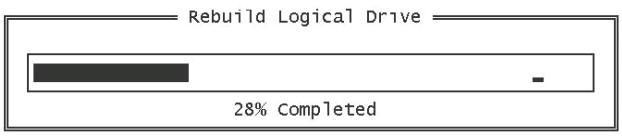

Step 6. Upon the completion of rebuild, the following message will prompt. Press **ESC** to dismiss the message.

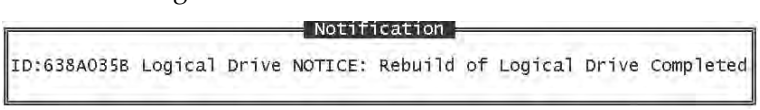

Step 7. You may now return to the "**View and Edit Logical Drives**" menu to check if the array status is stated as "**GOOD**."

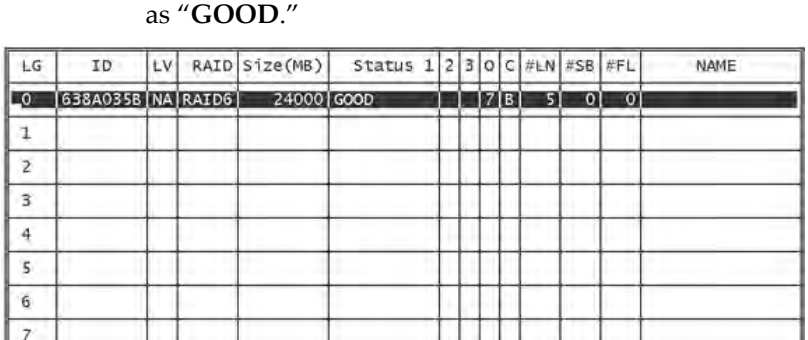

LG	ID	LV	RAID	Size(MB)	Status	1	2	3	0	C	#LN	#SB	#FL	NAME
0	638A035B	NA	RAID6	24000	GOOD				7	B	5	0	0	
1														
2														
3														
4														
5														
6														
7														

10.4 Regenerating Logical Drive Parity

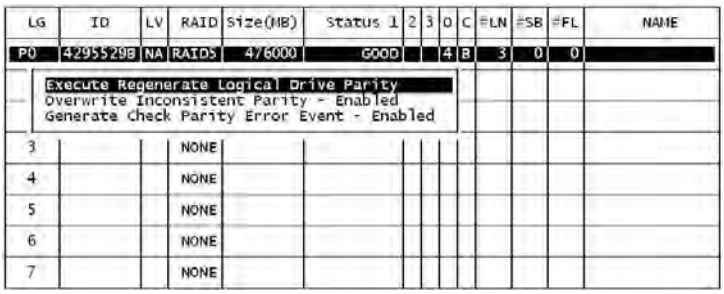

Parity regeneration is a function manually performed onto RAID-1/3/5/6 arrays to determine whether inconsistency has occurred with data parity.

You may perform the parity check directly without changing the two options below, or set preferred options and then press **[ENTER]** on "**Execute Regenerate Logical Drive Parity**" to begin the operation.

10.4.1 Overwrite Inconsistent Parity

Default is "enabled."

If an array's data parity is seriously damaged, restoring parity data by regenerating and overwriting the original data may cause data loss. Disable this option if you suspect parity data has been seriously corrupted.

10.4.2 Generate Check Parity Error Event

Default is "enabled."

When enabled, parity inconsistency will be reported as system events.

IMPORTANT!

If a regenerating process is stopped by a drive failure, the process cannot be restarted until the logical drive is successfully rebuilt by having its failed member replaced.

10.5 Disk Array Parameters

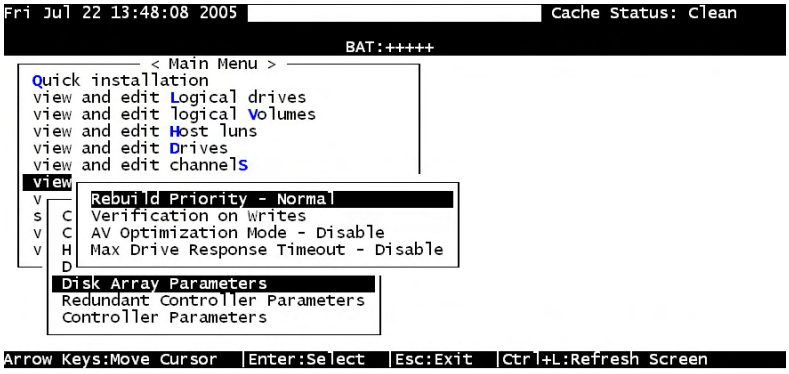

Step 1. Select "**View and Edit Configuration Parameters**" on the Main Menu and press **[ENTER].**

Step 2. Choose "**Disk Array Parameters**," then press **[ENTER]** again. The Disk Array Parameters menu will appear.

10.5.1 Rebuild Priority

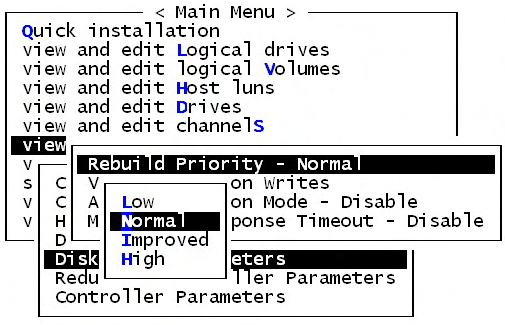

Step 1. Choose "Rebuild Priority," then press **[ENTER]**. A list of the priority selections (Low, Normal, Improved, or High) displays.

Step 2. Move the cursor bar to a selection, then press **[ENTER]**.

10.5.2 Verification on Writes

Errors may occur when a hard drive writes data. To avoid the write error, the controller can force hard drives to verify written data. There are three selectable methods:

- Verification on LD Initialization Writes
  Performs Verify-after-Write when initializing a logical drive

- Verification on LD Rebuild Writes
  Performs Verify-after-Write during the rebuild process

- Verification on LD Normal Drive Writes
  Performs Verify-after-Write during normal I/Os

Each method can be enabled or disabled individually. Hard drives will perform Verify-after-Write according to the selected method.

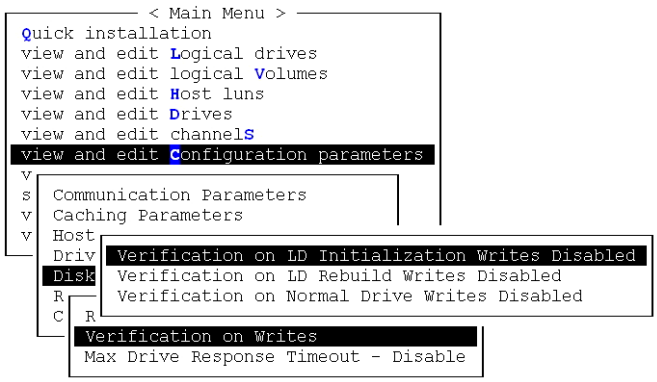

Step 1. Move the cursor bar to the desired item, then press **[ENTER]**.

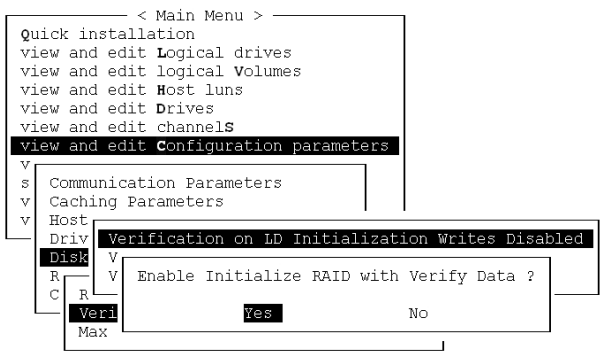

Step 2. Choose **Yes** in the confirm box to enable or disable the function. Follow the same procedure to enable or disable each method.

IMPORTANT!

The "**verification on Normal Drive Writes**" method will affect the "write" performance of your RAID system.

10.5.3 Maximum Drive Response Time

In situations such as drive failure or the occurrences of media errors, a read or write request returned after several hundreds milliseconds will be too long for AV applications for which choppy audio or dropped video frames are not acceptable.

Response Time in Read Scenarios

The Maximum Response Time option provides a timeout value for processing read/write requests. If delays caused by media errors are reported on a specific member of an array, the subsystem firmware immediately retrieves data by generating data from RAID parity and the data blocks on other members of the array. In this way, delays on read requests can be efficiently eliminated. Without the Response Time setting, firmware may wait several seconds for the hard drive to timeout.

Maximum Drive Response Time in Write Scenarios

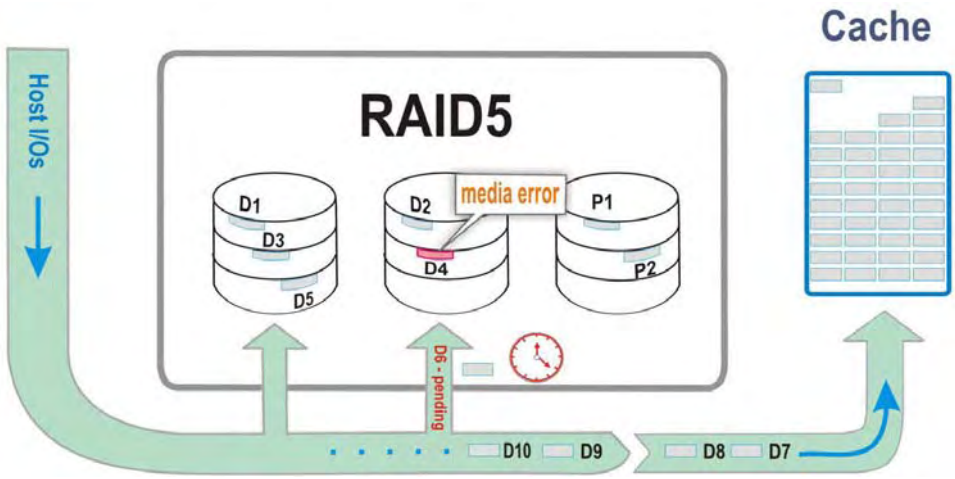

Figure 10-1: Delays on Write by Media Error

As shown above, the occurrences of media errors on a single disk drive can cause a performance drag within a few seconds. If media errors occur while servicing write requests, the following can occur:

1. A media error is encountered while RAID system firmware is conducting a write request (D4: data block #4).
2. It usually takes 3 to 4 seconds for a hard drive to return timeout state, and during that time the succeeding write requests (data blocks D7, D8, and onward) will be cached in system buffer and quickly fill the data cache.
3. Supposed the data cache capacity is 512MB, it is easily used up when hundreds of megabytes of write requests come streaming down from the application server.
4. When the cache is full, performance is quickly reduced and the benefit of write-back caching soon vanishes.

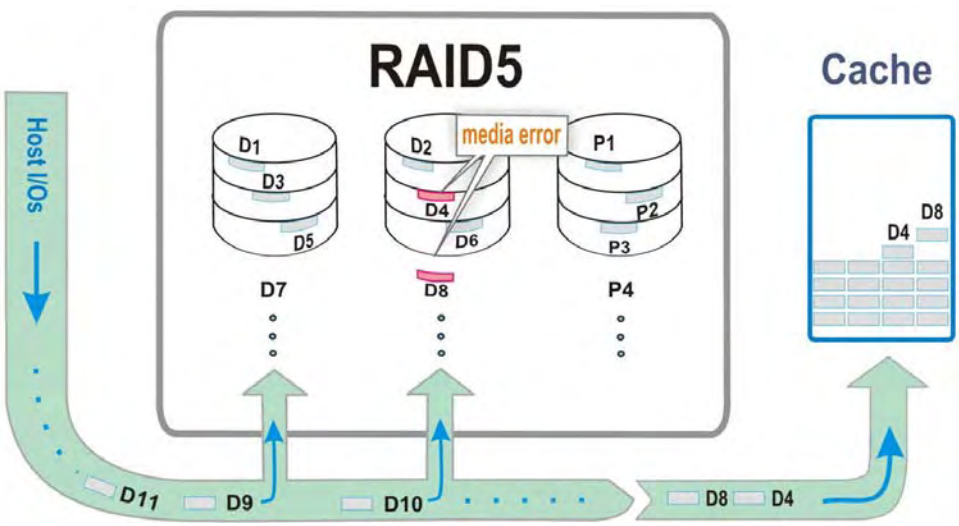

Figure 10-2: Caching the Affected Writes Only

The **Response Time** remedy is described as follows:

1. A response delay time is set in firmware: default is 160ms.
2. If a single disk drive cannot fulfill a write request within 160ms, firmware automatically proceeds with conducting write requests on other disk drives while also generating parity data.
3. **Only those writes** affected by media errors on an individual disk drive will be cached in memory so that the data cache will not be quickly overwhelmed. The data cache holds a comparatively small portion of write requests. If a logical drive contains 8 members, one of them is parity drive and media

errors are found on one member drive, caching data blocks to one disk drive only occupies 1/8 of cache capacity.

4. With the response time on write, RAID subsystems can ensure array performance with the occurrences of media errors without waiting for physical hard drives to resolve hardware errors.
5. If the drive carrying media errors does fail afterwards, data blocks cached in memory will be dumped and the rebuild begins.

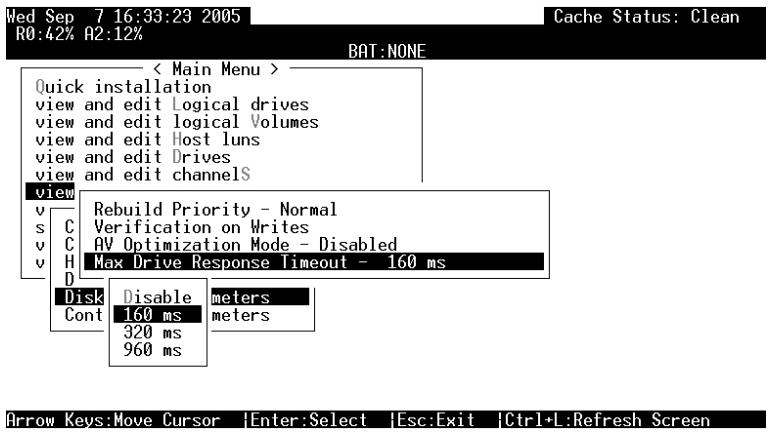

The **Response Time** option can be found in "**View and Edit Configuration Parameters**" -> "**Disk Array Parameters**" -> "**Max Drive Response Timeout.**"

Other Concerns:

To prepare the array for read-intensive applications, the following are recommended:

1. Default timeout as 160ms.
2. Arrays should not be partitioned.
3. The priorities for **Rebuild** or **Media Scan** operations should be set to "low."
4. Another timeout value, the "**Drive I/O Timeout**" which determines whether a drive has eventually failed to respond to I/O requests, is required as the first-level timeout.

Chapter

11

Array Expansion

The array expansion functions allow you to expand storage capacity without the cost of buying new equipment. Expansion can be completed on-line while the system is serving host I/Os.

11.1 Overview

What is RAID Expansion and how does it work?

Before the invention of RAID Expansion, increasing the capacity of a RAID system meant backing up all data in the disk array, re-creating the disk array configuration with new drives, and then restoring data back into system.

Infortrend's RAID Expansion technology allows users to expand a logical drive by adding new drives, or replacing drive members with drives of larger capacity. Replacing is done by copying data from the original members onto larger drives; the smaller drives can then be replaced without powering down the system.

11.2 Notes on Expansion

1. Expansion Capacity:

When a new drive is added to an existing logical drive, the capacity brought by the new drive appears as a new partition. For example, if you have 4 physical drives (36GB each) in a logical drive, and each drive's maximum capacity is used, the capacity of the logical drive will be 108GB. (One drive's capacity is used for parity, e.g., RAID 3). When a new 36GB drive is

added, the capacity will be increased to 144GB in two separate partitions (one is 108GB and the other 36GB).

CAUTION!

1. If an array has not been partitioned, the expansion capacity will appear as an added partition, e.g., partition 1 next to the original partition 0.
2. If an array has been partitioned, the expansion capacity will be added behind the last configured partition, e.g., partition16 next to the previously-configured 15 partitions.
3. If an array has been partitioned by the maximum number of partitions allow, the expansion capacity will be added to the last partition, e.g., partition 63. Partition change will invalidate data previously stored in the array.
4. See the diagram below for the conditions that might occur during array expansion.

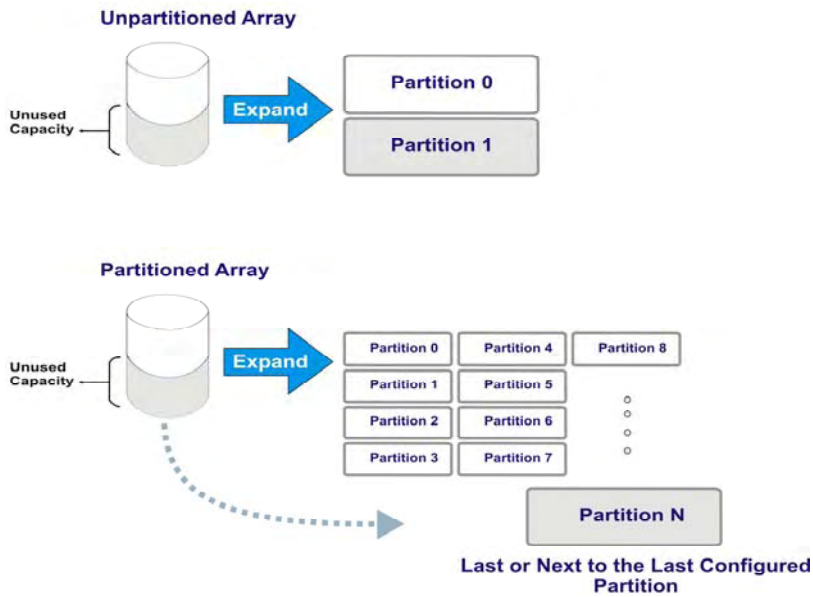

2. Size of the New Drive:

A drive used for adding capacity should have the same or more capacity as other drives in the array.

3. Applicable Arrays:

Expansion can only be performed on RAID 0, 1, 3, 5, and 6 logical drives. Expansion cannot be performed on logical configurations that do not have parity, e.g., NRAID or RAID 1.

NOTE:

Expansion on RAID0 is not recommended, because the RAID0 array has no redundancy. Interruptions during the expansion process may cause unrecoverable data loss.

4. Interruption to the Process:

Expansion should not be canceled or interrupted once begun. A manual restart should be conducted after the occurrence of a power failure or interruption of any kind.

Expand Logical Drive: Re-striping

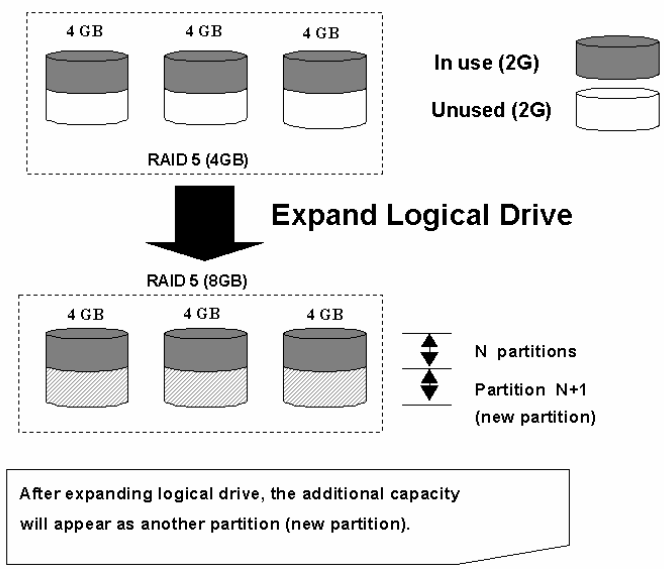

Figure 11 - 1: Logical Drive Expansion

RAID levels supported: **RAID 0, 1, 3, 5 and 6**

Expansion can be performed on logical drives or logical volumes under the following conditions:

1. There is unused capacity in a logical unit

2. Capacity is increased by using member drives of larger capacity (see Copy and Replace in the discussion below)

Data is recalculated and distributed to drive members or members of a logical volume. Upon the completion of the process, the added or the previously unused capacity will become a new partition. The new partition must be made available through host LUN mapping in order for a host adapter to recognize its presence.

11.3 Mode 1 Expansion: Adding Drives to a Logical Drive

Use drives with the same capacity as the original drive members. Once completed, the added capacity will appear as another partition (new partition). Data is automatically re-striped across the new and old members during the add-drive process. See the diagram below to get a clear idea:

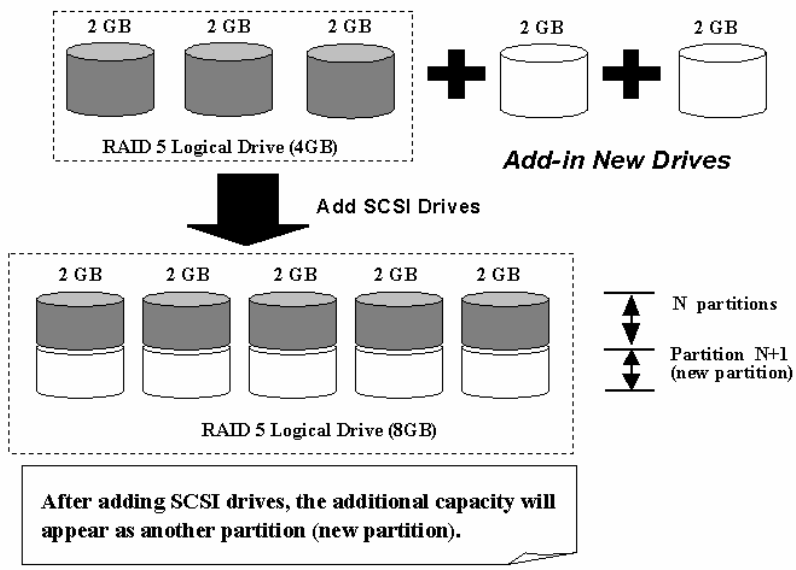

Figure 11 - 2: Expansion by Adding Drive

RAID levels supported: **RAID 0, 1, 3, 5, and 6.**

The new partition must be made available through a host ID/LUN.

Add Drive Procedure

Step 1. First select from the Main Menu, "**View and Edit Logical Drive**," and select a logical drive to add a new drive to.

NOTE:

The drive selected for adding should have a capacity no less than the original member drives. If possible, use drives of the same capacity because all drives in the array are treated as though they have the capacity of the smallest member in the array.

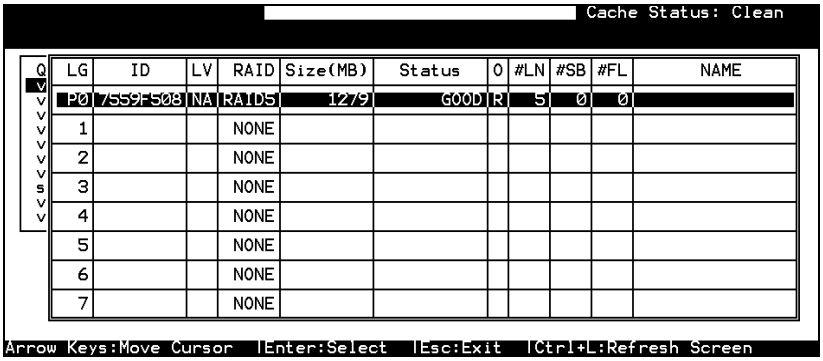

Step 2. Press **[ENTER]** to select a logical drive and choose "**Add Drives**" from the submenu. Proceed with confirming the selection.

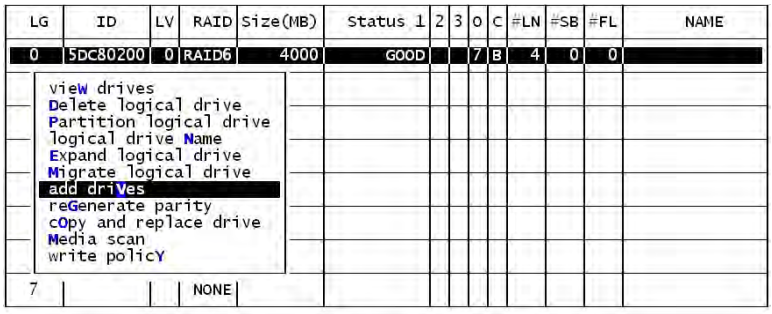

Step 3. Available drives will be listed. Select one or more disk drive(s) to add to the target logical drive by pressing **[ENTER].** The selected drive will be indicated by an asterisk "*" mark.

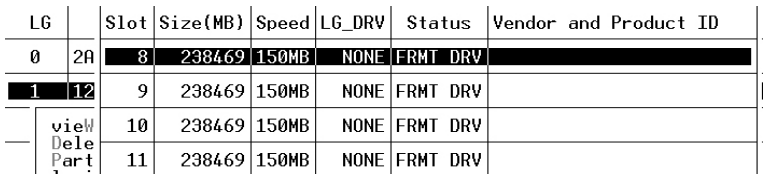

Step 4. Press **[ESC]** to proceed and the notification will prompt.

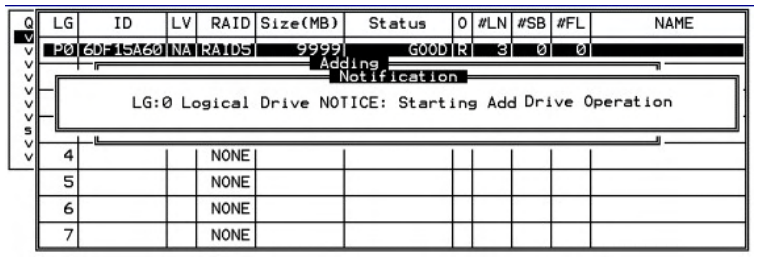

Step 5. Press **[ESC]** again to cancel the notification prompt; a status bar will indicate the percentage of progress.

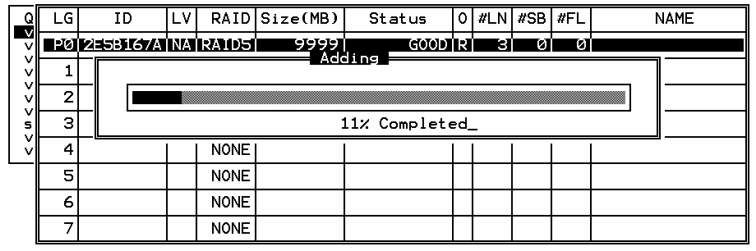

Step 6. Upon completion, there will appear a confirming notification. The capacity of the added drive will appear as an unused partition.

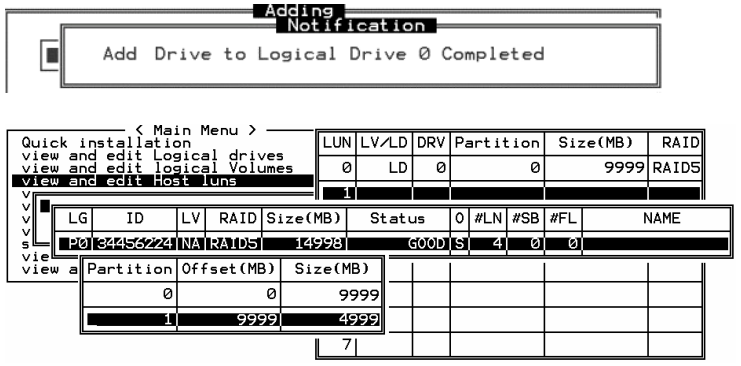

The added capacity will be automatically included, meaning that you do not have to "**expand logical drive**" later. Map the added capacity to another host ID/LUN to make use of it.

As diagrammed above, in "**View and Edit Host LUN**," the original capacity is 9999MB, its host LUN mapping remains unchanged, and the added capacity appears as the second partition.

IMPORTANT!

- Expansion by adding drives can not be canceled once started. If power failure occurs, the expansion will be paused and the controller will NOT restart the expansion when power comes back on. Resumption of the RAID expansion must be performed manually.

- If a member drive of the logical drive fails during RAID expansion, the expansion will be paused. The expansion will resume after the logical drive rebuild is completed.

11.4 Mode 2 Expansion:

Copy and Replace Drives with Drives of Larger Capacity

You may also expand your logical drives by copying and replacing all member drives with drives of higher capacity. Please refer to the diagram below for a better understanding. The existing data in the array is copied onto the new drives, and then the original members can be removed.

When all the member drives have been replaced, execute the "**Expand Logical Drives**" function to make use of the expansion capacity.

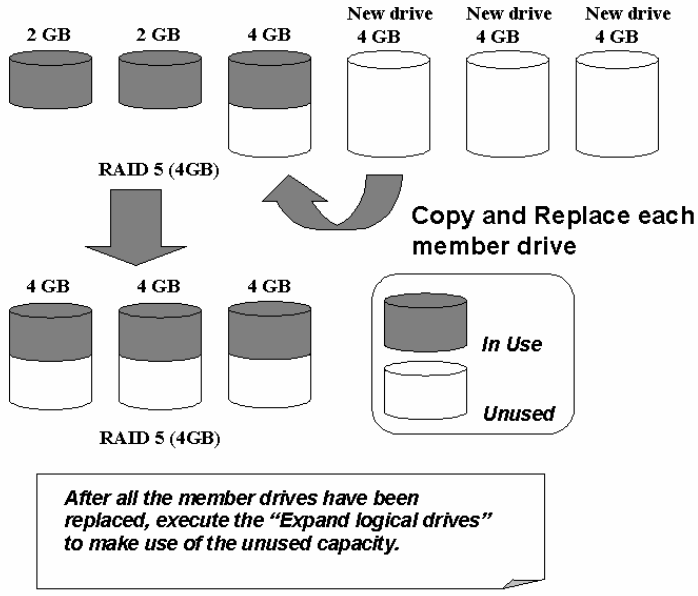

Figure 11 - 3: Expansion by Copy & Replace

RAID levels supported: **RAID 0, 1, 3, 5 and 6**

Copy and Replace Procedure

Step 1. Select from Main Menu "**View and Edit Logical Drives**." Select a target array, press **[ENTER]** and scroll down to choose "**Copy and Replace Drive**." Press **[ENTER]** to proceed.

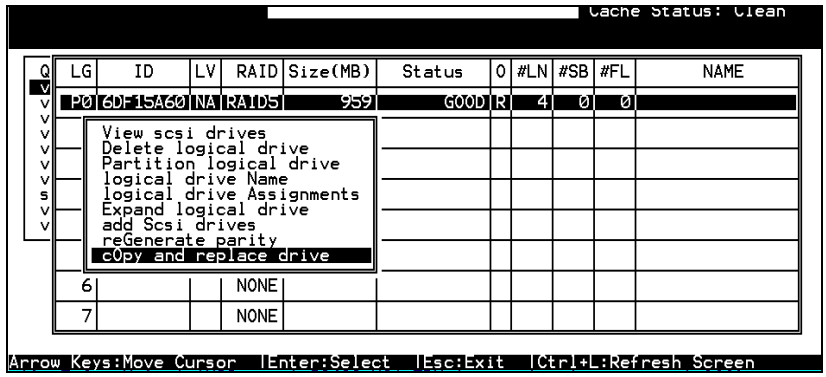

Step 2. The array members will be listed. Select the member drive (the source drive) you want to replace with a larger one.

Step 3. Select one of the members as the "**source drive**" (status indicated as **ON-LINE**) by pressing **[ENTER]**; a table of available drives will prompt.

Step 4. Select a "**new drive**" to copy the capacity of the source drive to. The channel number and ID number of both the "**Source Drive**" and the "**Destination Drive**" will be indicated in the confirm box.

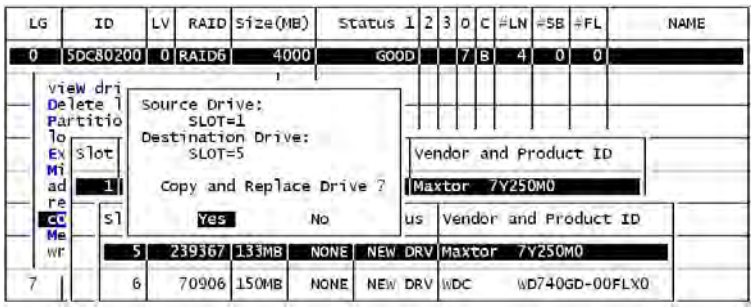

Step 5. Choose **Yes** to confirm and proceed.

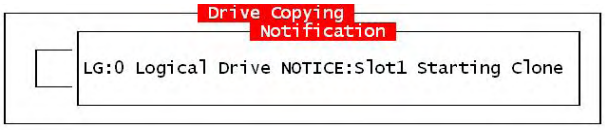

Step 6. Press **[ESC]** to view the progress.

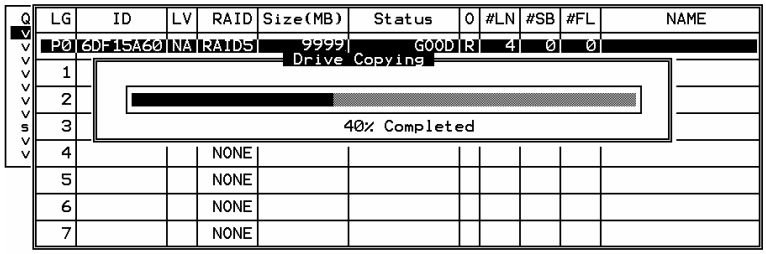

Step 7. Completion of the Copy and Replace process will be indicated by a notification message. Follow the same method to copy and replace every member drive. You may now perform "**Expand Logical Drive**" to make use of the added capacity, and then map the additional capacity to a host LUN.

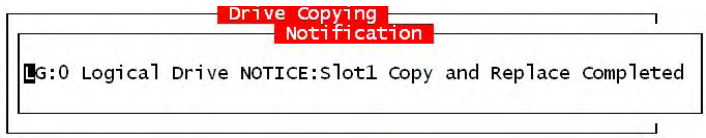

11.5 Making Use of the Added Capacity: Expand Logical Drive

In the following example, the logical drive is originally composed of three member drives and each member drive has the capacity of 1GB. "**Copy and Replace**" has been performed on the logical drive and each of its member drives has been replaced by a new drive with the capacity of 2GB. The next step is to perform "**Expand Logical Drive**" to utilize the additional capacity brought by the new drives.

Step 1. Select "**View and Edit Logical Drives**" from the Main Menu and select the logical drive with its members copied and replaced.

Step 2. Select "**Expand Logical Drive**" in the sub-menu and press **[ENTER]** to proceed. A confirming box will appear.

Step3. Proceed by pressing **[ENTER]** or entering any value no larger than the "maximum drive expand capacity" and press **[ENTER]**.

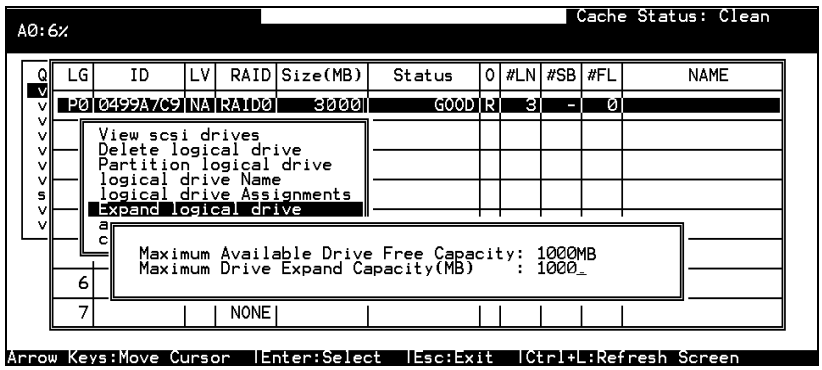

Step 4. Choose **Yes** to confirm.

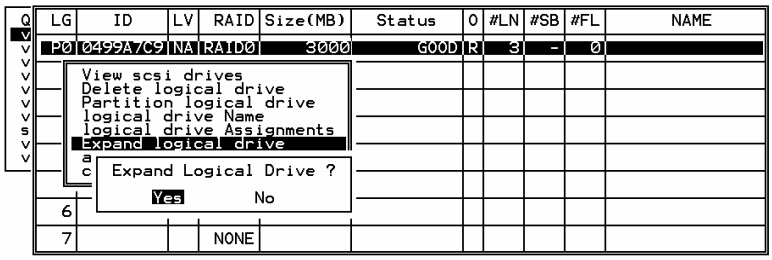

Step 5. Upon completion, you will be prompted by the notification message.

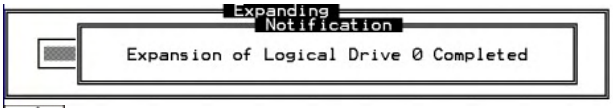

Step 6. Press **[ESC]** to return to the previous menu screen.

As shown below, the total capacity of logical drive has been expanded to 6GB.

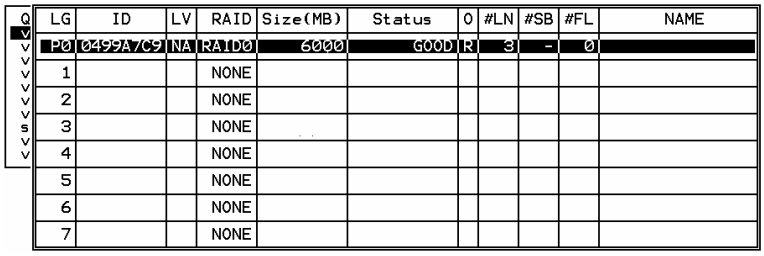

11.6 Expand Logical Volume

NOTE:

If the logical drive that has an expanded capacity is a member of a logical volume, make sure you expand all logical drives within the logical volume. A logical volume is made of logical drives that are "striped" together. Unless all logical drives within a logical volume have excessive capacity, you cannot expand a logical volume.

Step 1. To expand a logical volume, expand its logical drive member(s) and then perform "**Expand logical volume**."

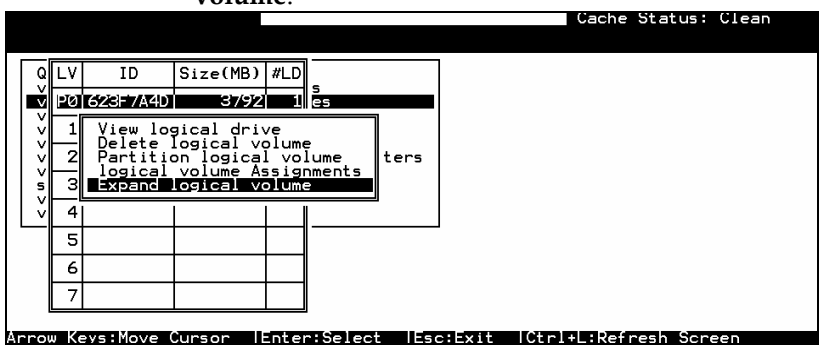

Step 2. When prompted by "**Expand Logical Volume?**", choose **Yes** to confirm and the process will be completed immediately.

11.7 Configuration Example: Volume Extension in Windows 2000®

Limitations When Using Windows 2000

1. These limitations apply only to the Windows NT Server or Windows 2000 Server Disk Management which includes the Extend Volume Set function; Windows NT Workstation does not support this feature. The Volume Set Expansion formats the new area without affecting files existing on the original volume.

2. The system drive (boot drive) of a Windows NT/2000 system cannot be expanded.

3. The drive to be expanded should be using the NTFS file system.

Example

The following example demonstrates the expansion of a 16988MB RAID 5 logical drive. The HyperTerminal emulation software that comes with Windows Server is used to connect to the RAID controller via RS-232C.

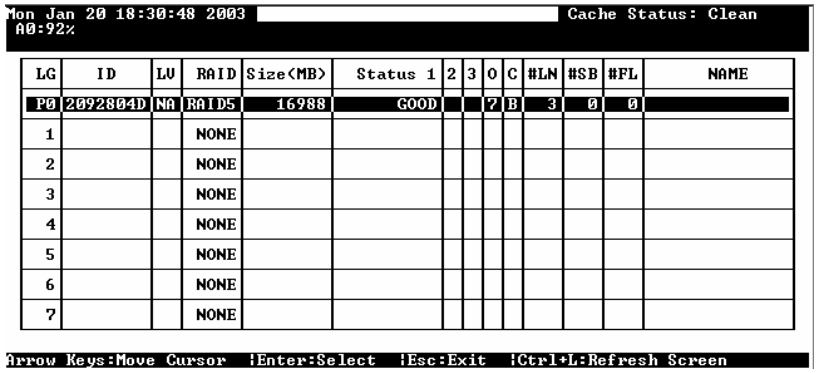

You can view information about this drive in the Windows 2000 Server's **Computer Management** -> **Storage** -> **Disk Management**.

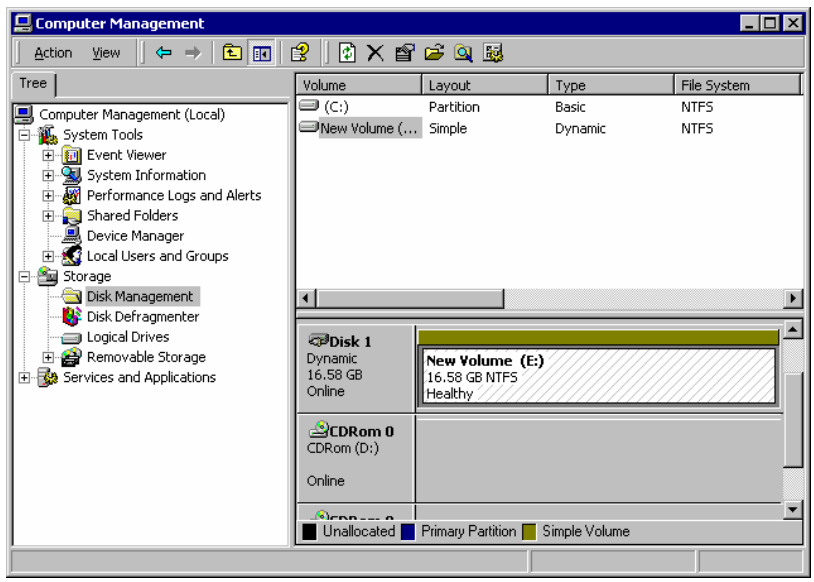

Place the cursor on Disk 1, right-click your mouse, and select "**Properties**." You will see that the total capacity for the Drive E: is about 16.5GB.

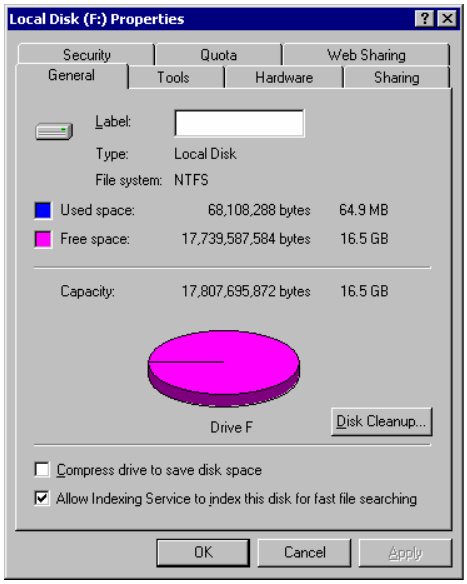

Follow the steps described in the previous section to "**add**" or "**copy & replace**" disk drives and perform "**Logical Drive Expansion**."

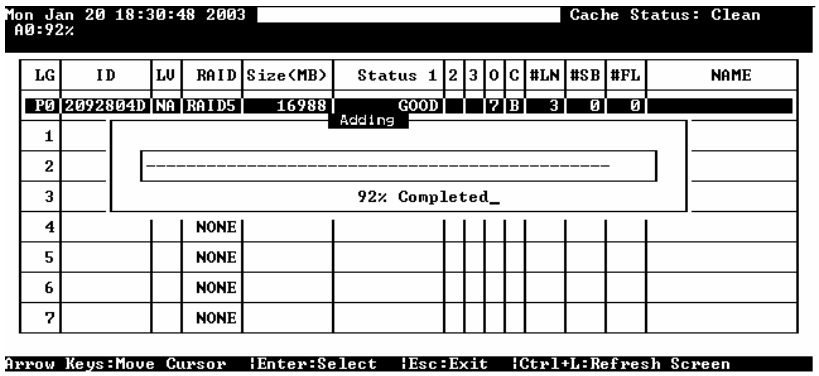

The 16.5GB logical drive has become a 25GB logical drive. Place the cursor on that logical drive, and then press **[ENTER]**.

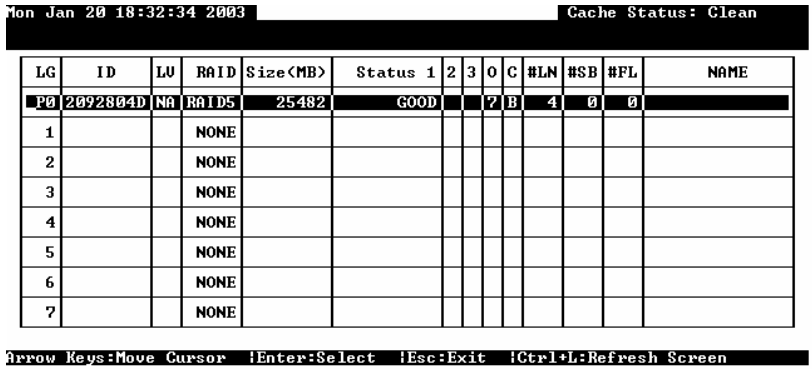

From the menu, select "Partition Logical Drive." You will see that the 25GB logical drive is composed of a 17GB partition and an 8.4GB partition.

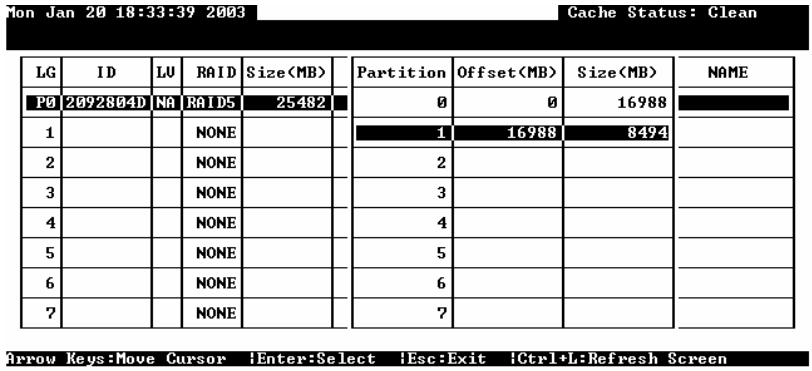

Map the new partition to a host LUN. The new partition must be "**mapped**" to a host LUN in order for the HBA (host-bus adapter) to recognize its presence. Once you have mapped the partition, reboot your Windows server. The HBA should be able to detect an additional disk during the initialization process.

Return to Windows 2000 Server's Disk Management. There now exists a Disk 2 with 8.3GB of free space. You may use the "rescan disks" command to bring up the new drive.

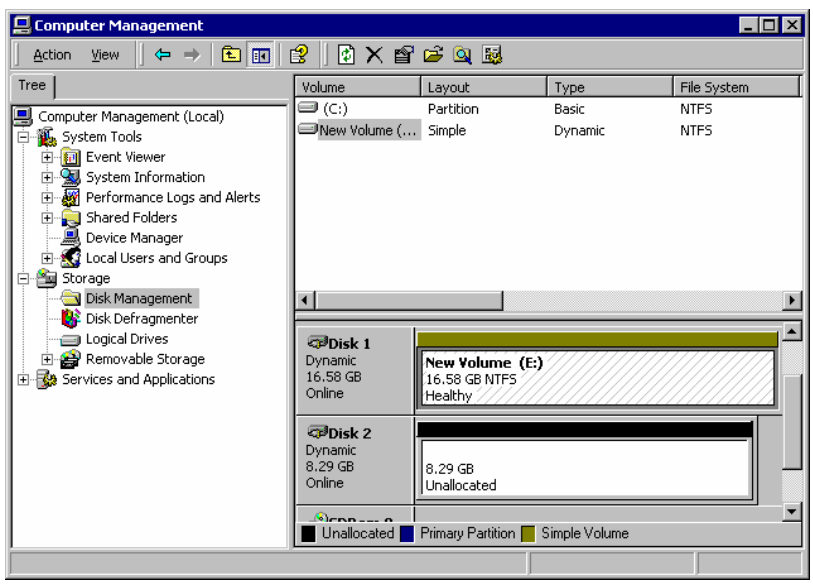

Select an existing volume (Disk1) and then right-click on the disk column. Select "**Extend Volume**" to proceed.

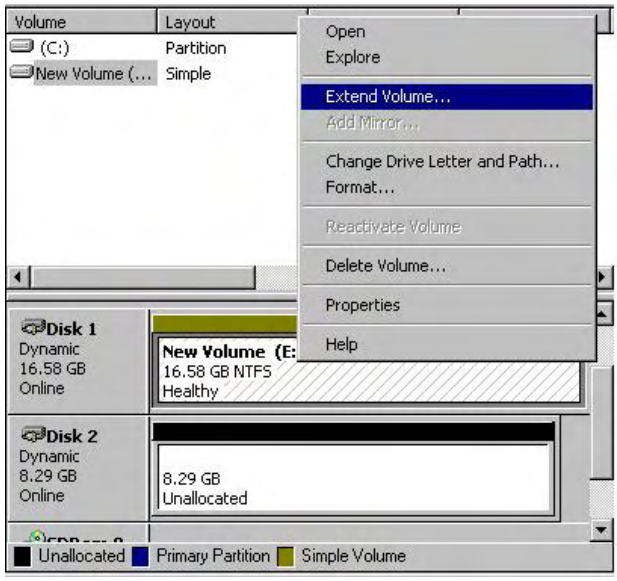
Volume
Layout
(C:)
Partition
New Volume (...
Simple
Open
Explore
Extend Volume...
Add Mirror...
Change Drive Letter and Path...
Format...
Reactivate Volume
Delete Volume...
Properties
Help
Disk 1
Dynamic
16.58 GB
Online
New Volume (E:
16.58 GB NTFS
Healthy
Disk 2
Dynamic
8.29 GB
Online
8.29 GB
Unallocated
Unallocated
Primary Partition
Simple Volume

The **Extend Volume Wizard** will guide you through the rest of the process.

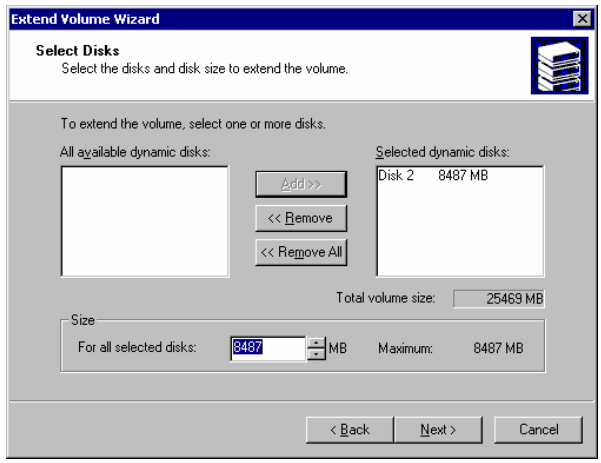

The screen will display that the volume set of Drive E: has been extended into a spanned volume by the 8.3GB in Disk2.

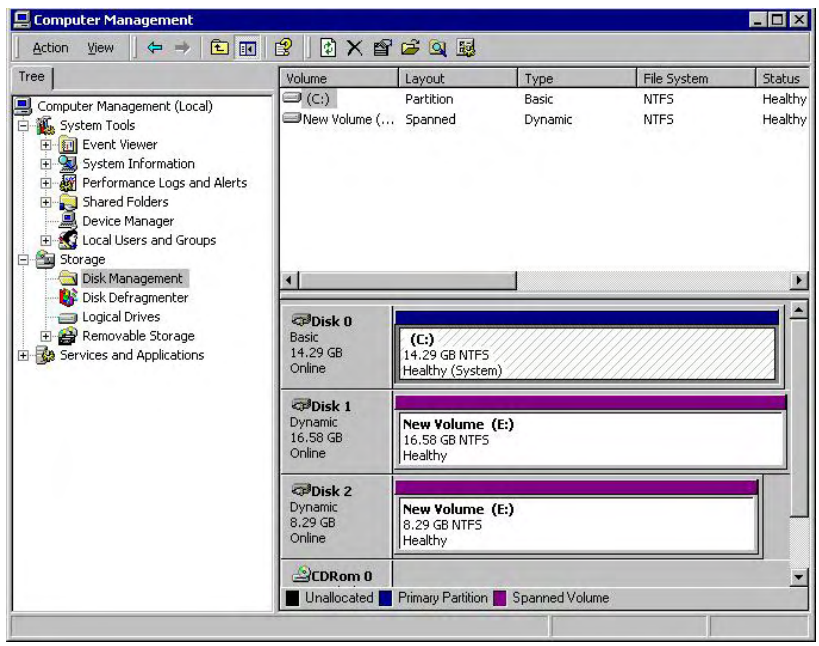

Logical Drive E: is now composed of two partitions with a total volume of 2500MB. To see this, hold down on the <Ctrl> key and select both Disk 1 and Disk2; then right-click your mouse and select "Properties."

Drive E: now has a capacity of about 25GB.

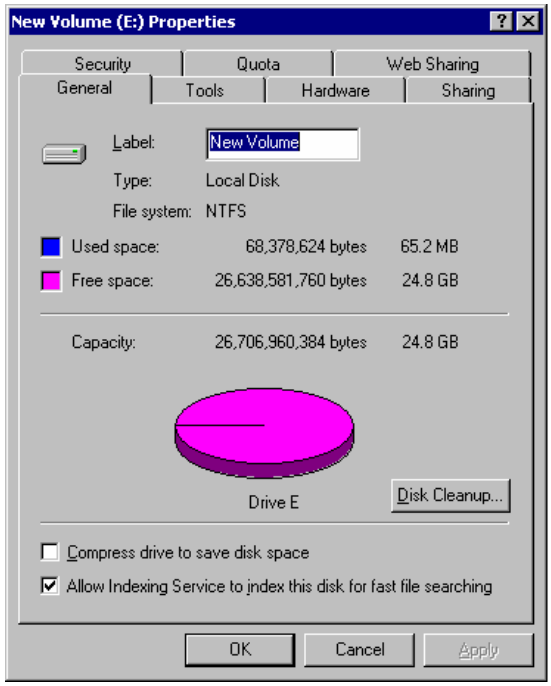
New Volume (E:) Properties
Security
Quota
Web Sharing
General
Tools
Hardware
Sharing
Label:
New Volume
Type:
Local Disk
File system:
NTFS
Used space:
68,378,624 bytes
65.2 MB
Free space:
26,638,581,760 bytes
24.8 GB
Capacity:
26,706,960,384 bytes
24.8 GB
Drive E
Disk Cleanup...
Compress drive to save disk space
Allow Indexing Service to index this disk for fast file searching
OK
Cancel
Apply

This page is intentionally left blank

Chapter

12

S.M.A.R.T. Support

12.1 Overview

With the maturity of technologies like S.M.A.R.T., drive failures can be predicted to certain degree. Before S.M.A.R.T., receiving notifications of drive bad block reassignments may be the most common omen that a drive is about to fail. In addition to the S.M.A.R.T.-related functions as will be discussed later, a system administrator can also choose to manually perform "Clone Failing Drive" on a drive which is about to fail.

This function provides system administrators a choice on when and how to preserve data from a failing drive. Although not necessary under normal conditions, you may also replace any drive at-will even when the source drive is healthy.

The "**Clone Failing Drive**" can be performed under the following conditions:

1. Replacing a failing drive either detected by S.M.A.R.T. or notified by the controller.

2. Manually replacing and cloning any drive with a new drive.

12.2 Clone Failing Drive

Unlike the similar functions combined with S.M.A.R.T., the "**Clone Failing Drive**" is a manual function. There are two options for cloning a failing drive: "**Replace after Clone**" and "**Perpetual Clone**."

12.2.1 Replace after Clone

Data on the source drive, the drive with predicted errors (or any selected member drive), will be cloned to a standby spare and replaced later by the spare. The status of the replaced drive, the original member drive with predicted errors, will be redefined as a "**used drive.**" System administrators may replace the "used drive" with a new one, and then configure the new drive as a spare drive.

Step 1. Locate the logical drive to which the drive with predictable errors belongs. Select the "**Clone Failing Drive**" function.

Step 2. Select "**Replace After Clone**." The controller will automatically start the cloning process using the existing "**stand-by**" (dedicated/global spare drive) to clone the source drive (the target member drive with predicted errors). If there is no standby drive (local/global spare drive), you need to add a new drive and configure it as a standby drive.

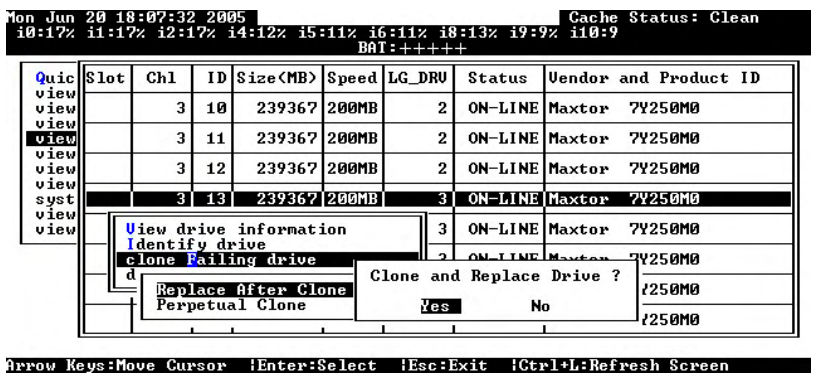

Step 3. The cloning process will begin with a notification message. Press **[ESC]** to proceed.

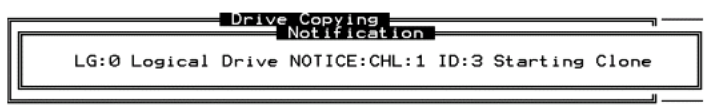

Step 4. The cloning process will be indicated by a status bar.

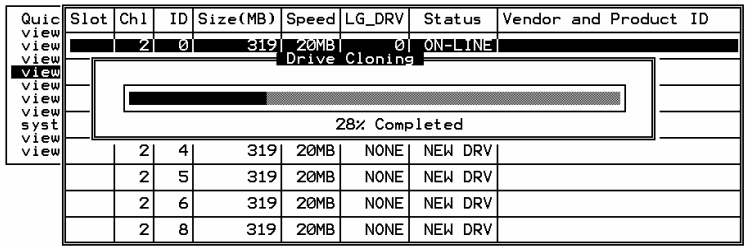

Step 5. You may quit the status bar by pressing **[ESC]** to return to the table of the connected drives. Select the drive indicated as "**CLONING**" by pressing **[ENTER]**.

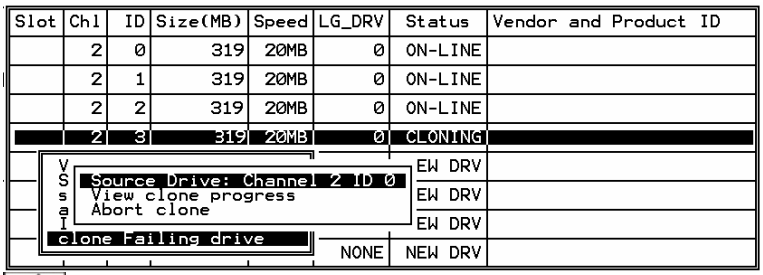

Step 6. Select "**Clone Failing Drive**" again to view the current status. You may identify the source drive and choose to "**View Clone Progress**," or "**Abort Clone"** if you happen to have selected the wrong drive.

Step 7. When the process is completed, you will be notified by the following message.

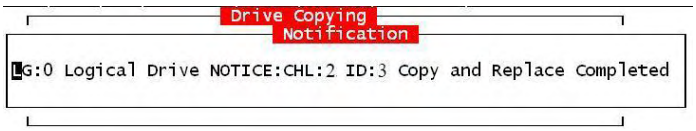

12.2.2 Perpetual Clone

The standby spare will clone the source drive, the member drive with predicted errors or any selected drive, without substituting it. The status of the spare drive will be displayed as "**clone drive**" after the cloning process. The source drive will remain a member of the logical drive. If the source drive fails, the clone drive can readily take its place in the array.

Step 1. In "**View and Edit Drives**," locate the member drive that shows predicted errors. Select "**Clone Failing Drive**," and choose "**Perpetual Clone**."

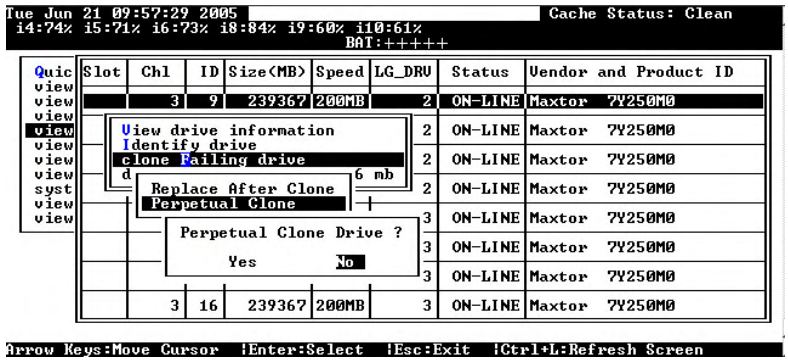

Step 2. The controller will automatically start the cloning process using the existing "**stand-by**" (local/global spare drive) to clone the source drive (the target member drive).

Step 3. The cloning process will begin with a notification message:

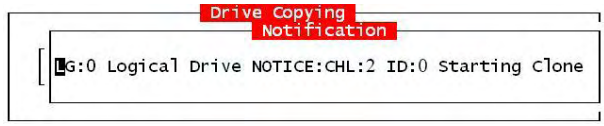

Step 4. Press **[ESC]** to view the current progress:

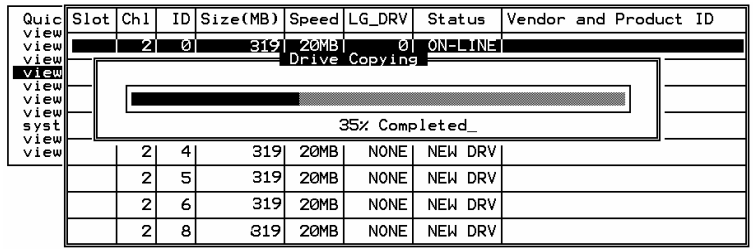

Step 5. You may quit viewing the status bar by pressing **[ESC]** to return to the previous menu. Select the drive indicated as "**CLONING**" by pressing **[ENTER]**. Select "**Clone Failing Drive**" again to view the progress. You may identify the source drive and choose to "**View Clone Progress**" or "**Abort Clone**" if you happen to have selected the wrong drive.

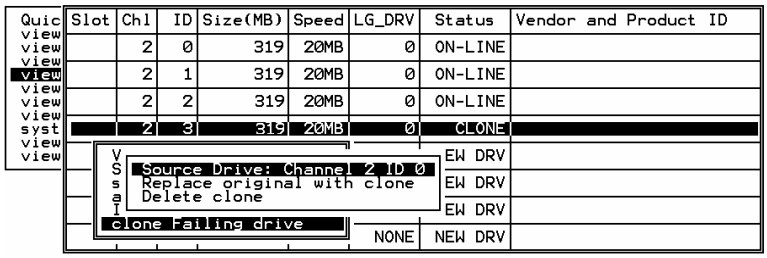

Step 6. The cloning progress will be completed by a notification message as displayed below:

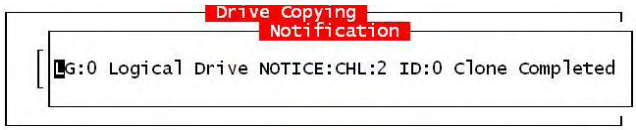

Step 7. You may press **[ESC]** to clear the notification message to see the drives' status after the cloning process. The source drive (Channel 1 ID 5) remains as a member of logical drive "0," and the "**stand-by**" drive (Channel 1 ID 2, the dedicated/global spare drive) has become a "**CLONE**" drive.

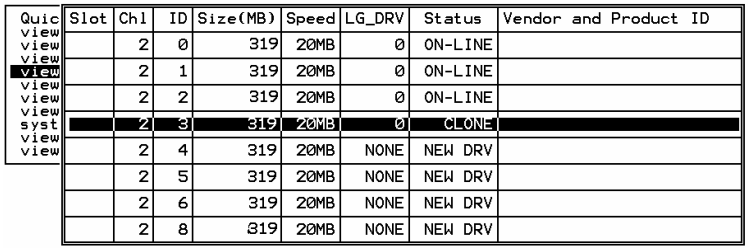

12.3 S.M.A.R.T. (Self-Monitoring, Analysis and Reporting Technology)

This section provides a brief introduction to S.M.A.R.T. as one way to predict drive failure and Infortrend's implementations with S.M.A.R.T. for preventing data loss caused by drive failure.

A. Introduction

Self-Monitoring, Analysis and Reporting Technology (S.M.A.R.T.) is an emerging technology that provides near-term failure prediction for disk drives. When S.M.A.R.T. is enabled, the drive monitors predetermined disk drives attributes that are susceptible to degradation over time.

If a failure is likely to occur, S.M.A.R.T. makes a status report available so that the host can prompt the user to backup data from the failing drive. However, not all failures can be predicted. S.M.A.R.T. predictions are limited to the attributes the drive can monitor which are selected by the device manufacturer based on the attribute's ability to contribute to predict degrading or fault conditions.

Although attributes are drive specific, a variety of typical characteristics can be identified:

- Head flying height
- Data throughput performance
- Spin-up time
- Re-allocated sector count
- Seek error rate
- Seek time performance
- Spin try recount
- Drive calibration retry count

Drives with reliability prediction capability only indicate whether the drive is "good" or "failing." In a SCSI environment, the failure decision occurs on the disk drive and the host notifies the user for action. The SCSI specification provides a sense bit to be flagged if the disk drive determines that a reliability issue exists. The system then alerts the user/system administrator.

B. Infortrend's Implementations with S.M.A.R.T.

Infortrend uses the ANSI-SCSI Informational Exception Control (IEC) document X3T10/94-190 standard.

There are four selections related to the S.M.A.R.T. functions in firmware:

Disabled

Disables S.M.A.R.T.-related functions

Detect Only:

When the S.M.A.R.T. function is enabled, the controller will send a command to enable all drives' S.M.A.R.T. function, if a drive predicts a problem, the controller will report the problem in an event log.

Detect and Perpetual Clone

When the S.M.A.R.T. function is enabled, the controller will send a command to enable all drives' S.M.A.R.T. function. If a drive predicts a problem, the controller will report the problem in an event log. The controller will clone the drive if a Dedicated/Global spare is available. The drive with predicted errors will not be taken off-line, and the clone drive will still act as a standby drive.

If the drive with predicted errors fails, the clone drive will take over immediately. If the problematic drive is still working and another drive in the same logical drive fails, the clone drive will resume the role of a standby spare and start to rebuild the failed drive immediately. This prevents a fatal drive error if yet another drive should fail.

Detect and Clone + Replace

The controller will enable all drives' S.M.A.R.T. function. If a drive predicts a problem, the controller will report the problem in the form of an event log. The controller will then clone the problematic drive to a standby spare and take the problematic drive offline as soon as the cloning process is completed.

Fail Drive

Before using this function, you should be ready with a hot spare or a replacement drive so that a logical drive having a member

disbanded can be quickly rebuilt. A disk drive can become unstable or dragging the array performance before being considered as a failed drive by your RAID system. If there are signs showing a member drive is seriously degraded, (such as recurring reports of slow responses), you can use this option to disband a faulty drive from a logical drive.

WARNING!

The Fail Drive option can impose a danger in the situation when other members of a logical drive carry immanent defects. In the extreme cases, similar defects may be found in disk drives of the same lot by the same manufacturer. If you fail a member in a RAID5 array and another member encounters media errors during the rebuild process, you will lose data.

NOTE:

- If you are using drives of different brands in your RAID system, as long as they are ANSI-SCSI Informational Exception Control (IEC) document X3T10/94-190-compatible, there should not be any problems working with the controller/subsystem.

12.4 Configuration Procedure

12.4.1 Enabling the S.M.A.R.T. Feature

Follow the procedure below to enable S.M.A.R.T. on all drives.

Step 1. First, enable the "**Periodic Drive Check Time**" function. In "**View and Edit Configuration Parameters" -> "Drive-side Parameters" -> "Periodic Drive Check Time**," choose a time interval.

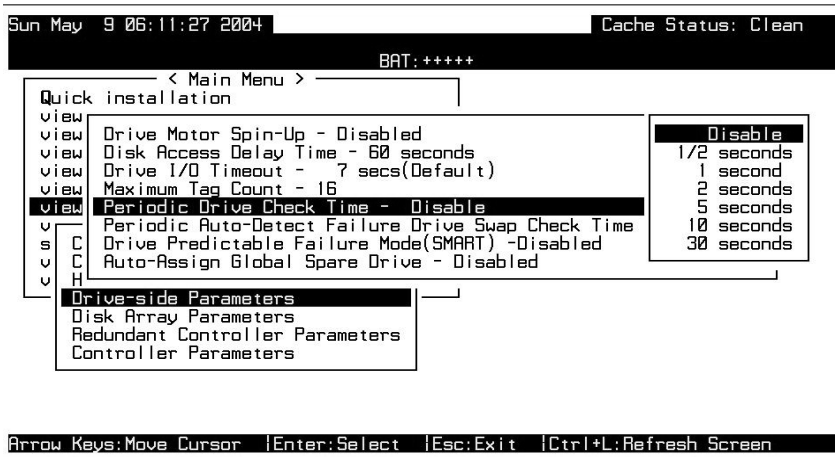

Step 2. In "**View and Edit Configuration Parameters" -> "Drive-side Parameters" -> "Drive Predictable Failure Mode <SMART>**," choose one from "**Detect Only**," "**Detect, Perpetual Clone**" and "**Detect, Clone+Replace**."

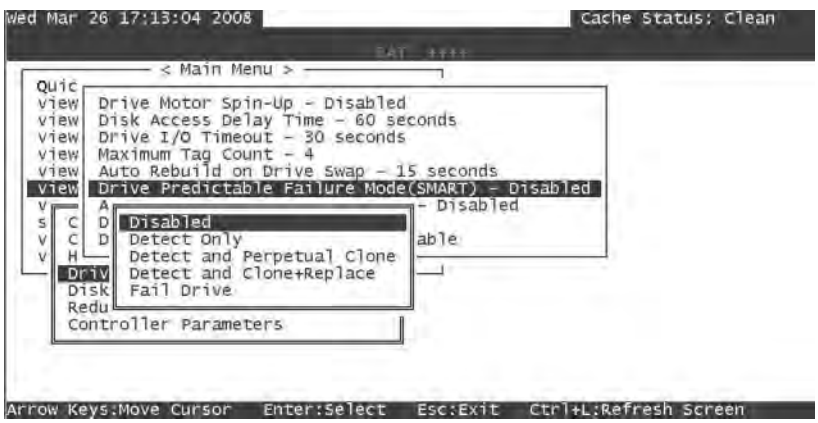

12.4.2 Using S.M.A.R.T. Functions

Step 1. Enable "**S.M.A.R.T.**" on the RAID controller. Make sure your drives do support S.M.A.R.T. so that your system will work properly.

Step 2. **Select a S.M.A.R.T.-related option**

Step 2-1. **The "Detect Only" Setting**

Step 2-1-1. In "**View and Edit Configuration Parameters" -> "Drive-side Parameters" -> "Drive Predictable**

Failure Mode <SMART>," choose "**Detect Only**."

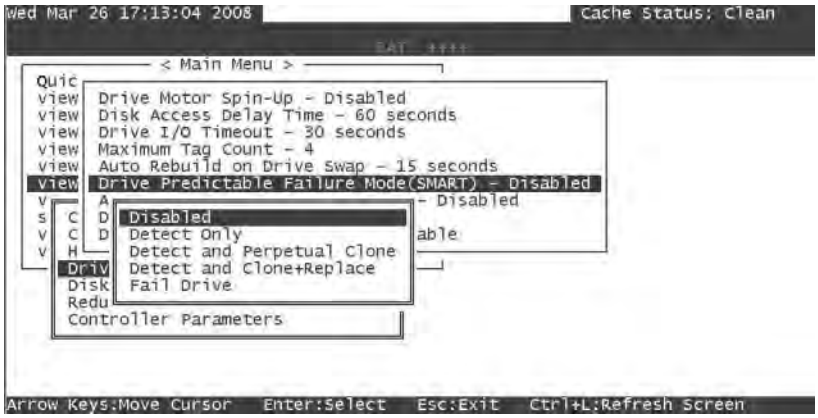

Whenever a drive detects symptoms of predictable drive failure, the controller will issue an error message.

Step 2-2. The "Detect, Perpetual Clone" Setting

Before selecting this option, you should make sure you have at least one spare drive for the logical drive (either Local Spare or Global Spare Drive).

Step 2-2-1. In "**View and Edit Configuration Parameters" -> "Drive-side Parameters" -> "Drive Predictable Failure Mode <SMART>**," choose "**Detect, Perpetual Clone**."

Step 2-2-2. When a drive (logical drive member) detects predictable drive errors, the controller will "clone" the drive with a spare drive. You may enter the "**View and Edit Drives**" menu and click on the spare drive (either Local or Global one). Choose from the menu items if you want to know the status of the source drive, the cloning progress, or to abort cloning.

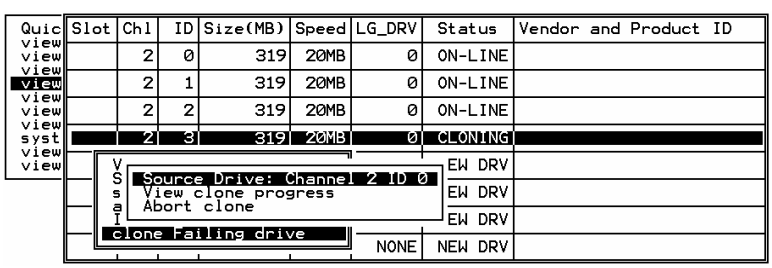

NOTE:

- As a precaution against the untimely failure of yet another drive, when configured as "**perpetual clone**," the spare drive will only stay mirrored to the source drive (the drive with signs of failure), but not replace it until the source drive actually fails.

- While the spare drive is mirroring the source drive, any occurrence of drive failure (when there are no other spare drives) will force the spare drive to give up the mirrored data and resume its original role – it will become a spare drive again and start rebuilding the failed drive.

Step 2-3. The "Detect, Clone + Replace" Function

Before enabling this option, make sure you have at least one spare drive to the logical drive (either **Local Spare Drive** or **Global Spare Drive**).

Step 2-3-1. In "**View and Edit Configuration Parameters" -> "Drive-side Parameters" -> "Drive Predictable Failure Mode <SMART>**," select "**Detect, Clone+Replace**."

When a drive (a logical drive member) detects the predictable drive failure, the controller will "**clone**" the drive with a spare drive. After the "clone" process is completed, it will replace the source drive immediately. The source drive will be identified as a "used drive."

Step 2-3-2. If you want to see the progress of cloning, press **[ESC]** to clear the notification message and see the status bar.

The source drive's status will be defined as a "used drive" and will be immediately replaced and pulled offline. This drive should be replaced with a new one as soon as possible.

This page is intentionally left blank.

Chapter

13

Implementations for AV Applications

This chapter introduces new firmware functions that optimize array performance for AV applications.

IMPORTANT!

Due to the wide variety of I/O demands by different AV applications, detailed parameters such as read-ahead or cache threshold parameters can be otherwise implemented only by communicating with our technical support. This chapter only presents two generic configuration options. More options will be available for specific applications as customized features.

NOTE:

All exemplary screens are captured from a hyper terminal management console.

AV Optimization Mode

The **AV optimization** option is applied for the emerging Audio/Video streaming applications such as the single-stream NLE (Non-Linear Editing), and the multi-stream VOD/MOD environments.

The AV Optimization Mode setting provides two configurable options: **Fewer Streams** and **Multi-Streaming**.

Fewer Streams: Read-ahead Performance

Applications such as an NLE (Non-Linear Editing) station may issue an I/O request for audio/video files of the sizes ranging from 1GB to 10GB or even larger.

Shown below is a RAID3 array configured in a 256KB stripe size. With only one 512KB outstanding I/O targeting at a large sequential file, the first I/O falls on two member drives while triggering a sequence of read-aheads at the same time. Read-aheads then occur across all member drives to make use of the combined disk performance.

The first I/O hit will be quickly returned and the read-aheads that ensue will be cached in memory. I/Os are then delivered through the read-aheads that are already stored in the fast data cache. As the result, applications featuring very few streams will be efficiently serviced through read-aheads in cache with minimized latency.

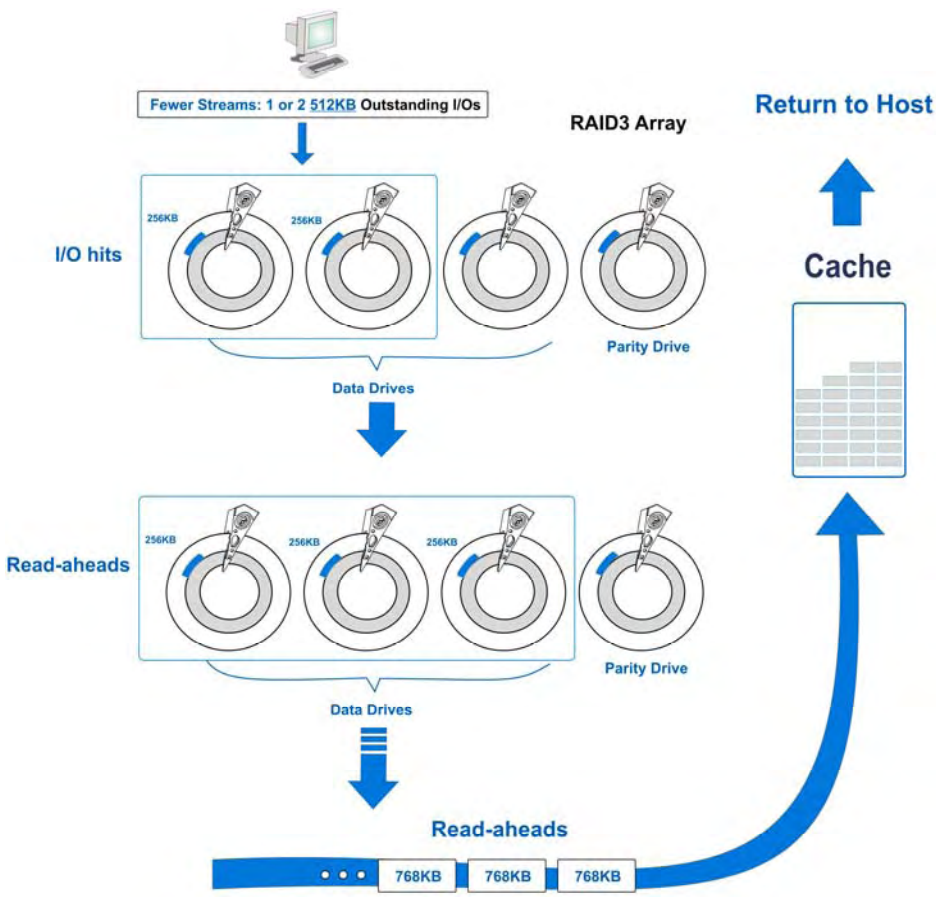

Figure 13-1: Read-ahead Operation for Applications Featuring Fewer Streams

With the **Fewer Streams** setting, the related **Maximum Drive Response Time** is automatically set to 160ms to prevent interruptions by media errors.

Multi-Streaming: Simultaneous Access Performance

The **Multi-Streaming** option is designed for applications featuring shorter-length and concurrent requests coming in the swarm of outstanding I/Os, e.g., low-bit-rate clips in VOD or MOD Media Broadcasting.

Shown below is a RAID3 array configured in a 512KB stripe size. With multiple, say, 16 outstanding I/Os targeting at different data files, I/Os fall simultaneously on different member drives. As the result, each hard drive's actuator arms can quickly move to the next location to fulfill another I/O request.

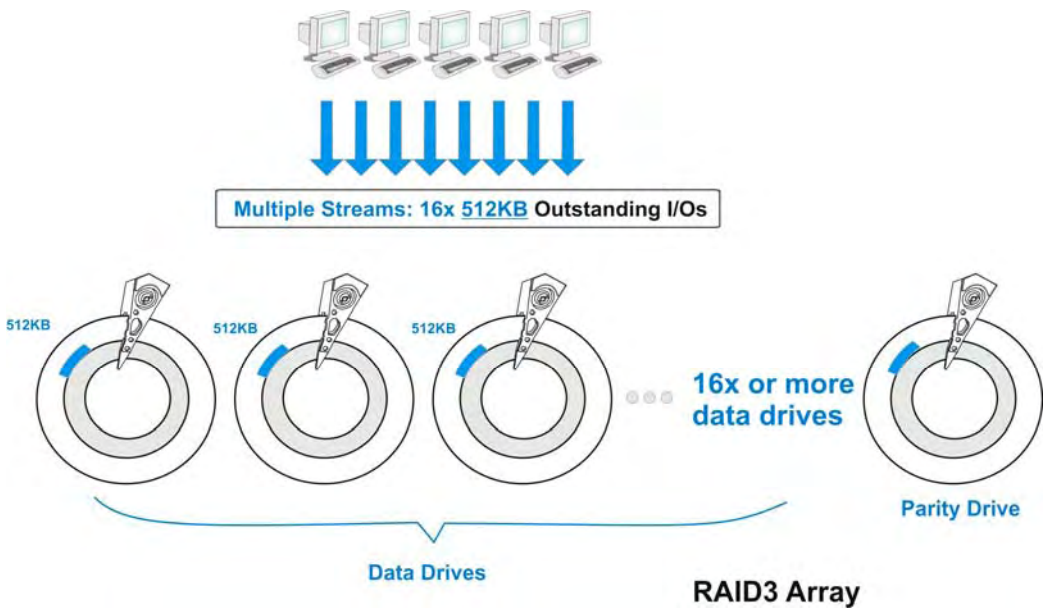

Figure 13-2: Multi-Streaming Applications

The **Multi-Streaming** option automatically configures the **Maximum Drive Response Time** to 960ms because read latency cause less-serious problems with the smaller, randomly-generated file requests in VOD/MOD than the large, sequential files in NLE applications.

The **Multi-Streaming** applications require the following:

1. A logical drive consisting **sufficient** number of disk drives so that I/Os can fall simultaneously on different members. Even though the real-world applications do not always make a

perfect fit, configuring an array using an equal or slightly larger stripe size will ensure each individual outstanding I/O can fall within the range of a data drive's strip size (or chunk size).

2. Properly tune the application I/O transfer size.

3. Appropriate stripe size of your RAID arrays.

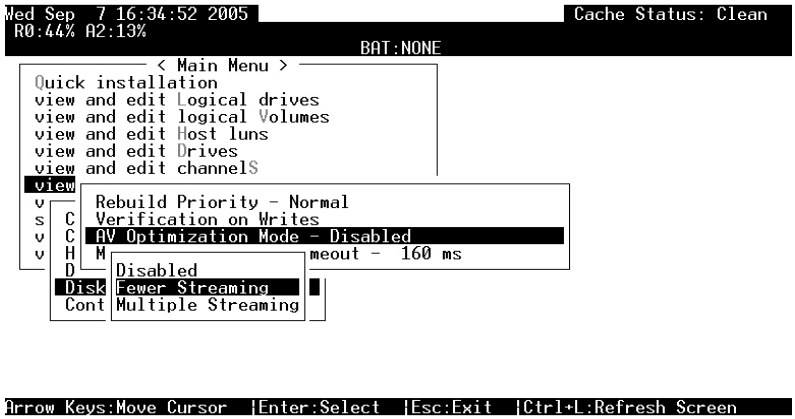

The AV Optimization mode can be found under "**View and Edit Configuration Parameters**" -> "**Disk Array Parameters**" -> "**AV Optimization Mode**."

NOTE:

The Maximum Drive Response Timeout bundled within the AV Optimization mode will over-rule any value you previously set in the similar menu found under Main Menu -> "**View and Edit Configuration Parameters**"-> "**Disk Array Parameters**."

Chapter

14

Redundant Controller

Sample topologies using redundant controllers can be found in the following discussions or in the ***Installation and Hardware Reference Manual*** that came with your RAID subsystems. The proceeding discussions will focus on the working theories and the configuration procedures for readying a redundant controller system.

This chapter covers the following topics:

Configuration Concerns

- **14.1 Configuration Concerns**

Requirements

- **14.2 Requirements**

Configuration Processes

- **14.3 Array Configuration Processes in Dual-controller Mode**

Troubleshooting

- **14.4 Troubleshooting Controller Failure**

Related Functions

- **14.5 Configurable Parameters Related to Redundant Controllers**

Working Theory

- **14.6 Operation Theory**

Configuration Samples

- **14.7 Configuration Samples**

14.1 Configuration Concerns

Listed below are the configuration concerns and phenomena you will encounter when configuring a redundant controller subsystem:

1. By system default, **Controller A** is always the **primary** RAID controller. **Controller B** in the lower slot serves as the **secondary**. If Controller A fails and is replaced afterward, firmware returns the Primary role to the replacement controller after a system reset.

2. The traditional mapping method co-exists with the new, cross-controller access available with the firmware 3.48 release.

Firmware before rev. 3.48	A RAID configuration can only be managed by one controller, i.e., mapping to IDs owned by one controller.
Firmware rev. 3.48	A RAID configuration can be alternatively managed by both RAID controllers, i.e., mapping to IDs owned by two controllers.

CAUTION!

Firmware revisions 3.48 and later support the **cross-controller** ID mapping. The cross-controller mapping allows you to associate a logical drive with BOTH controller A and controller B IDs. However, mapping to both controllers' IDs is only beneficial when it is difficult making the fault-tolerant links between RAID controllers and host HBAs, e.g., using SAS-to-SAS RAID systems. Currently, there is no external SAS switch available on the market. For Fibre-host systems, fault-tolerant links can easily be made with the help of external bypass such as Fibre Channel switches.

For details of fault-tolerant link connections, please refer to your system Hardware Manual.

One benefit of the cross-controller access is that when a host link fails, I/Os can travel through the counterpart controller, the RCC links, and then back to the RAID controller originally managing the array. The I/O load will still be managed by two controllers in the event of host link failure.

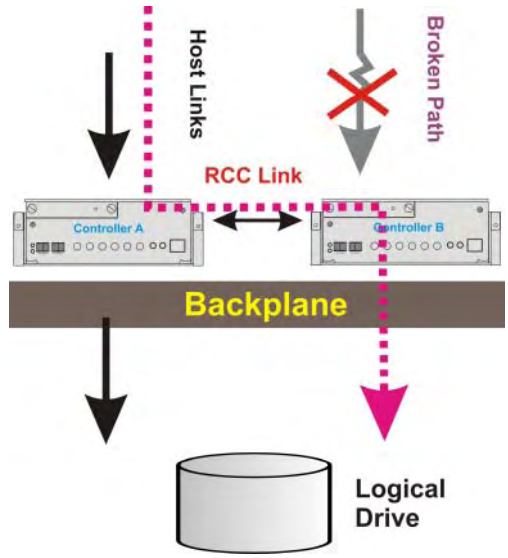

Figure 14-1 Routing Access when a Host Link Fails

3. If your subsystem comes with an LCD, the upper right corner of LCD will display a "**P**" or "**S**," meaning "**Primary**" or "**Secondary**" respectively. You may press the arrow keys together for two seconds to switch between the display of the **Primary** or **Secondary** controller status.

4. The controller partners synchronize each other's configurations at frequent intervals through the communications channel(s). And the synchronization act consumes part of the system resource.

14.2 Requirements

Requirements for Redundant-Controller Configurations

Communications Channels

- Controller Communications (Cache Synchronization) Paths:

Controller	RCC
Subsystem	Pre-configured RCC routes over the system backplane; may be SCSI, Fibre, or SATA data paths. These data paths cannot be re-assigned.
1U controller head	**ASIC266**: "Dedicated RCC" or "Drive+RCC." **ASIC400**: pre-configured RCC routes; no need to assign.

- If controllers are running with write-back caching, a battery module on each controller is highly recommended.

Out-of-Band Configuration Access

- **RS-232C serial port cable** (for terminal interface operation) connection. Normally a Y-cable will be included with dual-controller subsystems. The Y-cable ensures a valid link in the event of single controller failure.
- **Ethernet connection:** If management through Ethernet is preferred, connect the Ethernet interface from both controllers to your local network. In the event of controller failure, the IP address assigned to the Primary Controller will be inherited by the surviving controller. In this way, the Ethernet port connection (management session) will be interrupted. An operator may have to re-enter the IP address to re-connect the controller/subsystem to a management console.

Limitations

- Both controllers must be exactly the same. Namely, they must operate with the same firmware version, the same size of cache memory, the same number/configuration of host and drive channels, etc. If battery backup is preferred, both should be equipped with a battery module.

- If a RAID controller fails and needs to be replaced, it is often the case that the replacement controller may carry a newer revision of firmware. It is advised you provide information such as firmware revision number, boot record version, etc. to your system vendor before acquiring for a replacement controller.

- For a subsystem featuring Fibre host channels and if the onboard hub is not enabled, connection through Fibre switches will be necessary for configuring fault-tolerant paths between host and RAID storage.

  In the event of data path failure, an intelligent FC switch should be able to direct data flow through an alternate path. In this case, multipathing software should be necessary to manage the data flow through the fault-tolerant paths that are strung between host and RAID storage.

- Your RAID subsystem may not come with sufficient numbers of Controller A and Controller B IDs. You will then need to manually create Controller A or Controller B IDs.

Configurable Parameters

Active-to-Active Configuration

Since firmware revision 3.48, users can freely map a logical configuration to both the Controller A and Controller B IDs [putting forth different LUN views of a logical storage unit to different initiators (HBAs on servers)]. The I/O load to a logical drive can be dynamically shared by partner controllers.

The traditional mapping method requires at least two logical units which are separately managed by a RAID controller. Each logical unit is associated either with Controller A or Controller B IDs.

The dual-active configuration engages all system resources to performance. Users may also assign all logical configurations to one controller and let the other act as a standby (active-standby).

Cache Synchronization (Mirrored Cache)

The Write-back caching significantly enhances controller performance. However, if one controller fails in the redundant-controller configuration, data cached in its memory will be lost and data inconsistency will occur when the surviving controller takes over and attempts to complete the writes.

Cache synchronization distributes cached writes to both controllers and each controller stores an exact replica of the cache content on its counterpart. In the event of controller failure, the unfinished writes will be completed by the surviving controller.

14.3 Array Configuration Processes in Dual-controller Mode

General Firmware Configuration Procedures

Below are the basic procedures for readying a redundant-controller subsystem:

Step 1. **Controller Unique Identifier**

The EonStor subsystems usually come with a default identifier. If the default is lost for some reasons, provide a unique identifier for each RAID subsystem.

"**View & Edit Peripheral Devices**" → "**Controller Parameters**" → "**Controller Unique Identifier**."

Step 2. **Create Controller A and Controller B IDs**

Step 2-1. "**View and Edit Channels**"→ Choose a host channel.

Step 2-2. "**View and Edit ID**"→ Select an existing ID.

Step 2-3. Under "**Add/Delete Channel ID**" → "**Controller A/Controller B**" → Select an ID from the pull-down list.

Step 2-4. Reset the controller for the configuration to take effect.

Step 3. **Create Logical Configurations of Drives**

Step 3-1. Under "**View and Edit Logical Drives**"→ Select a Logical Drive entry.

Step 3-2. Select a RAID level.

Step 3-3. Select member drives

Step 3-4. Configure other parameters, e.g., stripe size.

Step 3-5. Assign logical drive either to the Controller A or Controller B.

Step 3-6. Create Logical Drive.

Step 4. **Map Each Logical Configuration of Drives to Controller A and/or Controller B IDs on host channel(s)**

Step 4-1. Under "**View and Edit Host LUN**"→ Choose a "**Channel-ID-Controller**" Combination.

Step 4-2. Select a "**Logical Drive**" or "**Logical Volume**" and then the "**Logical Partition**" within the logical unit → "**Map to Host ID/LUN**" (Create Host LUN Entry).

Step 4-3. Repeat the process to avail a logical drive through multiple host IDs so that host can access the array through different host ports.

Detailed Procedures:

14.3.1 Controller Unique ID

This value will be used to generate a controller-unique WWN node name, port names, Ethernet port MAC address, and to identify the controller during the failover process.

EonStor systems come with a default ID. It is recommended to use it. If the unique ID is accidentally lost, you can create a new ID using the following procedure:

Enter "**View and Edit Config Parms**"-> "**Controller Parameters**". Use the up or down arrow keys to find "**Ctlr Unique ID- xxxxx**".

```
View and Edit
Peripheral Dev..
```

```
Ctlr Unique (Hex)-
ID-
```

Enter a hex number from 0 to FFFFF and press **[ENTER]**. The value you enter should be different for each RAID subsystem.

The unique identifier setting can be accessed from "**View and Edit Configuration Parameters**" -> "**Controller Parameters**" -> "**Controller Unique ID**."

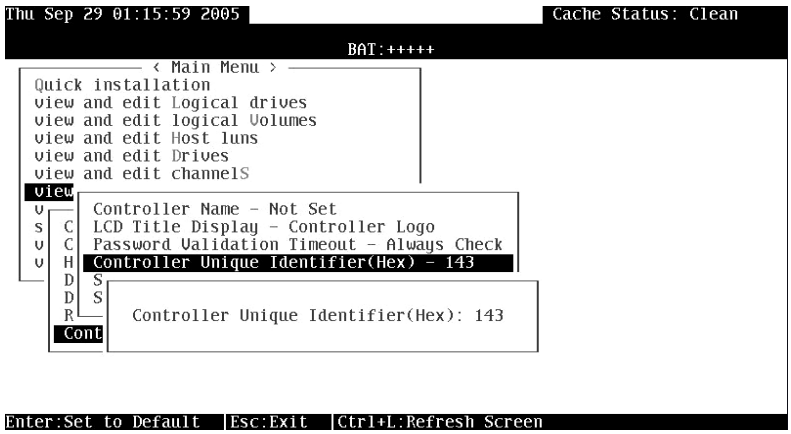

14.3.2. Creating Controller A and Controller B IDs

The dual-controller operation requires that you manually create more Controller A and Controller B IDs.

LCD Keypad

Step 1. In "**View and Edit Channels**", press **[ENT]** to select a host channel.

```
CHL=0 ID=0
Slot A Ctlr  ..
```

Step 2. Use the up or down arrow keys to select "**Set Channel ID**". A pre-configured ID will appear, press **[ENT]** to proceed.

Step 3. Use the up or down arrow keys to select "**Add Channel ID**" and then press **[ENT]** for two seconds on the "**Slot A**" or "**Slot B?**" option to proceed.

```
Add Channel
ID    ..
```

```
Slot A
Controller    ?
```

Step 4. When prompted by this message, use the arrow keys to select an ID. Press **[ENT]** to confirm.

```
Add CHL=0 ID=113
Slot A Ctlr  ?
```

Step 5. A message will prompt to remind you to reset the controller. Press **[ENT]** to reset the controller or press **[ESC]** to return to the previous menu. The ID change will only take effect after a system reset.

```
Change Setting
Do Reset Ctlr ?
```

Terminal

Step 1. Enter "**View and Edit Channels**," then press **[ENTER]** and select the host channel on which you wish to create Controller A or Controller B IDs.

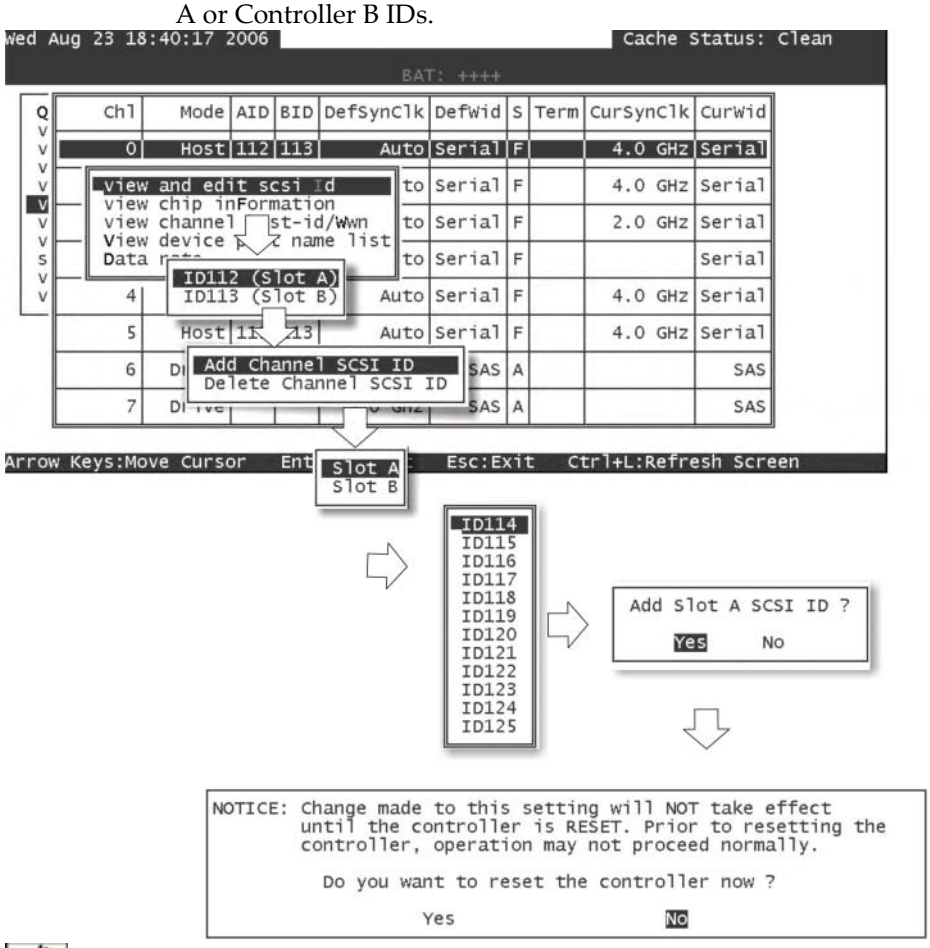

Step 2. Enter the "**View and Edit IDs**" sub-menu by pressing **[ENTER].**

Step 3. Press **[ENTER]** on an existing ID.

Step 4. Select "**Add Channel SCSI ID**." Press **[ENTER]** to confirm.

Step 5. Select either "**Slot A**" or "**Slot B**" controller to create IDs that will be managed by a designated RAID controller.

Step 6. A pull-down list will display all available IDs. Use your arrow keys to select an ID and press **[ENTER]** to confirm.

The configuration change will only takes effect after a system reboot.

14.3.3. Logical Drive Assignments (Dual-Controller Configuration)

A logical drive or a logical volume can be assigned to either Controller A or Controller B. By default, a logical drive is automatically assigned to Controller A, the controller installed in the upper slot (also the Primary controller by factory default). To divide the workload, you may manually assign a logical drive to Controller B.

NOTE:

By default, logical drives will always be assigned to the Slot A controller. They can be manually assigned to the Slot B controller if the host computer is also connected to the Secondary controller.

LCD Keypad

Step 1. Press **[ENT]** key for two seconds to enter the firmware utility's Main Menu.

Step 2. Use the arrow keys to navigate through the configuration menus. Choose "**View and Edit Logical Drives**", then press **[ENT]**.

```
View and Edit
Logical Drives
```

Step 3. Create a logical drive or choose an existing logical drive, then press **[ENT]** to see the logical drive menu. The creation procedure is detailed in previous chapters.

Step 4. Choose "**Logical Drive Assignment..**," then press **[ENT]**.

```
Logical Drive
Assignment..
```

Step 5. The message "**Redud Ctlr LG Assign Slot B?**" will appear. Press **[ENT]** for two seconds to confirm.

```
Redud Ctlr LG
Assign Slot B?
```

Step 6. Map the logical drive (or any logical unit) to a host ID or a LUN number under controller B ID. The host channel must have a "**Slot B**" ID. If not available, Slot B IDs can

be manually added to a host channel.

Step 1. Access "**View and Edit Logical Drives**" from the Main Menu. Create a logical drive by selecting members and then a selection box will appear on the screen.

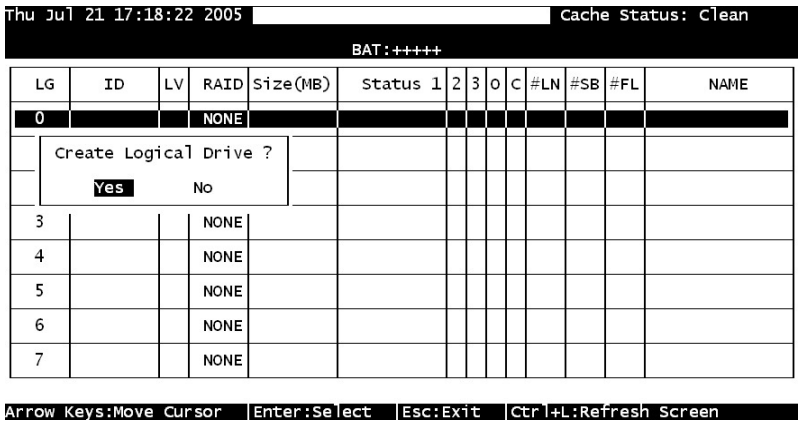

Step 2. For the first logical drive on the RAID subsystem, simply select the first logical drive entry, **LG 0**, and press **[ENTER]** to proceed. You may create as many as 32 logical drives or more using drives in a RAID subsystem or in an expansion enclosure.

Step 3. When prompted to "**Create Logical Drive?**," select ***Yes*** and press **[ENTER]** to proceed. Please refer to the previous chapters for options specific to individual logical drives.

Step 4. Access "**View and Edit Logical Drives**" from the Main Menu. Move the cursor bar to "**Logical Drive Assignments**" and press **[ENTER]** if you want to assign logical drive to another controller.

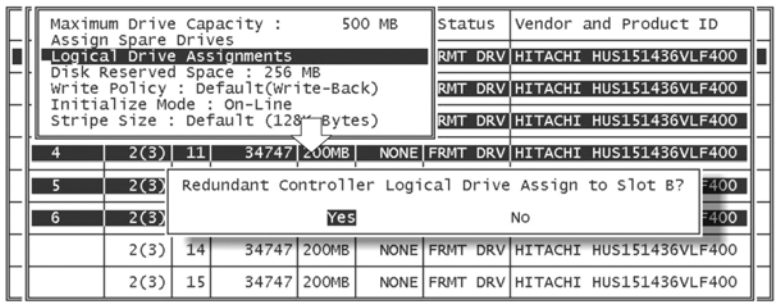

Logical drive assignment can also be changed after a logical drive is created.

Step 5. Select an existing logical drive from the "**View and Edit Logical Drive**" menu, then press **[ENTER]** to see the logical drive menu.

Step 6. Select "**Logical Drive Assignments**," then press **[ENTER]**. Choose **Yes** and press **[ENTER]** to confirm reassignment to an alternate controller.

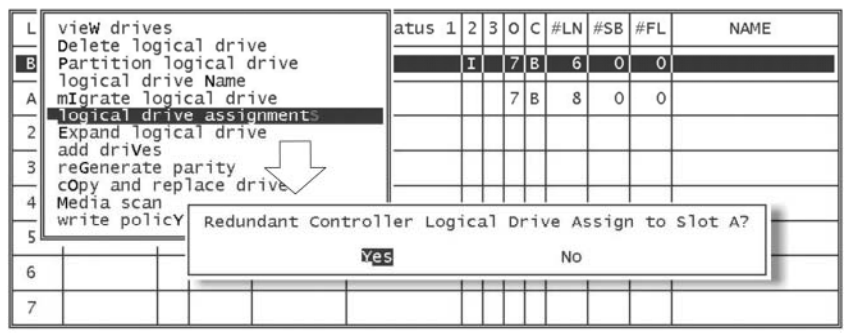

The reassignment is evident from the Logical Drive Status screen. "**B0**" indicates that the logical drive is Logical Drive #0 assigned to the **Slot B** controller.

NOTE:

- You cannot reassign a logical drive until it is disassociated with host ID/LUNs (remove the previous LUN mapping).

14.3.4. Mapping a Logical Drive/Logical Volume to Host LUNs

TIPS:

1. Before proceeding with the mapping process, draw an abstract diagram of your configurations to help clarify the relationships among physical and logical components.
2. Before the mapping process, check if you have properly configured logical drives, logical drive assignment, and host IDs. Changing host LUN mapping and re-configuring a RAID array may also require corresponding efforts on the management software running on host.

LCD Keypad

Step 1. Choose "**View and Edit Host Luns**" from Main Menu and press **[ENT]** to proceed.

```
View and Edit
Host Luns
```

Step 2. Use the arrow keys to navigate through the list of existing IDs and press **[ENT]** to select one of them.

```
Map Slot B
CH=0 ID= 000  ?
```

Step 3. Use the arrow keys to select mapping to a "**Logical Drive**" or "**Logical Volume**" to host LUNs. If the logical unit has been partitioned, map each partition to different IDs/LUNs.

```
Map to
Logical Drive ?
```

Step 4. Use the arrow keys to choose a LUN number and press **[ENT]** to confirm.

```
CH0 ID0 LUN0
No Mapped
```

Step 5. Press **[ENT]** again to confirm.

```
Map Host LUN   ?
```

Step 6. Use the arrow keys to select a logical drive/logical volume if there are many.

```
LG0 RAID5 DRV=3
  9999MB GD SB=0
```

Step 7. Press **[ENT]** and select a partition if the logical unit has been partitioned into individual capacity volumes.

```
LG=0 PART=0
     999MB   ?
```

```
CH0 ID9 LUN0 Map
to LG0 PRT0?
```

Step 8. Press **[ENT]** again to confirm.

Step 9. Press **[ENT]** to proceed.

Step 10. Press **[ENT]** to confirm.

Step 11. This message indicates that the logical unit has been successfully mapped to the ID/LUN combination. Use the arrow keys to continue mapping other logical units or press **[ENT]** to delete the mapped LUN.

```
Map Host LUN   ?
```

```
Map Slot B
CH=0 ID= 0   ?
```

```
CH0 ID9 LUN0
Mapto LG0 PRT0
```

Step 12. Repeat the process to map all logical units to host ID/LUNs.

Terminal

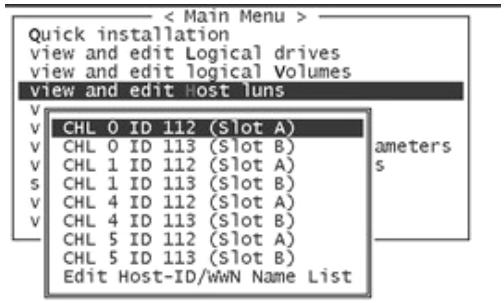

Step 1. Enter the "**View and Edit Host LUNs**" menu. A list of existing host IDs will be displayed.

Step 2. Select a host ID (either a **Controller A** or **Controller B** ID) by pressing [**ENTER**]).

Step 3. Select the type of logical unit (**Logical Drive** or **Logical Volume**) you wish to associate with the host ID.

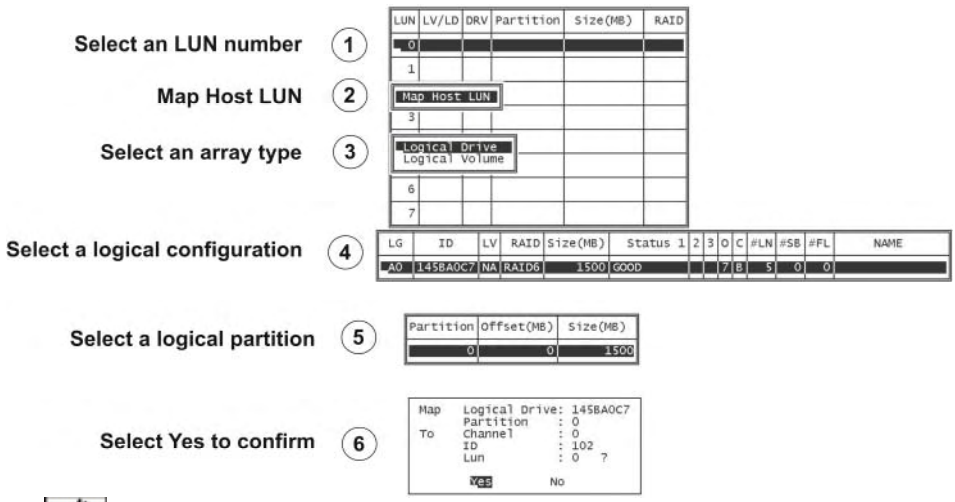

Step 4. Select an LUN number under the host ID.

Step 5. All logical units will be listed. Select one of them by pressing [**ENTER**] on it.

Step 6. When selected, all logical partitions under the logical unit will be listed. Select a partition.

Step 7. A confirm box will appear. Verify the details and press [**ENTER**] on **Yes** to complete the mapping process.

Step 8. Repeat this process until you finish mapping all logical units to the host IDs you prefer. Repeat the process to map a logical unit to two host ID/LUNs if you want it to appear on two data paths.

14.4 Troubleshooting Controller Failure

What will happen when one of the controllers fails?

If one of the controllers fails, the surviving controller will automatically take over within a few seconds.

NOTE:

- Although the surviving controller will keep the system running, you should contact your system vendor for a replacement controller as soon as possible. Your vendor should be able to provide an appropriate replacement.
- You should provide your vendor the serial number of the failed controller and the firmware version currently running on your system.
- Some operating systems (SCO, UnixWare, and OpenServer, for example) will not automatically retry with I/Os shortly delayed while the controller is taking over.

LCD Keypad

The red **ATTEN LED** on the LCD panel will light up, and the message "**Redundant Ctlr Failure Detected**" will appear on the LCD. Users will also be notified by audible alarm and messages sent over event notification methods such as Email, LAN broadcast, etc.

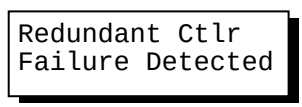

Terminal

1. When one controller fails, the other controller will take over in a few seconds.
2. There will be an alert message that reads "**Redundant Controller Failure Detected**."

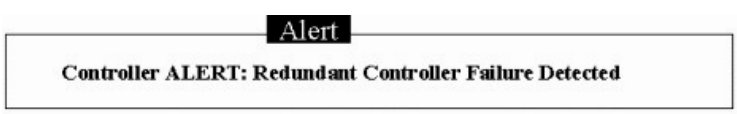

3. Users will be notified by audible alarm and the messages through event notification methods such as Email, LAN broadcast, etc.
4. After a controller takes over, it will act as both controllers. If the Primary Controller fails, the Secondary Controller manages the logical drives originally managed by the Primary Controller.

When and how is the failed controller replaced?

Remove the failed controller **AFTER** the "**working**" controller has taken over. For the ventilation concern in RAID enclosures, it is better to leave a failed controller in place before a replacement arrives.

NOTE:

- If you need to replace a failed controller, **DO IT WHEN THE SYSTEM IS POWERED ON AND IS MANAGED BY THE SURVIVING CONTROLLER!** Problems may occur if you replace a failed controller when system is powered down. If you power up both the surviving controller and the replacement together, they may contend for the role of the Primary (dominating) controller.

  If you encounter this problem you may follow the procedure below to correct the fault:

  LCD Keypad

View and Edit Peri. Device	Redundant Ctrl Primary Scanning	Enable redundant As Primary?

  NOTICE: Change Made to this

  Terminal

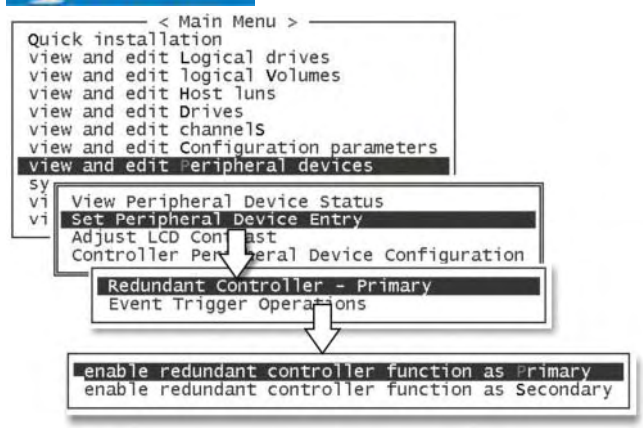

Step 1. Stop host I/Os.

Step 2. Power down the system and remove the surviving controller.

Step 3. Power on and enter Main Menu -> **View and Edit Peri. Device**

-> **Set Peri. Device Entry** -> "**Redundant Controller**" and change the controller role.

Step 4. You may then install both controllers into their original positions and power on the RAID enclosure.

- Check if the replacement is properly installed.

Redundant controller subsystems are designed to withstand a single controller failure. If the replacement does not initialize properly, try the following:

When the replacement is connected, the "Auto-Failback" process should start automatically. If the replacement controller does not initialize, you may execute the following steps to bring the new controller online.


Step 1. Press **[ENT]** for two seconds on the existing controller to enter the Main Menu.


Step 2. Use the arrow keys to select "**View and Edit Peripheral Dev..**," then press **[ENT]**.

```
View and Edit
Peripheral Dev
```


Step 3. Choose "**Set Peripheral Device Entry..**," then press **[ENT]**.

```
Set Peripheral
Devices Entry ..
```


Step 4. Select "**Redundant Ctlr Function__**," then press **[ENT]**.

```
Redundant Ctlr
Function__
```


Step 5. The message "**Redundant Ctlr Slot A/Slot B Degraded**" will appear on the LCD.

```
Redundant Ctlr
Slot A Degraded
```


Step 6. Press **[ENT]** and the message "**Deassert Reset on Failed Ctlr?**" will appear.

```
Deassert Reset
on Failed Ctlr?
```


Step 7. Press **[ENT]** for two seconds and the controller will start to scan for the new controller and bring it online.

```
Redundant Ctlr
Primary Scanning
```

Step 8. The new controller will then start to initialize.

```
Initializing...
  Please Wait...
```

Step 9. Once initialized, the replacement controller should assume the role of the Secondary Controller, and if the replacement is installed into the upper slot, it will restore its Primary role after a system reboot.

```
                v3.**
■■■■■■
```

Terminal

When the new controller is connected, the existing controller will automatically start initializing the replacement controller. If the replacement controller failed to initialize, try the following:

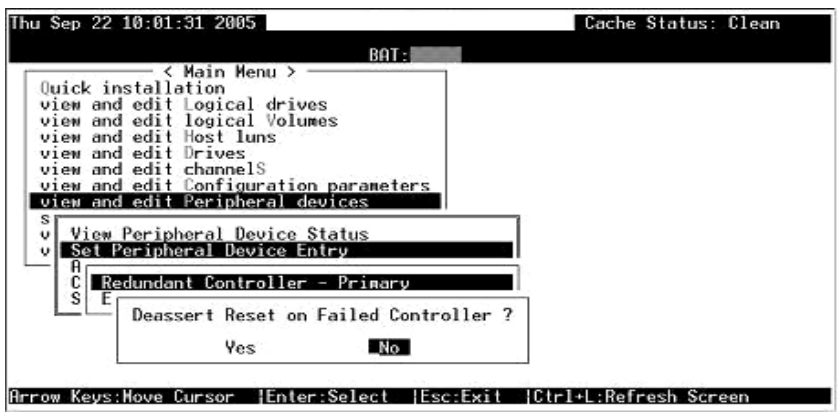

If the replacement has been initialized successfully, you may proceed to examine the system status. From the Main Menu, select "**View and Edit Peripheral Devices**" and then "**View Peripheral Device Status**" to see that the new controller is being scanned.

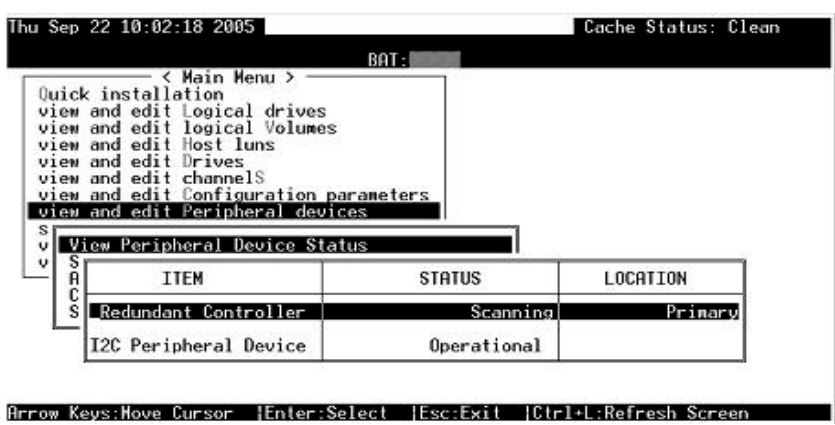

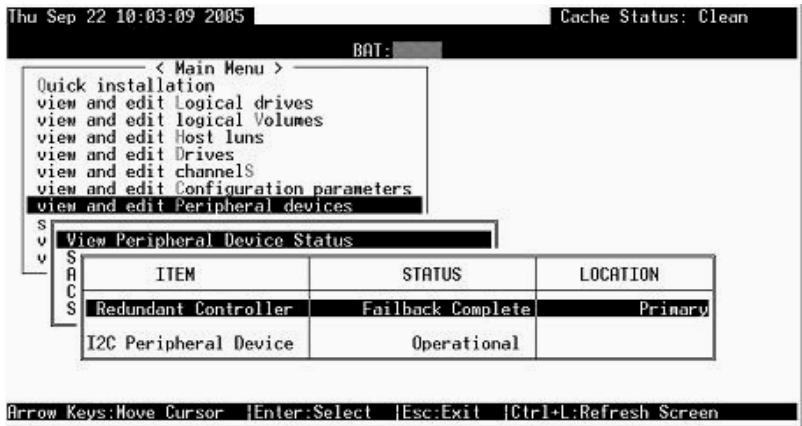

When the scanning is completed, the status will change to "**Failback Complete**."

14.5 Configurable Parameters Related to Redundant Controllers

14.5.1 RCC (Redundant Controller Communications Channel) Status

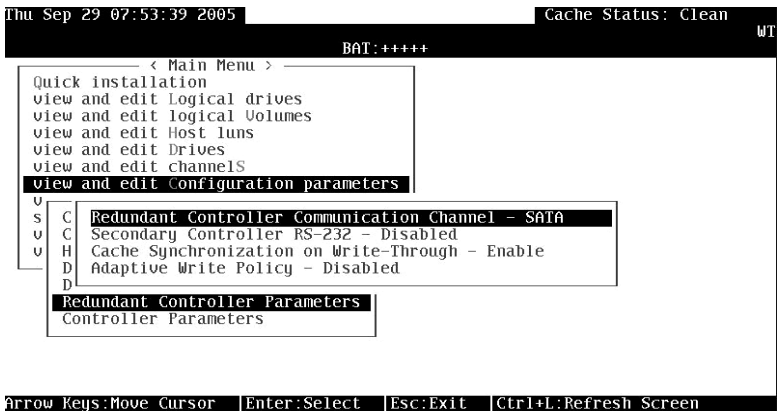

This item is for display only, showing the current communications routes.

14.5.2 Adaptive Write Policy

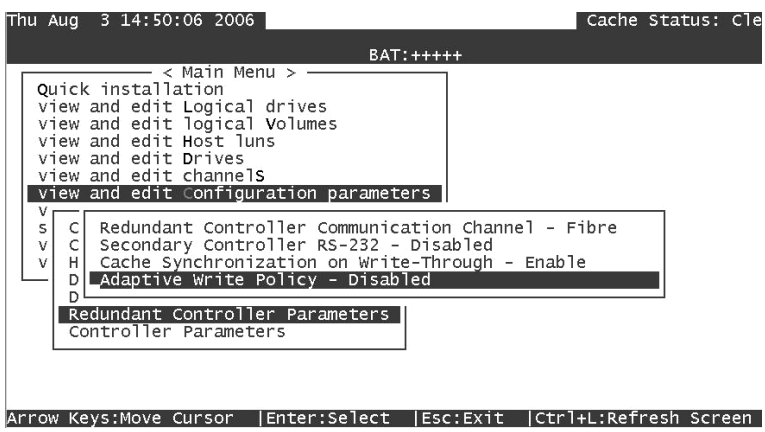

Firmware is embedded with intelligent algorithms to detect and to adapt the array's caching mode to the characteristics of I/O requests. The adaptive capability is described as follows:

1. When enabled, the **Adaptive Write Policy** optimizes array performance for sequential writes.

2. The adaptive policy temporarily disables an array's write-caching algorithm when handling sequential writes. Write-caching can be unnecessary with sequential writes for that write requests can be more efficiently fulfilled by distributing writes directly onto disk drives following the receiving order.

3. The adaptive policy changes the preset write policy of an array when handling I/Os with heterogeneous characteristics. If firmware determines it is receiving write requests that come in a sequential order, the write-caching algorithm is disabled on the target logical drives.

   If the subsequent I/Os are fragmented and are received randomly, firmware automatically restores the original write-cache policy of the target logical drives.

Adaptation for the Redundant Controller Operation

4. If arrays managed by a redundant-controller configuration are configured to operate with write-back caching, cached data will be constantly synchronized between the partner controllers. Upon receiving sequential writes, firmware disables write-caching on target arrays and also the synchronized cache operation because the synchronization also consumes some of the controllers' processing power.

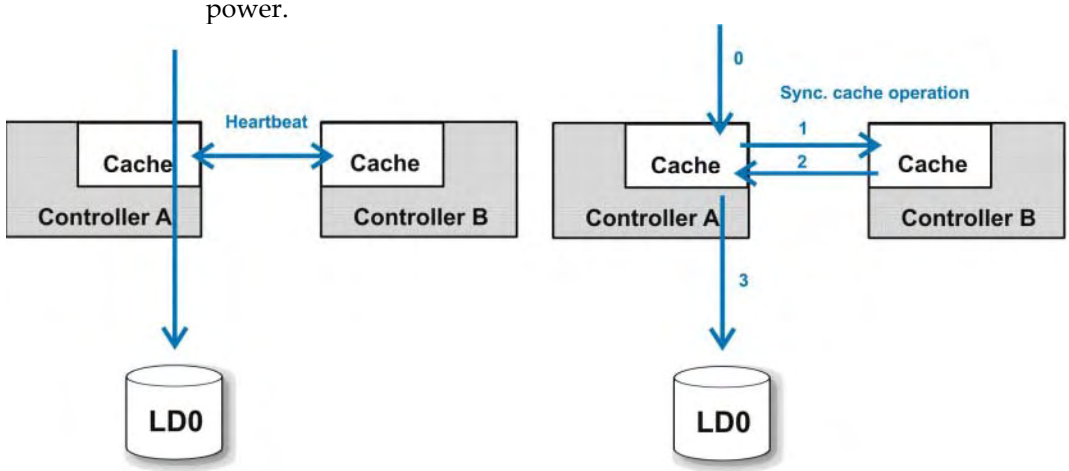

Figure 14-2 Write-through and Write-back caching with Synchronized Cache Efforts

IMPORTANT!

- Every time you change the Caching Parameters, you must reset the controller for the changes to take effect.

- The Adaptive Write Policy is applicable to subsystems working in the normal condition. If, for example, a drive fails in an array, firmware automatically restores the array's original write policy.

14.5.3 Cache Synchronization on Write-Through

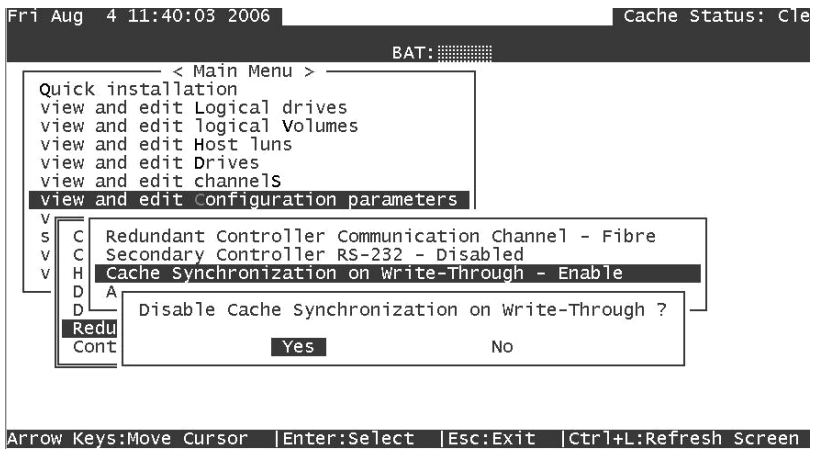

If your redundant controller system is not operating with Write-back caching, you can disable the synchronized cache communications between RAID controllers. Your system can be spared of the efforts to mirror and transfer data between partner controllers. This increases array performance for subsystems that operate without write caching.

Note that the configuration changes made to the RAID subsystem firmware will still be synchronized between the partner controllers.

14.6 Operation Theory

14.6.1 The Inter-Controller Relationship

For subsystems running firmware release **3.48** or later, the **Primary**/**Secondary** controller role is determined by a controller's position in a RAID enclosure. The new principle helps ensure the fixed location of a dominating, "Primary," controller. Other aspects of array management, ID/LUN mapping and array operation remain basically unchanged.

The new principle defines the RAID controller installed in Slot A, the upper controller slot, as the Primary controller. The factory configuration ensures that the **Slot A** controller always behaves as a **Primary controller**. In the following condition, a slot A controller temporarily serves as a Secondary controller:

1. If the Slot A controller fails, the original Slot B (Secondary) controller takes over and becomes the Primary controller.
2. When the slot A controller is replaced by a new controller, the new controller temporarily serves as the Secondary controller.
3. Once the subsystem is reset or powered down and powered on again, firmware returns the Primary role to the replacement controller in slot A.

14.6.2 Grouping Hard Drives and LUN Mapping

Configuration Rules

Listed below are the basics about configuring RAID arrays in a redundant-controller system:

1. All configuration utilities are managed by the Primary RAID (normally the controller A) controller.

2. Controller B status can also be displayed on a terminal or LCD screen. Management screen of a specific RAID controller is indicated by a flashing digit, <A> or <B> respectively on an LCD screen. Messages generated by different controllers will also be noted as shown below.

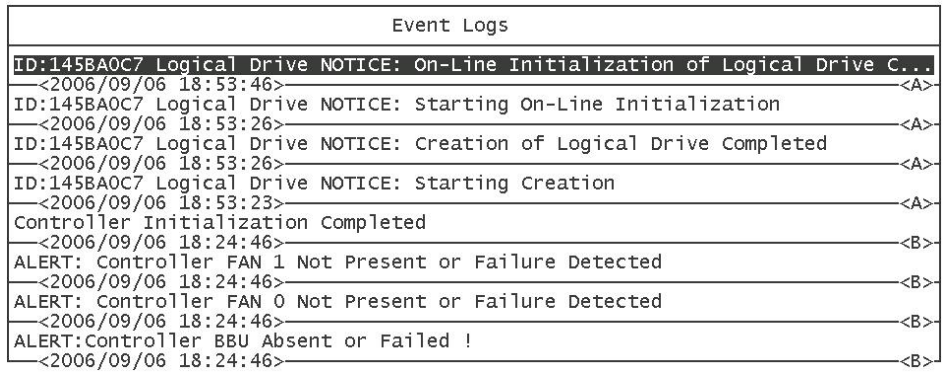

3. In dual-controller mode, two controllers behave as one, and there is no need to repeat the configuration on another controller. The array configuration profile is automatically synchronized between partner controllers.

4. Disk drive and array configuration processes are the same for subsystems using single or dual-active controllers.

5. Using logical drives as the basic configuration units, system workload can be shared by partner RAID controllers. Logical units can be manually assigned to different controllers (Controller A or Controller B and consequently Primary or Secondary) to facilitate the active-active load-sharing configuration.

6. Host channel IDs are designated either as a Controller A or as a Controller B ID. The controller A/B IDs then function as the designators for workload assigned to different RAID controllers.

7. Each logical drive can be configured in a different RAID level.

8. Several logical drives can be striped together to compose a larger logical volume. A logical volume then becomes the basic configuration unit for host LUN mapping and capacity management.

9. 9-1. Each of the logical units (a logical drive, logical volume, or one of their partitions) can be made available on one or more host ports using the host LUN mapping function.

   9-2. Each of them can be "**mapped**" or "**associated**" with one or more host ID/LUNs. Each of these associated host ID/LUNs appears to the host operating system as a virtual storage volume.

   The idea is diagrammed as follows:

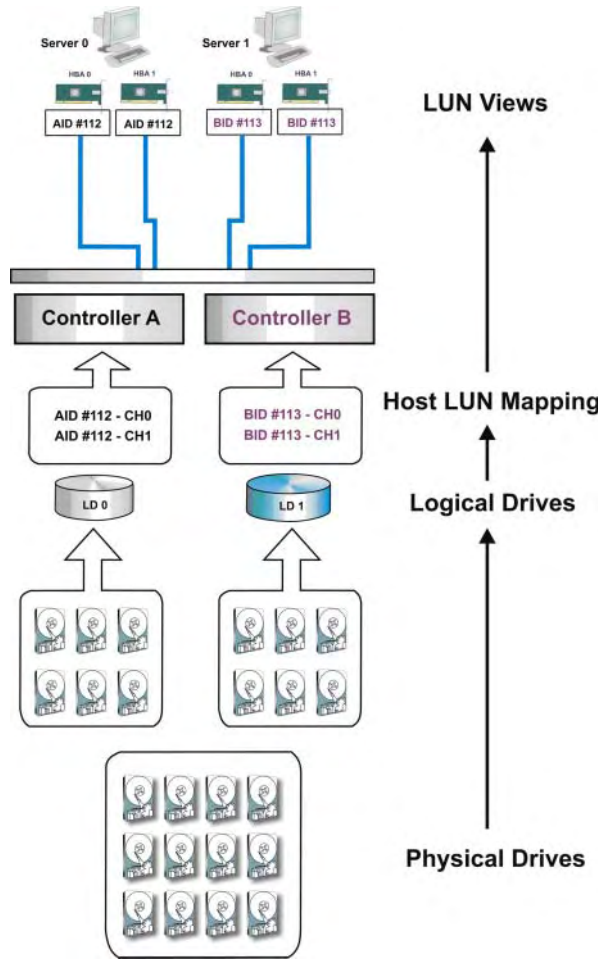

Figure 14 - 3 Grouping Hard Drives and Host LUN Mapping

10. As diagrammed below, array composition can be very flexible. You may divide a logical drive or logical volume into several partitions, or use the entire logical drive as a single partition, with or without the support of spare drives.

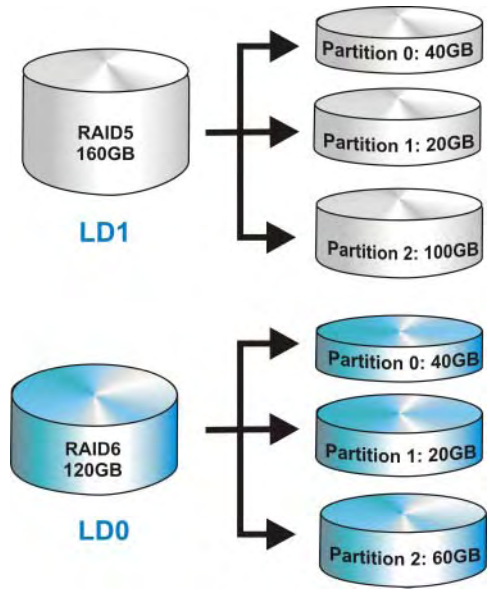

Figure 14-4 Partitioning of Logical Units

11. Each logical unit can be associated (mapped) with one or more host IDs (pre-configured as a Controller A or a Controller B ID) or the LUN numbers under these host IDs.

14.6.3 Host LUN Mapping

Design Concerns

1. When it comes to building a reliable storage solution, redundancy is a virtue. We assume that an environment running mission-critical applications should consist of redundant RAID controllers and multi-pathing software that manage fault-tolerant data paths.
2. Carefully configure your RAID arrays and select the appropriate settings such as stripe size and write policy. Reconfiguration takes time and may require you to move or back-up your data.
3. Create at least two logical drives (LD0 and LD1) and associate (map) them equally with Controller A IDs (AID) and Controller B IDs (BID). Doing so you get the maximum work power from both of the RAID controllers.
4. Logical RAID units are manually associated with Controller A or B IDs that reside on the host channels.
5. Disable some configuration options for they might cause data inconsistency if module failures should occur. For example, disabling the use of buffers on individual disk drives may let you lose some performance, yet it is relatively safer for the drive

buffers may hold cached writes during a power outage and cause data inconsistency.

The configuration can be found in firmware's embedded utility through Main Menu -> View and Edit Configuration Parameters -> Drive-side Parameters -> Drive Delayed Write.

Mapping for Fault-tolerant Links

- **The purpose for mapping a logical drive to multiple IDs is diagrammed below:**

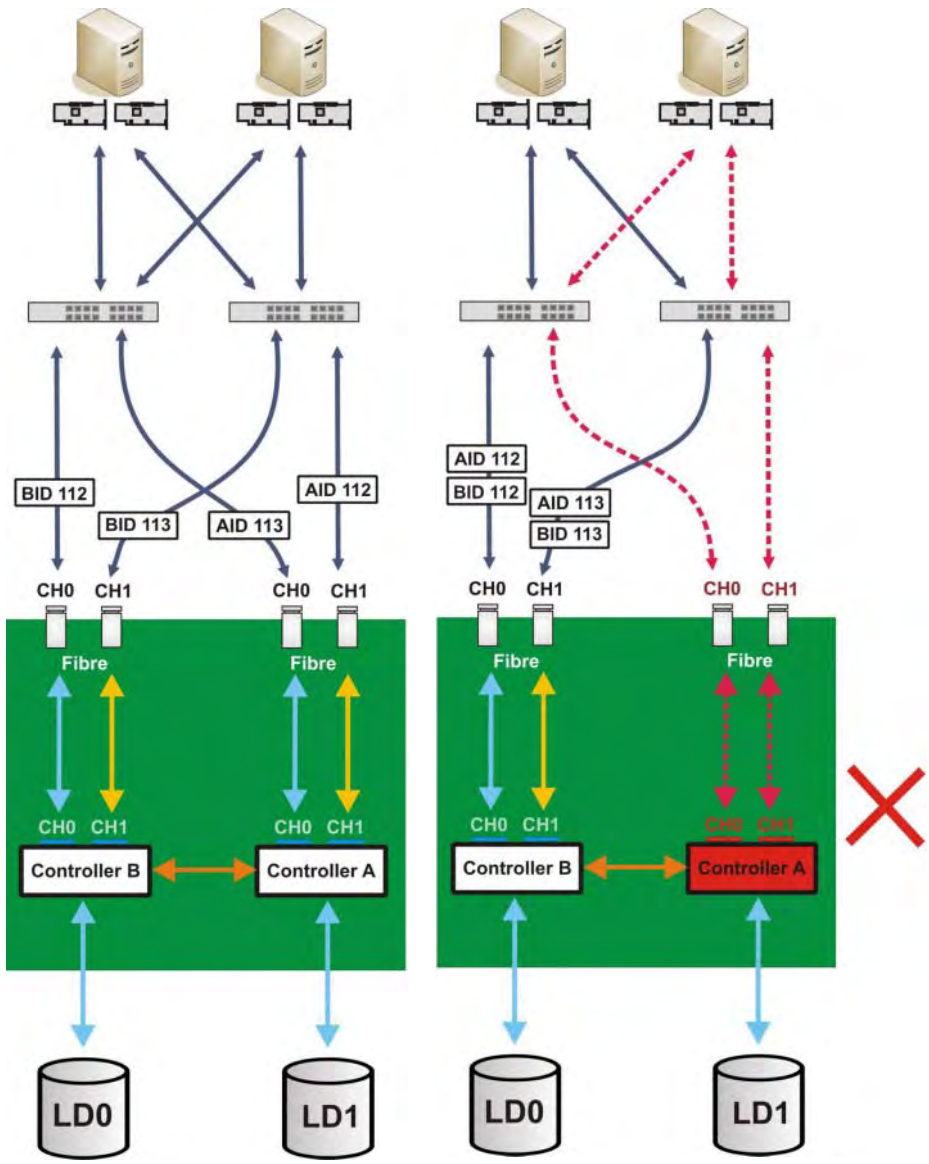

Figure 14-5 Dual-controller Subsystem in Normal and Degraded Modes

In the event of single RAID controller failure, all IDs managed by the failed controller will be taken over by the surviving controller. See the locations of mapped IDs in the above diagram.

If an application server can access the arrays through fault-tolerant paths, multi-path management software, such as Infortrend's EonPath, should be available.

Shown below is a condition with a broken host link. The host computer can still access the array (LD1) through an alternate data link. Even if one of the FC switches fails, access to data can still be continued:

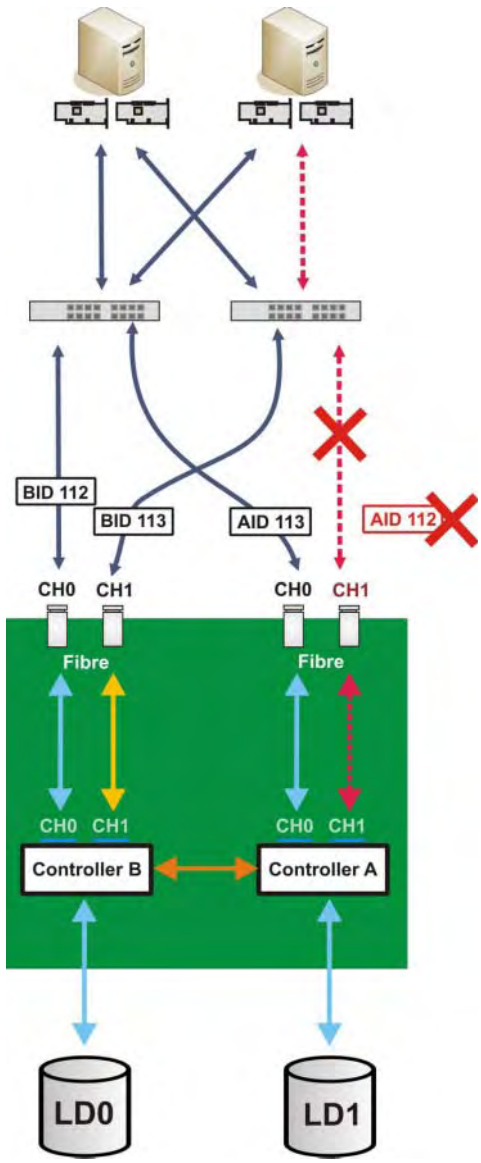

Figure 14-6 Dual-controller Subsystem with a Broken Link

Mapping Using the FW3.48 Cross-controller Mapping

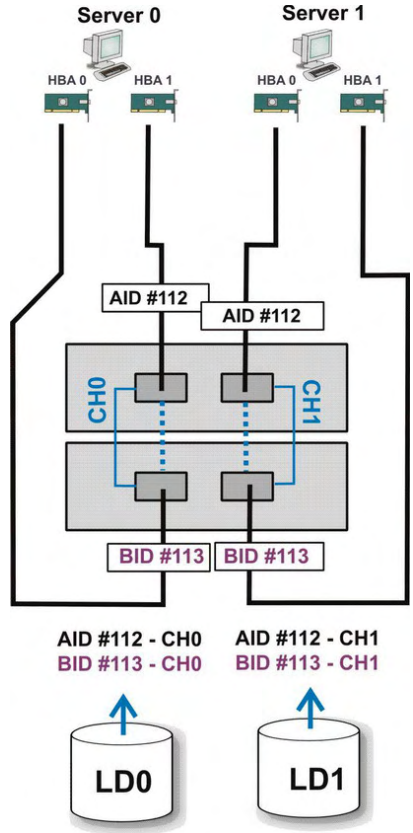

Figure 14-7 Mapping Array Volumes to Different Controllers

- As diagrammed above, each logical drive is associated with two different channel IDs managed by different RAID controllers (AID + BID). This mapping method also ensures continuous host access in the situation when no port bypass is available, e.g., FC switches.

  Note the following when configuring fault-tolerant configurations:

  - Multi-pathing management software should be installed on the host computers to manage the access to the same array volume via two different I/O paths.
  - Each channel ID (or an LUN under target ID) will appear as one virtual storage volume to the host operating system.

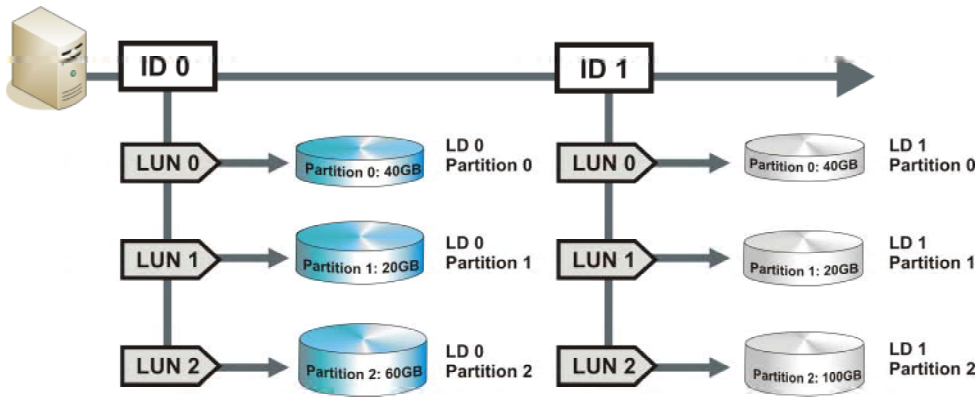

Figure 14-8 Mapping System Drives (Mapping to LUNs)

- Shown below is a host channel bus teamed with multiple IDs/LUNs that are associated with logical partitions.

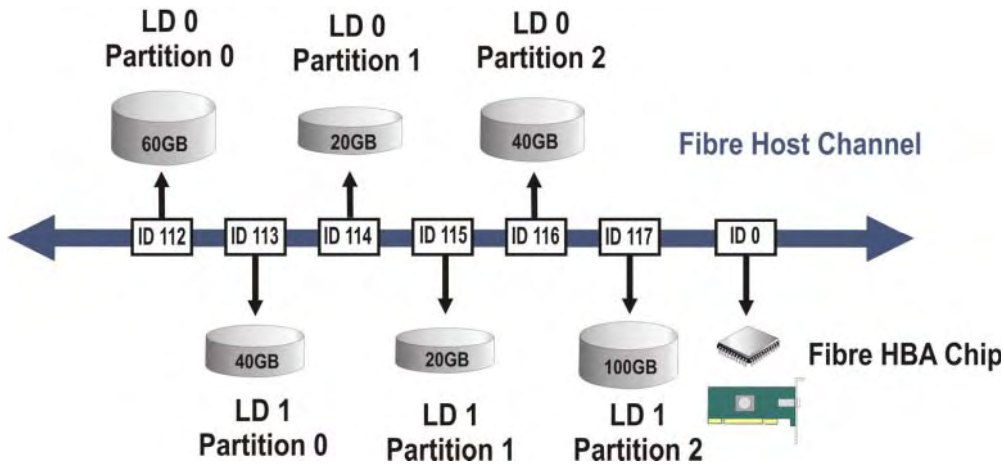

Figure 14-9 Mapping System Drives (an Example with Multiple IDs on FC-AL bus)

- Some older operating systems/HBA cards do not read multiple LUNs under a target ID. As diagrammed above, you may have the host channel to present several IDs and map logical configurations to these IDs. Each of these IDs can be identified as "Controller A ID" or "Controller B ID." As a rule for most operating systems, each configuration unit will then be mapped to LUN 0 under each ID.

14.6.4 Fault Tolerance

Why Using a Redundant Controller Configuration?

1. Hardware failures can occur. A simple parity error can sometimes cause a RAID system to completely hang.
2. Having two controllers working together will guarantee that at least one controller will survive catastrophes and keep the system working.
3. When dealing with high-availability applications, redundancy is always a virtue. This is the logic behind having redundant controllers – to minimize the chance of down time for a storage subsystem.

A redundant-controller system uses two controller modules to manage the storage arrays. It requires two identical controllers to work together and both must be working normally. During normal operation, each controller serves its I/O requests. If one controller fails, the existing controller will temporarily take over for the failed controller. The failover and failback processes are completely transparent to the host (sometimes with the help of intelligent FC switch firmware) and require only minimum efforts to restore the original configuration.

Controller Failover and Failback

In an unlikely event of controller failure, the surviving controller will acknowledge the situation and disconnect from the failed controller. The surviving controller will then act as both controllers and serve all the I/O requests from host.

System failover is transparent to host. System vendors should be contacted for an immediate replacement of the failed unit.

Auto-Failback

Once the failed controller is removed and a replacement controller is installed, the existing controller will acknowledge the situation. The existing controller should automatically attempt to combine with the replacement controller.

When the initialization process of the replacement controller is completed, the replacement controller should always inherit the status of the Secondary controller.

IMPORTANT!

For a subsystem powered by the ASIC266 or later chipset, reset the subsystem if the replaced controller resides in slot A. If the replacement controller in slot A is successfully combined, a system reset should restore its status as the Primary controller.

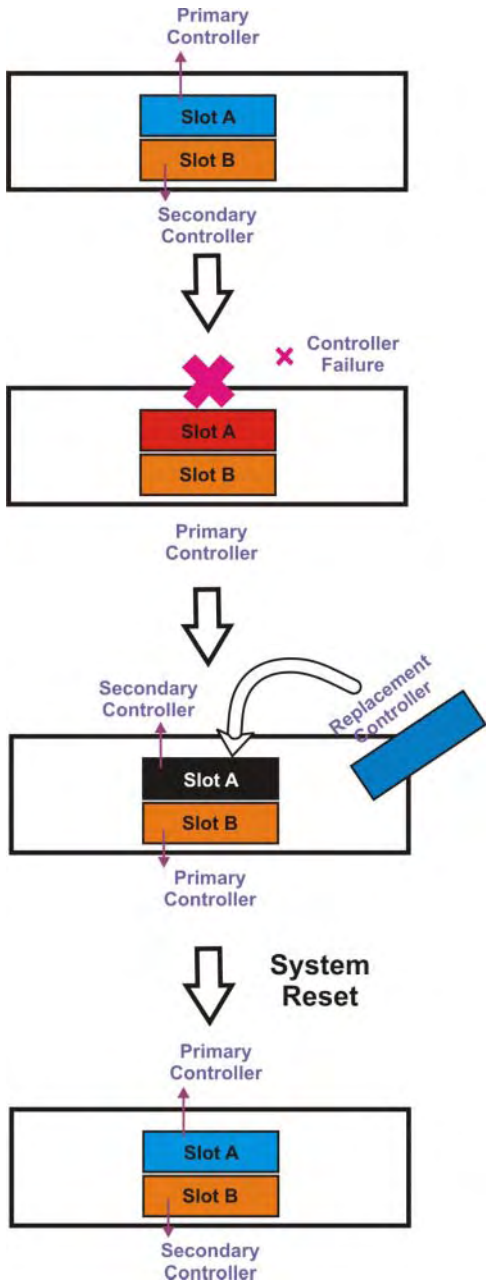

Figure 14-10 Controller Primary-Secondary Relative Positions

A typical controller failover and failback process is diagrammed on the left.

State 1: Subsystem operating normally. Slot A controller is the Primary controller by factory default.

State 2: Slot A controller fails. Slot B controller inherits the Primary role.

State 3: The failed controller in Slot A is replaced by a healthy replacement. The replacement controller becomes the Secondary controller temporarily.

State 4: If the subsystem resets later, the Slot B controller returns the Primary role to the Slot A controller.

If the subsystem is reset later, the controller installed in the Slot A position will obtain the Primary controller status. The Slot B controller then resumes the Secondary role. The replacement controller will obtain all related configuration parameters from its counterpart.

Controller Failure

Controller failure is managed by the surviving controller (regardless of its original role as Primary or Secondary). The surviving controller disconnects from its counterpart while gaining access to all signal paths. The existing controller then proceeds with the ensuing event notifications and take-over process.

Symptoms

- The LCD screen displays controller failure message.
- The surviving controller sounds an alarm.
- The "ATTEN" LED is flashing on the front panel.
- The surviving controller sends event messages to notify of controller failure (indicating its partner has failed).

14.7 Configuration Samples

Design Concerns

1. We assume that an environment running mission-critical applications should consist redundant RAID controllers and multi-pathing software that manages networking devices, such as FC switches or HBAs in fault-tolerant pairs.

2. Carefully configure your RAID arrays and select the appropriate array settings such as stripe size and write policy. Reconfiguration takes time and may require you to move or back-up your data.

3. Create at least two logical drives (LD0 and LD1) and associate (map) them equally with Controller A IDs (AID) and Controller B IDs (BID). Doing so you get the maximum work power from both of the RAID controllers. For more details on creating AIDs/BIDs and LUN mapping processes, please refer to the discussions later in this chapter.

4. Logical RAID units are manually associated with Controller A or B IDs that reside on the host channels.

5. Disable some configuration options for they might cause data inconsistency if module failures should occur. For example, disabling the use of buffers on individual disk drives may let you lose some performance, yet it is relatively safer for drive buffers may hold cached writes during a power outage and cause data inconsistency.

   The configuration option can be found in firmware's embedded utility through **Main Menu** -> **View and Edit Configuration Parameters** -> **Drive-side Parameters** -> **Drive Delayed Write**.

There are similar concerns with the mirrored cache between the RAID controllers. Make sure compensatory measures are applied, e.g., use of battery backup modules or UPS devices.

Pros and Cons of Various Configurations

Configuration		Pros and Cons
1	**Simple DAS w/o Hub**	Applies to single logical drive over flexible cabling.
2	**DAS w/ Hubbed Ports**	DAS without FC switches; total host-side bandwidth can be halved by combining two host ports into a common host loop.
3	**SAN w/ FC Switches**	Applies to multi-server SAN; requires external FC switches.
5	**Multi-pathing w/ Clustered Servers**	High redundancy on server side and on the storage side. I/O path re-routing is partially managed by FC switches.

Table 14-1 Pros and Cons of Sample Topologies

1. Simple DAS without Hub (FW 3.48 Cross-controller Mapping Method)

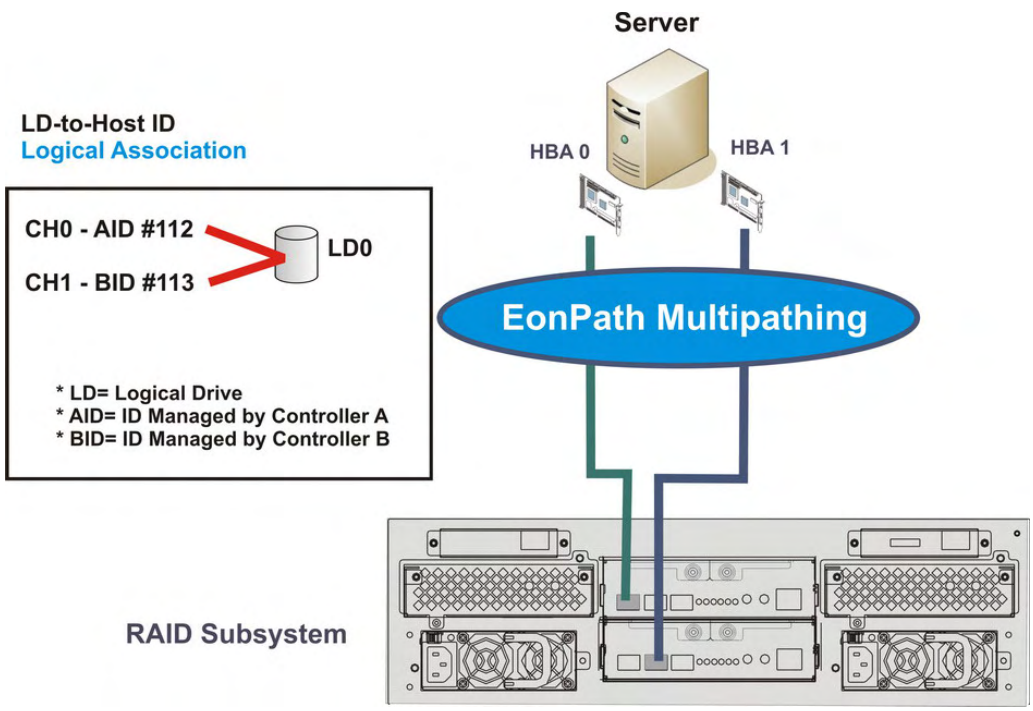

Figure 14-11: Simple DAS without Hub

Configuration Tasks

Tasks	Logical Drive	LUNs	Channel	AID	BID
Map LD0 to an AID on channel #0.	LD0	0	0	112	N/A
Map LD0 to a BID on channel #1 for redundant-path access.		0	1	N/A	113

This configuration applies to a dual-controller subsystem directly attached to a host computer without intermediate networking devices. A logical drive is associated with different Controller IDs (Controller A and Controller B IDs) on separate host channels and different RAID controllers.

In the event of cabling or controller failure, host can still access the array.

NOTE:

- You may use different channel IDs than are shown here in the sample topologies, IDs used in the sample configurations are mostly default numbers in firmware. As long as the IDs are carefully selected according to the configuration rules, there is no limitation on using different host channel IDs.
- A logical drive is associated with both a Controller A and a Controller B ID. This methodology applies when no onboard or external bypass is available.
- You may use the onboard hub to combine two host ports into a common host loop. Then you may not need the cross-controller mapping.

2. DAS with Hubbed Ports

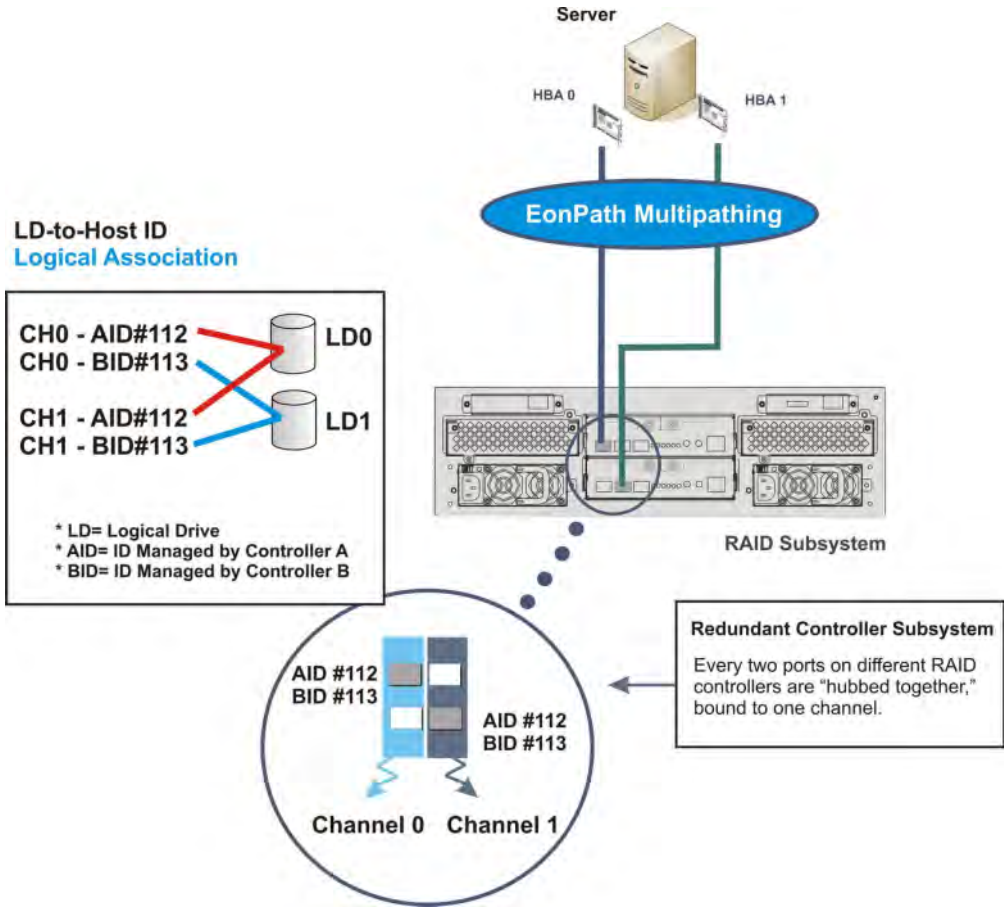

Figure 14-12: DAS with Hubbed Ports

Configuration Tasks

Tasks	Logical Drive	LUNs	Channel	AID	BID
Map LD0 to an AID on channel #0.	LD0	0	0	112	N/A
Map LD0 to an AID on channel #1 for path redundancy.		0	1	113	N/A
Map LD1 to a BID on channel #0.	LD1	0	0	N/A	113
Map LD1 to a BID on channel #1 for path redundancy.		0	1	N/A	112

1. This configuration applies to a redundant-controller subsystem directly attached to a host computer without intermediate FC switch devices.

2. Two logical drives are separately associated with Controller A or Controller B IDs on separate host channels.

3. The onboard hub combines the host ports of the same channel number on different controllers into a single host loop. When the onboard hub is enabled, both the Controller A and Controller B IDs can be presented through any of the SFP ports in a host loop.

4. See your Hardware Manual for how to enable the onboard hub.

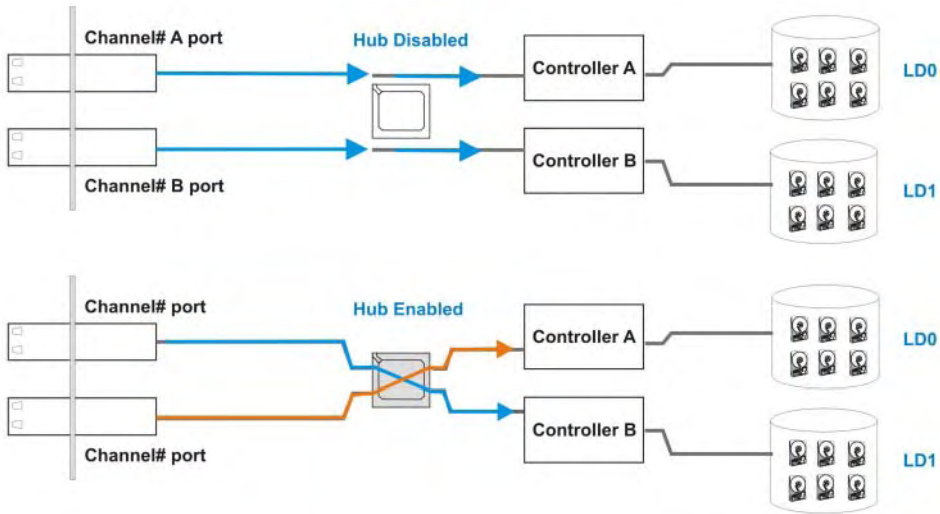

Figure 14-13: The Onboard Hub Function

3. SAN with FC Switches

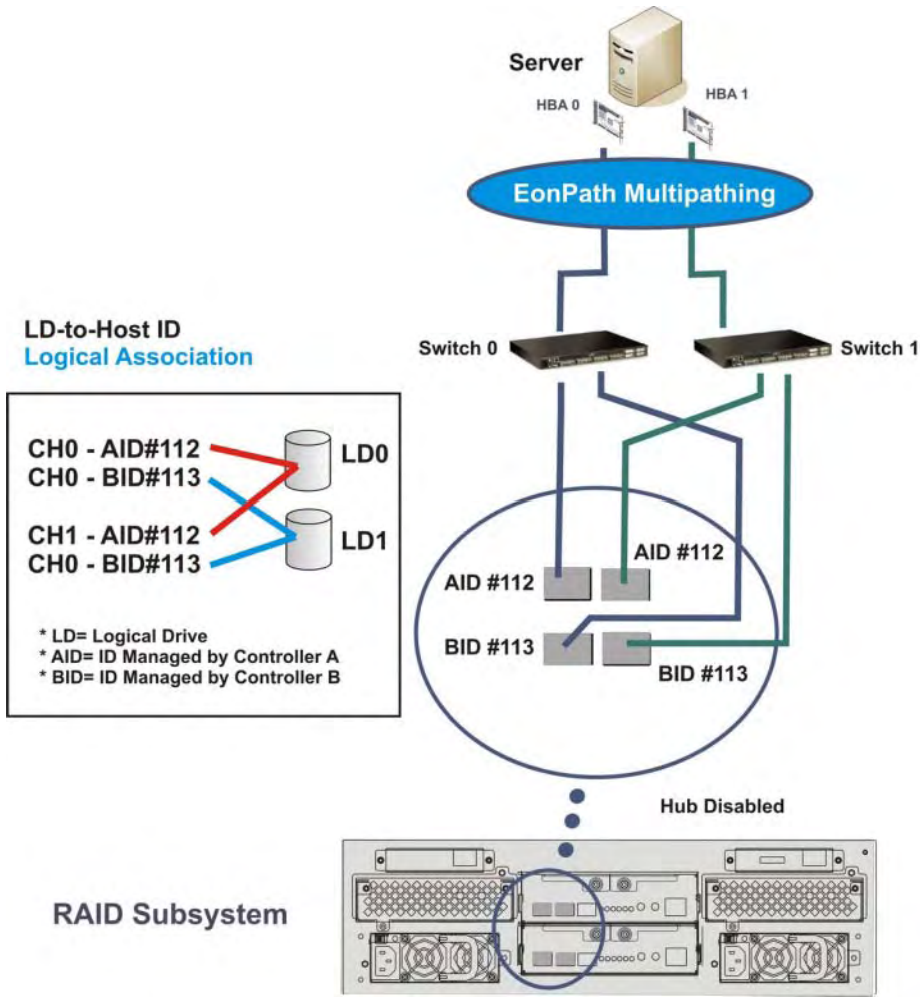

Figure 14-14: SAN with FC Switches

Shown above is a configuration using FC switches to facilitate the connections with multiple SAN servers. For the reason with diagram's simplicity, only one server is displayed.

Configuration Tasks

Tasks	Logical Drive	LUNs	Channel	AID	BID
Map LD0 to an AID on channel #0.	LD0	0	0	112	N/A
Map LD0 to an AID on channel #1 for path redundancy.		0	1	113	N/A
Map LD1 to a BID on channel #0.	LD1	0	0	N/A	113
Map LD1 to a BID on channel #1 for path redundancy.		0	1	N/A	112

This configuration applies to a redundant-controller subsystem attached to switched fabric and then to application server(s).

Fault Tolerance is achieved through the following:

1. Logical drives are separately associated either with the Controller A IDs or Controller B IDs on separate host channels.

2. In the event of a controller failure, the surviving controller inherits IDs from the failed controller. Host IDs managed by a failed controller are automatically passed down to a surviving RAID controller. For instance, Controller A IDs will be managed by the Controller B if Controller A fails.

3. In the event of cabling failure, an array is access through the alternate data path through an alternate host ID.

4. Through the intermediate FC switches or switch zoning, cable/controller failure can be managed by re-routing host I/Os to a valid link.

5. When attached to switched fabrics, the subsystem's onboard hub function should be disabled.

4. Multi-pathing with Clustered Servers (FW3.48 Cross-controller Mapping Method)

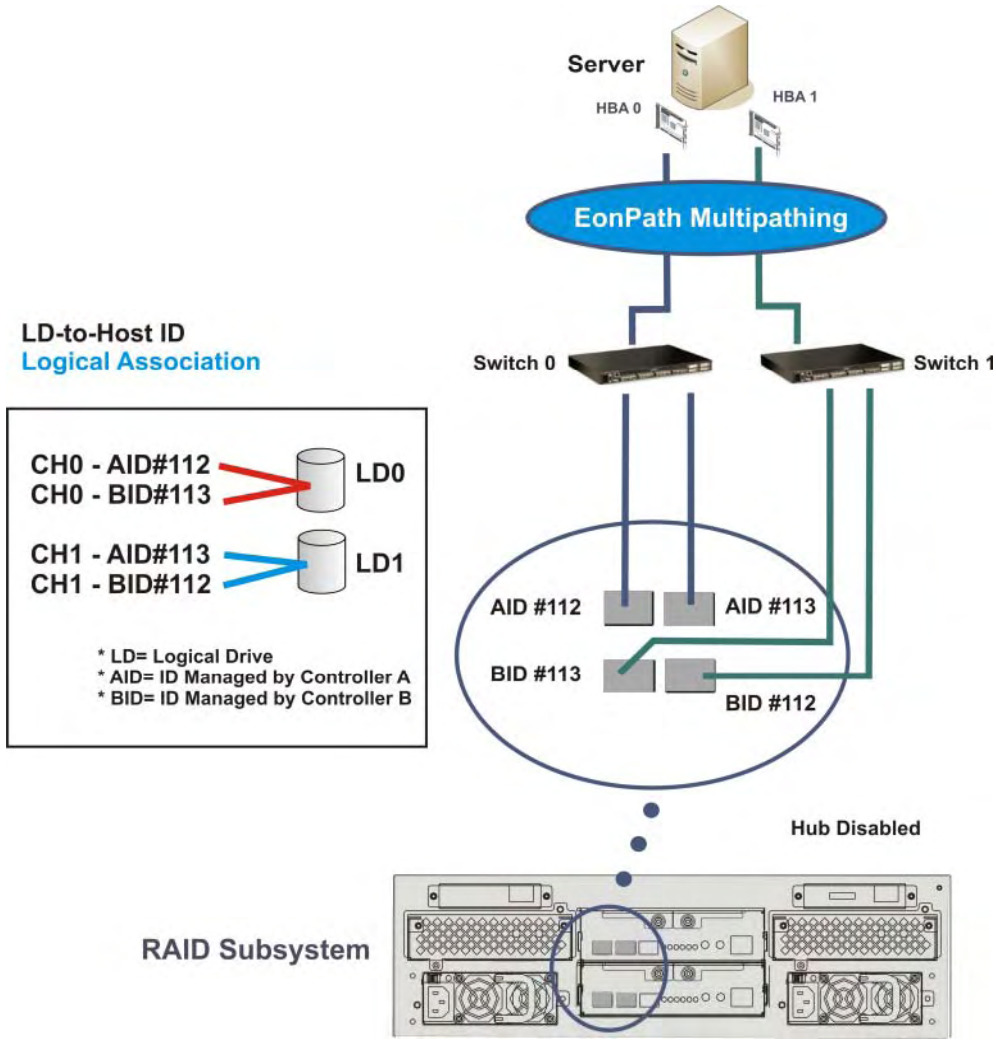

Figure 14-15: SAN with Clustered Servers

Configuration Tasks

Tasks	Logical Drive	LUNs	Channel	AID	BID
Map LD0 to an AID on channel #0.	LD0	0	0	112	N/A
Map LD0 to a BID on channel #0 for path redundancy.		0	0	N/A	113
Map LD1 to an AID on channel #1.	LD1	0	1	113	N/A

Map LD1 to a BID on channel #1 for path redundancy.		0	1	N/A	112

The EonPath multi-pathing software is installed on both of the clustered servers to manage the fault-tolerant data paths.

This page is intentionally left blank.

Appendix

A

Firmware Functionality Specifications

A.1 Basic RAID Management:

Specification	Feature
RAID levels	0, 1(0+1), 3, 5, 6, 10, 30, 50, 60, JBOD and NRAID. Levels 10, 30, 50, and 60 are the multi-level RAID defined as the logical volume implementations; logical volumes consist of logical drives of different RAID levels that are striped together.
Maximum number of logical drives	16 as default, up to 32 with a 1GB or above memory size
Maximum logical drive capacity	64TB
RAID level dependency to each logical drive	Independent. Logical drives configured in different RAID levels can co-exist in a logical volume and within a RAID subsystem
Maximum number of logical drive members	128 with 512MB memory size 112 with 256MB memory size; (theoretical number, not recommended for the difficulties with backup, rebuild, and management tasks)
Configurable stripe size	16KB, 32KB, 64KB, 128KB, 256KB, 512KB, or 1024KB per logical drive
Configurable Write Policy (write policy per array)	Write-Back or Write-Through per logical drive. This policy can be modified later.
Logical drive identification	Unique, controller randomly generated logical drive ID; Logical drive name user-configurable for ease of identification in a mult-array configuration
Maximum number of partitions for each logical drive	8 as default, up to 64 with a 1GB memory size

Maximum number of logical volumes	16 as default, up to 16 with a 1GB or above memory size
Maximum number of LUNs Mappable	Default is 128, up to 1024 with a 1GB or above memory size
Maximum number of LUNs per host ID	Up to 32, user configurable
Maximum number of Media Scan task schedules	16
Concurrent I/O	Supported
Tag Command Queuing (TCQ)	Supported
Native Command Queuing (NCQ)	Supported
Dedicated spare drive	Supported, hereby defined as the spare drive specifically assigned to a logical drive. Also known as Local Spare
Global spare drive	Supported, the spare drive that serves all logical drives (as long as it is equal in size or larger than logical drive members)
Global spare auto-assign	Supported, applies to non-configured drive(s); safeguards the array if a spare has been used in the previous array rebuild and users forget to configure a new drive as a spare.
Enclosure spare drive	A Spare that participates in the rebuild of the failed drive within the same enclosure.
Co-existing Dedicated (Local), Enclosure-specific, and Global spare drives	Supported
Auto-rebuild onto spare drive	Supported
Auto-scan of replacement drive upon manually initiated rebuild	Supported
One-step rebuild onto a replacement drive	Supported
Immediate logical drive availability	Supported; Logical arrays are immediately ready for Host I/Os. Initialization task is completed in the background except when the logical array is stated as “INCOMPLETE” or “BAD;” e.g., has a failed member right after the creation.
Auto-rebuild onto failed drive replacement	Supported. With no spare drive, the subsystem will auto-scan the failed drive and starts rebuild automatically once the failed drive has been replaced.
Concurrent rebuild / expansion	Multiple logical drives can proceed with a Rebuild/Regenerating Parity, and/or Expansion/Initialization/Add Drive operation at the same time.

	NOTE: Regenerate Parity and Rebuild cannot take place on a logical drive at the same time. Create, Expand, and Add Drive operations cannot take place on a logical drive at the same time.
Background firmware download	Firmware can be downloaded during active I/Os, and takes effect after system reboot.
Auto recovery from logical drive failure (configuration on drives)	Supported. If a user accidentally removed the wrong drive to cause the 2nd drive failure of a one-drive-failed RAID5 / RAID3 logical drive, fatal error may occur. However, you may force the subsystem to reaccept the logical drive by switching off the subsystem, installing the drive back to its original drive slot, and then power on the subsystem. You may have the chance to restore the logical drive into the one-drive-failed status.

NOTE:

To ensure smooth operation, sufficient cache memory buffer is required for configurations made up of numerous logical units. An intelligent trigger mechanism is implemented with firmware version 3.47 and later. If a subsystem/controller comes with a DIMM module of the size equal or larger than 1GB, firmware automatically enlarges the maximum numbers of logical units.

	DIMM size < 1G	DIMM size >= 1G
Max. no. of LD	16	32
Max. no. of LV	8	16
Max. partitions per LD/LV	16	64
Max. no. of LUN	128	1024

A.2 Advanced Features:

Media Scan	Supported. Verify written data on drives to avoid bad blocks from causing data inconsistency. If bad blocks are found, data can be reconstructed by comparing and recalculating parity from adjacent drives (RAID1/3/5/6). The “Reconstruction Writes” are followed by “Write Verification” operation.

Bad Block Handling in degraded mode	A method for handling low quality drives. The operation is performed on both the logical drive in degraded mode or those that are being rebuilt. If bad blocks should be encountered during Rebuild, Add Drive, Host Write, or Regenerate Parity operation, the controller will first attempt to reconstruct affected data and those unrecoverable bad blocks are stated as bad and the controller return to host. Users have the option to abandon data on the unrecoverable sectors to continue rebuild in a degraded mode. Low quality drive handling comes with transparent resetting of hung hard drives.
Transparent reset of hung HDDs	Supported
Auto cache flush on critical conditions **(caching mode dynamic switch)**	When critical conditions occur, e.g., component failure, or BBU under charge, cached data will be flushed and the write policy will be changed to write-through mode. Configurable “Trigger Events” for Write-through/Write-Back Dynamic Switch. The configuration can also be set with the OEM “lappend” utility.
RAID parity update tracking and recovery	Yes, to avoid write holes.
Host-side Ordered Tag support	Supports write commands with embedded Ordered Tags.
Drive identification (flash drive function)	Supported. Force a drive to light on its activity indicator for users to visually recognize its position in a configuration consisting of numerous disk drives.
Drive information listing	Supported. Drive vendor name, model number, firmware revision, capacity (blocks), serial number, narrow/wide and current sync. speed
Drive read/write test	Supported
Configuration on disks (Drive Roaming)	Supported. The logical drive information is recorded on drive media. The logical drives can still be accessed if using different Infortrend RAID controllers/subsystems, e.g., drives removed and installed in a different subsystem.
Save/ restore NVRAM to / from disks	Supported. Save all the settings stored in the controller NVRAM to the logical drive members. Now this feature comes with an option whether to restore the previously saved password in case an administrator changed the password some time before.

Save / restore NVRAM to / from a file	Supported. Save all the settings stored in the controller NVRAM to a file (via GUI manager) on user's computer. Now this feature comes with an option whether to restore the previously saved password in case an administrator changed the password some time before.
Host-side 64-bit LBA support	Supports array configuration (logical drive, logical volume, or a partition of them) of a capacity up to 64TB.
Host LUN geometry: user configurable default geometry (Solaris OSes)	On FW 3.48, this feature comes with preset combinations of head, cylinder, and sector variables.
User configurable geometry range:	Sector: 32, 64, 127, 255 or Variable Head: 64, 127, 255 or Variable Cylinder: <1024, <32784,<65536 or Variable
Drive motor spin-up	Supported. The controller will send spin-up (start unit) command to each drive at the 4 sec. intervals.
Drive-side tagged command queuing	Supported. User adjustable up to 128 for each drive.
Host-side maximum queued I/O count	User adjustable up to 1024
Maximum concurrent host LUN connection	User adjustable up to 1024
Number of tags reserved for each Host-LUN connection	User adjustable up to 256
Controller shutdown	Flushes cached contents upon the detection of critical conditions, e.g., a high temperature condition persists for a long time.
Drive I/O timeout	User adjustable
I/O channel diagnostics	Supported; please contact your dealer for more details.
Drive Spindown Idle Delay	Allows users to spin down hard drives when the subsystem has been idle for certain period of time.
Maximum Drive Response Time (Guaranteed Latency I/O)	User adjustable from 160 to 960ms. If a disk drive fails to return data on read requests before the timeout value is exceeded, the array immediately generates data from the parity data and the other members of a logical drive.

A.3 Caching Operation:

Write-back cache	Supported.
Write-through cache	Supported.

Supported memory type	DDR memory for enhanced performance. Fast Page Memory with Parity for enhanced data security.	
Read-ahead operation	Intelligent and dynamic read-ahead operation for processing sequential data requests.	
Multi-threaded operation	Yes, internal parameters adjusted in accordance with the number of outstanding I/Os.	
Scatter / Gather	Supported	
I/O sorting	Supported. Optimized I/O sorting for enhanced performance.	
Adaptive Write-back/Write-through switching	For a better performance when handling large sequential writes, firmware temporarily disables write-back cache and the synchronized cache operation between partner controllers if operating with dual-active RAID controllers. Firmware automatically restores the write-back mode when encountering random and small writes later.	
Periodic Cache Flush	Firmware can be configured to flush the cached contents in memory at every preset interval: 1. If data integrity is of the concern, e.g., the lack of a battery backup protection. 2. Cache flush on preset intervals to avoid the latency when cache memory is full due to write delays.	
Variable stripe size	**RAID0**	128
	RAID1	128
	RAID3	16
	RAID5	128
	RAID6	128
Caching Optimization		
• Cache buffer sorting prior to cache flush operation.		
• Gathering of writes during flush operation to minimize the number of I/Os required for parity update.		
• Elevator sorting and gathering of drive I/Os.		
• Multiple concurrent drive I/Os (tagged commands).		
• Intelligent, predictive multi-threaded read-aheads.		
• Multiple, concurrent host I/O threads (host command queuing).		

A.4 RAID Expansion:

On-line RAID expansion	Supported. Capacity brought by array expansion is immediately ready for Host I/Os when its status changes from “EXPAND” to “INITIALIZING.” Initialization task is then completed in the background except when the logical array is stated as “INCOMPLETE” or “BAD;” e.g., has a failed member right after creation.
Mode-1 RAID expansion - add drive	Supported. Multiple drives can be added concurrently. Though not recommended, Add Drive can even be performed in the degraded mode.
Mode-2 RAID expansion – copy and replace drives	Supported. Replace members with drives of larger capacity.
Expand capacity with no extra drive bays required	Supported in Mode 2 RAID expansion, which provides “Copy and Replace Drive” function to replace drives with drives of greater capacity. Protect your investment for there is NO need for hardware upgrade, e.g., adding a new enclosure for the extra drives.
Operating system support for RAID expansion	No. No operating system driver required. No software needs to be installed for this purpose.

A.5 S.M.A.R.T. Support:

Copy & replace drive	Supported. User can choose to clone a member drive showing symptoms of defects before it fails.
Drive S.M.A.R.T. support	Supported, with intelligent error handling implementations.
User selectable modes for S.M.A.R.T.	1. Detect only 2. Perpetual Clone: using a hot-spare to clone the drive reporting SMART errors; the hot-spare remains a clone drive 3. Clone + Replace: using a hot-spare to replace the drive reporting SMART errors; the drive reporting errors is pulled offline

A.6 Redundant Controller:

Active-active redundant controller	Supported
Synchronized cache	Supported. Through one or multiple, dedicated synchronizing channels on a common backplane or external cabling. Synchronized cache over SCSI channels, Fibre loops, or SATA channels is supported. Synchronized cache can be disabled via a UI option when using write-through mode in a redundant controller configuration to prevent performance trade-offs.
Write-back cache enabled in redundant controller mode	Yes, with synchronized cache connection and mirrored cache between controllers.
Automatic failover	Yes (user's interaction necessary; e.g., to restart the software management console)
Automatic failback	Yes (user's interaction necessary)
Controller hot-swap	▪ No need to shut down the failed controller before replacing the failed controller. ▪ Support online hot-swap of the failed controller. There is no need to reset or shutdown the failed controller. One controller can be pulled out during active I/Os to simulate the destructive controller failure.
Parity synchronization in redundant controller write-back mode to avoid write-hole	Supported.
Redundant controller communication over Fibre loops	Dedicated loops or distribution over drive loops selectable. Preset SATA paths or SAS domains on EonStor subsystems.
No single-point-of-failure	Supported.
Automatic engagement of replacement controller	Supported.

Dynamic cache memory allocation	Yes. Cache memory is dynamically allocated, not fixed.
Environment management	Supported. SAF-TE, S.E.S., ISEMS (I^2C interface), or S.E.S. over SAS links; and on-board controller voltage/temp monitor are all supported in both single and redundant controller mode. In the event of controller failure, services can be taken over by the surviving controller.
Cache battery backup	Supported. Battery backup modules support cached data during a power outage in both single controller and redundant controller modes. With EEPROM battery modules, firmware will be aware of the life expectancy of battery cells. This feature is supported since ASIC400 models.
Load sharing	Supported. Workload can be flexibly divided between different controllers by assigning logical configurations of drives (LDs/LVs) to different RAID controllers.
User configurable channel mode	Supported. Channel modes configurable (SCSI or Fibre) as HOST or DRIVE on specific models such as the EonRAID 2510FS and EonStor F16F.
Require a special firmware for redundant controller?	No.

A.7 Data Safety:

Data Service: Snapshot	The SANWatch software provides management interfaces for the Snapshot functionality. Please refer to SANWatch User's Manual for details.
Regenerate parity of logical drives	Supported. Can be manually executed to ensure that bad sectors do not cause data loss in the event of drive failure.
Scheduled Media Scan	Revision 3.34 and above allows Media Scan to be scheduled starting at a specified start time and repeated at regularly timed intervals. The start time and time intervals can be selected from drop-down menus. Start time is manually entered using its numeric representatives in the following order [MMDDhhmm[YYYY]], and it reads the date and time set for the controller's real-time clock. The selectable time intervals (the Execution Period) range from one (1) second to seven (7) weeks. Each such schedule can be defined to operate on individual hard drives, all members of a specified logical drive, or members of selected

	logical drives. Each schedule can include up to five (5) logical drives. The RS-232C terminal interface and RAIDWatch revision 2.0 support this functionality.
Bad block auto-reassignment	Supported. Automatic reassignment of bad block
Battery backup for cache memory	Supported. The battery backup unit supports cache memory when power failure occurs. The unwritten data in the cache memory can be committed to drive media when power is restored.
Verification on normal writes	Supported. Performs read-after-write during normal write processes to ensure data is properly written to drives.
Verification on rebuild writes	Supported. Performs read-after-write during rebuild write to ensure data is properly written to drives.
Verification on LD initialization writes	Supported. Performs read-after-write during logical drive initialization to ensure data is properly written to drives.
Drive S.M.A.R.T. support	Supported. Drive failure is predictable with reference to the different variables detected. Reaction schemes are selectable from Detect only, Perpetual Clone and Copy + Replace. These options help to improve MTBF.
Clone failing drive	Users may choose to clone data from a failing drive to a backup drive manually.
Automatic shutdown on over-temperature condition	Controller automatically enters an idle state (stops answering I/O requests) upon the detection of high-ambient temperature for an extended period of time.

A.8 System Security:

Password protection	Supported. All configuration changes require the correct password (if set) to ensure system security. Password protection is also bundled with all user interfaces.
User-configurable password validation timeout	Supported. After certain time in absence of user interaction, the password will be requested again. This helps to avoid unauthorized operation when user is away.
SSL-enabled RAIDWatch Agents	Agents communicate to the controller through limited set of authorization options.

A.9 Environment Management:

SAF-TE/S.E.S. support	Supported. The SAF-TE/S.E.S. modules can be connected to the drive channels. The RAID controller will detect errors from SAF-TE/S.E.S. devices or notify drive failures via SAF-TE/S.E.S. • Both SAF-TE/S.E.S. via drive and device-self-interfaced methods are supported. • Redundant SAF-TE/S.E.S. devices are supported • Multiple S.E.S. devices are supported
Dynamic on-lining of enclosure services	Once an expansion unit (JBOD) with supported monitoring interface is combined with a RAID system, its status will be automatically polled.
SAF-TE/S.E.S. polling period	User configurable (50ms, 100ms, 200ms, 500ms, 1~60sec)
ISEMS (Infortrend Simple Enclosure Management Service)	Supported via an I2C serial bus.
Multiple SAF-TE/S.E.S. modules on the same channel	Supported.
Multiple SAF-TE /S.E.S. modules on different channels	Supported.
Mapping SAF-TE/S.E.S. device to host channel for use with host-based SAF-TE/S.E.S. monitoring	Supported.
Event Triggered Operation	When any of the following happens, the firmware disables write-back caching to minimize the chance of losing data: - Battery, controller, cooling fan, or PSU failure - The upper temperature thresholds are exceeded - Low battery charge - UPS AC loss or low battery charge The triggering factors are user-configurable
Dual-speed cooling fan control	Yes, firmware triggers high rotation speed in the event of elevated temperature or component failure, e.g., a fan failure.
Dual-LED drive status indicators	Supported. Both single-LED and dual-LED drive status indicators are supported.
SAF-TE/ S.E.S. temperature value display	Supported. Display the temperature value provided by enclosure SAF-TE/S.E.S. module (if available).
On-board controller	Supported. Monitors the 3.3V, 5V, and 12V

voltage monitors	voltage status. Event triggered thresholds user configurable.
On-board controller temperature sensors	Supported. Monitors the CPU and board temperature status. Event trigger threshold user configurable.
Enclosure redundant power supply status monitoring	Supported. SAF-TE/S.E.S./ISEMS
Enclosure fan status monitoring	Supported. SAF-TE/S.E.S/ISEMS
Enclosure UPS status monitoring	Supported. SAF-TE/S.E.S/ISEMS
Enclosure temperature monitoring	Supported. SAF-TE/S.E.S/ISEMS

A.10 User Interface:

RAIDWatch on-board	Out-of-band configuration and monitoring via Ethernet. Browser accessible configuration options by installing RAIDWatch to reserved space on drive via ftp. FW 3.48 and later revisions support a browser-invoked, http-based RAIDWatch GUI that requires no installation efforts.
RS-232C terminal	Supports terminal modes: ANSI, VT-100, ANSI Color. Provides menu-driven user-friendly text-based interface.
Graphical user interface (Java-based GUI manager)	Provides user-friendly graphical interface. Communicates with RAID controller via Out-of-band Ethernet, In-band SCSI, In-band Fibre or SNMP traps.
SSH support	Secure Shell over Telnet supported
External interface API for customized host-based management	Supported.
LCD front panel	Provides easy access for user instinct operation.
Buzzer alarm	Warns users when any failures or critical events occur.

A.11 High Availability:

Custom inquiry serial number	Custom Inquiry Serial Number (for support of multi-pathing software like Veritas, QLogic, etc).
Continuous rebuild	Rebuild automatically continues if power outage or operator errors occur during a rebuild.
Asymmetric Logical Unit Access (or later known as Target Port Group Service)	Support for multipath drivers to select an optimal I/O path and for more flexible utilization of internal I/O paths in the event of path failure or controller failover/failback.

This page is intentionally left blank.

Appendix

B

System Functions: Upgrading Firmware

B.1 Upgrading Firmware

The RAID controller's firmware resides in flash memory that can be updated through the COM port, LAN port, or via In-band SCSI/Fibre. New releases of firmware are available in the form of a DOS file in the "pub" directory of Infortrend's FTP site or on a 3.5" diskette. The file available at the FTP site is usually a self-extracting file that contains the following:

FW30Dxyz	Firmware Binary (where "xyz" refers to the firmware version)
B30Buvw	Boot Record Binary (where "uvw" refers to the boot record version)
README.TXT	Read this file first before upgrading the firmware/boot record. It contains the most up-to-date information which is very important to the firmware upgrade and usage.

These files must be extracted from the compressed file and copied to a directory in boot drive.

B.1.1 Background RS-232C Firmware Download: Single Controller

Host I/Os will not be interrupted during the download process. After the download process is completed, user should find a chance to reset the controller for the new firmware to take effect.

B.1.2 Note for Redundant Controller Firmware Upgrade:

A controller used to replace a failed unit in a dual-controller system is often running a newer release of firmware version. To solve the contention, make sure the firmware on a replacement controller is downgraded to that running on the surviving controller.

IMPORTANT!

- Allow the downloading process to finish. Do not reset or turn off the computer or the controller while it is downloading the file. Doing so may result in an unrecoverable error that requires the service of the manufacturer.

- When upgrading the firmware, check the boot record version that comes with it. If the boot record version is different from the one installed on the surviving controller previously, the new boot record binary must be installed.

NOTE:

- The COM 2 serial port cannot be used to download firmware.

B.2 Upgrading Firmware Using RAIDWatch Manager

B.2.1 Establish the In-band SCSI connection in RAIDWatch Manager

Please refer to ***RAIDWatch User's Manual*** for details on establishing the management session for RAIDWatch Manager.

B.2.2 Upgrade Both Boot Record and Firmware Binaries

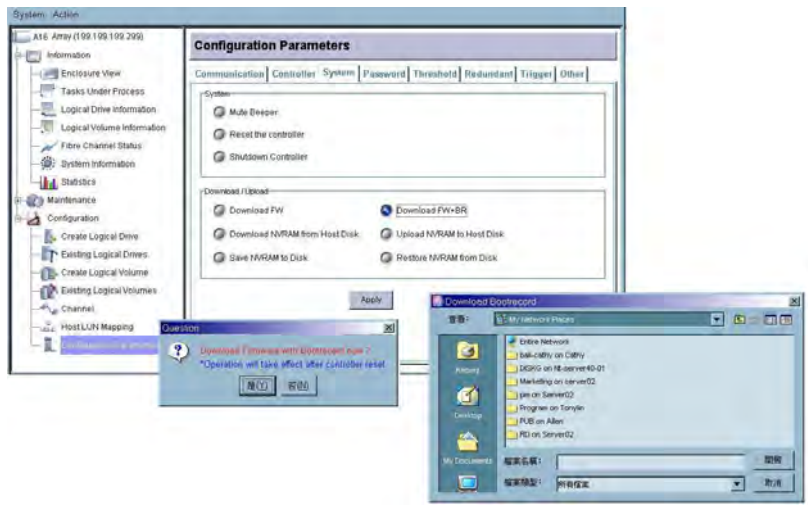

Figure B-1: Upgrading Boot Record and Firmware Binaries Using RAIDWatch Program

1. Connect to the RAID system locally or from a remote computer using RAIDWatch Manager. While connected to the RAID system, there will be icon(s) with IP address specified on the left of the menu screen. Select by double-clicking the icon of the RAID system which firmware is to be upgraded. Select the "Configuration Parameters" icon and then select the "System" tab. Single-click the "Download FW+BR" check circle. Click the Apply button to proceed. A message prompt should display. Click Yes and then a file location prompt will appear.

2. Provide the boot record binary filename, the RAIDWatch Manager will start to download the boot record binary to the controller.

3. After the boot record download is completed, provide the firmware filename to the RAIDWatch Manager. It will start to download the firmware to RAID controller.

4. Shutdown the application server(s) or stop the application(s) which are accessing the RAID, then reset the RAID subsystem/controller in order to use the new downloaded firmware. **With firmware release 3.21 and above**, host I/Os

will not be interrupted by the download process. Users may find a chance to stop host I/Os and reset the controller for new firmware to take effect.

B.2.3 Upgrade the Firmware Binary Only

Figure B-2: Upgrading Firmware Binaries Using RAIDWatch Program

1. Connect to the RAID system locally or from a remote computer using RAIDWatch Manager. While connected to the RAID system, there will be icon(s) with IP address specified on the left of the menu screen. Select by double-clicking the icon of the RAID subsystem which firmware is to be upgraded. Select the "Configuration Parameters" icon and then select the "System" tab. Single-click the "Download FW" check circle. Click the Apply button to proceed. A message prompt should display. Click Yes and then a file location prompt will appear.

2. Provide the firmware filename to the RAIDWatch Manager. It will start to download the firmware to RAID controller.

3. Shutdown the application server(s) or application(s) which are accessing the RAID, then reset the controller in order to use the new downloaded firmware.

B.3 Upgrading Firmware Using RS-232C Terminal Emulation

The firmware can be downloaded to the RAID controller/subsystem by using an ANSI/VT-100 compatible terminal emulation program. Whichever terminal emulation program is used must support the ZMODEM file transfer protocol. The following example uses the HyperTerminal in Windows NT®. Other terminal emulation programs (e.g., Telix and PROCOMM Plus) can perform the firmware upgrade as well.

B.3.1 Establishing the connection for the RS-232C Terminal Emulation

Please refer to ***Chapter 1*** and also your hardware manual for details on establishing the connection.

B.3.2 Upgrading Both Boot Record and Firmware Binaries

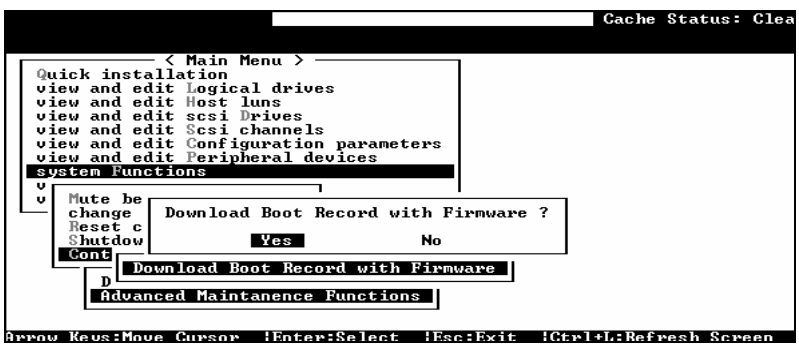

1. From the Main Menu, scroll down to "System Functions."

2. Go to "Controller Maintenance."

3. Choose "Advanced Maintenance."

4. Select "Download Boot Record and Firmware."

5. Set ZMODEM as the file transfer protocol of your terminal emulation software.

6. Send the Boot Record Binary to the controller. In HyperTerminal, go to the "Transfer" menu and choose "Send

file." If you are not using Hyper Terminal, choose "Upload" or "Send" (depending on the software).

7. After the Boot Record has been downloaded, send the Firmware Binary to the controller. In HyperTerminal, go to the "Transfer" menu and choose "Send file." If you are not using Hyper Terminal, choose "Upload" or "Send" (depending on the software).

8. When the Firmware completes downloading, the controller will automatically reset itself. For a newer version of firmware, you need to manually reset the subsystem/controller for the new firmware to take effect.

B.3.3 Upgrading the Firmware Binary Only

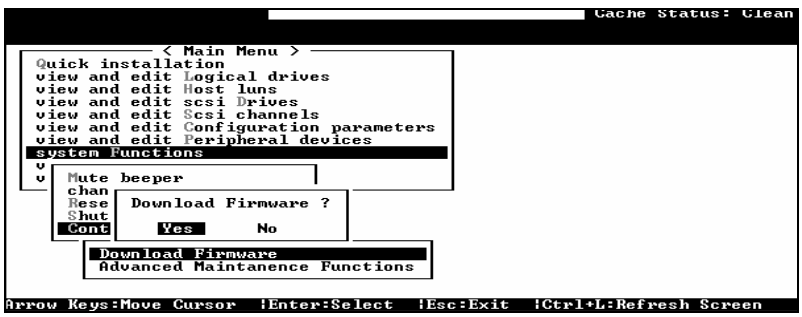

1. From the Main Menu, scroll down to "System Functions."

2. Go to "Controller Maintenance."

3. Choose "Download Firmware."

4. Set ZMODEM as the file transfer protocol of your terminal emulation software.

5. Send the Firmware Binary to the controller. In Hyper Terminal, select "Send file." If you are not using Hyper Terminal, choose "Upload" or "Send" (depending on the software).

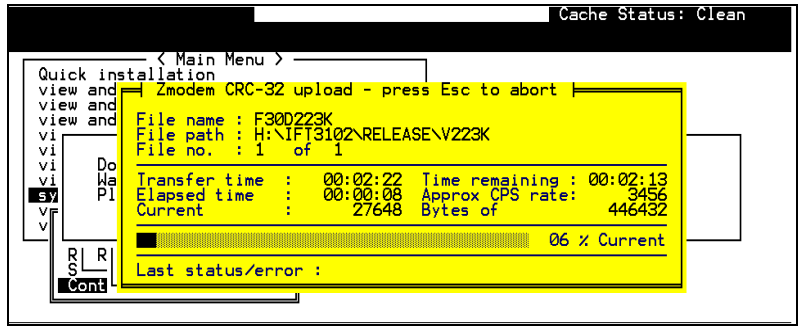

6. When the Firmware completes downloading, the controller will automatically reset itself. If the subsystem/controller is running firmware later than version 3.21, you need to manually reset the subsystem/controller for the new firmware to take effect.

This page is intentionally left blank.

Appendix

C

Event Messages

The controller events can be categorized as follows according to the severity levels:

Critical	Errors that need to attend to immediately
Warning	Errors
Notification	Command processed message sent from Firmware

The RAID subsystem records all system events from power on, temporarily record them in cache with up to 1,000 entries. To power off or to reset the controller will cause an automatic deletion of all the recorded event logs.

Firmware rev. 3.61 enables an event log to be saved in logical drives' 256MB reserved space and hence users can also see the events before system reset.

The RAIDWatch manager can be used to record events on multiple subsystems especially when controller reset or power-off is an expected action. The events can also be seen from RAIDWatch's Configuration Client utility. Associated details can be found in the RAIDWatch user's manual and online help.

Descriptions below may contain abbreviations. Abbreviations and Capitalized letters are preserved for the coherency with the event messages shown on the terminal session.

IMPORTANT!

Firmware revision 3.63 added event strings for I2C slave devices. The I2C device error is indicated by a bracketed message such as "**(XXX Status Failure Detected!)**" by the end of a peripheral device event. The cause of an error can be an I2C bus arbitration error, bus timeout, or a slave device content error.

Event Index

C.1 Logical Drive Events

C.1.1 Critical:

LG:0 Logical Drive ALERT: CHL:0 ID:132 Drive Missing
LG:0 Logical Drive ALERT: CHL:0 ID:132 Drive Failure
LG:0 Logical Drive ALERT: Rebuild Aborted
LG:0 Logical Drive ALERT: Rebuild Failed
LG:0 Logical Drive ALERT: Parity Regeneration Aborted
LG:0 Logical Drive ALERT: Parity Regeneration Failed
LG:0 Logical Drive ALERT: Creation Aborted
LG:0 Logical Drive ALERT: Creation Failed
LG:0 Logical Drive ALERT: Initialization Failed
LG:0 Logical Drive ALERT: Expansion Aborted
LG:0 Logical Drive ALERT: Expansion Failed
LG:0 ALERT: CHL:0 ID:132 Media Scan Failed
LG:0 ALERT: CHL:0 ID:132 Media Scan Aborted
LG:0 Logical Drive ALERT:CHL:0 ID:132 Clone Failed
LG:0 Logical Drive ALERT:CHL:0 ID:132 Clone Aborted
LG:0 Logical Drive ALERT: Logical Drive Bad Block Table FULL
LG:0 Logical Drive ALERT: Logical Drive Bad Block Table BAD
LG:0 Logical Drive ALERT: Logical Drive On-Line Init Table BAD
LG:0 Logical Drive ALERT: Bad Data Block Marked 000000084
LG:0 Logical Drive ALERT: UNPROTECTED Block Marked 000000084
LG:0 Logical Drive ALERT: Bad Data Block Encountered 000000084
LG:0 ALERT: Inconsistent Parity Encountered Block 000000084
LG:0 Logical Drive ALERT: Cache Data Purged
LG:0 Logical Drive ALERT: Fatal Fail
LG:0 Logical Drive ALERT: Invalid Array
LG:0 Logical Drive ALERT: Incomplete array
CHL:0 ID:0 ALERT: Media Error Unrecoverable-0x0

C.1.2 Notification:

LG:0 Logical Drive NOTICE: Continue Rebuild Operation
LG:0 Logical Drive NOTICE: Starting Rebuild
LG:0 Logical Drive NOTICE: Rebuild of Logical Drive Completed
LG:0 Logical Drive NOTICE: Continue Parity Regeneration Operation
LG:0 Logical Drive NOTICE: Starting Parity Regeneration
LG:0 Logical Drive NOTICE: Parity Regeneration of Logical Drive Completed
LG:0 Logical Drive NOTICE: Starting Creation
LG:0 Logical Drive NOTICE: Creation of Logical Drive Completed
LG:0 Logical Drive NOTICE: Starting On-Line Initialization
LG:0 Logical Drive NOTICE: On-Line Initialization of Logical Drive Completed
LG:0 Logical Drive NOTICE: Starting Off-Line Initialization

LG:0 Logical Drive NOTICE: Off-Line Initialization of Logical Drive Completed
LG:0 Logical Drive NOTICE: Starting On-Line Expansion
LG:0 Logical Drive NOTICE: On-Line Expansion of Logical Drive Completed
LG:0 Logical Drive NOTICE: Starting Off-Line Expansion
LG:0 Logical Drive NOTICE: Off-Line Expansion of Logical Drive Completed
LG:0 Logical Drive NOTICE: Continue Add Drive Operation
LG:0 Logical Drive NOTICE: Continue Migrate Operation
LG:0 Logical Drive NOTICE: Starting Add Drive Operation
LG:0 Logical Drive NOTICE: Starting Migrate Operation
LG:0 Logical Drive NOTICE: Add Drive Operation Paused
LG:0 Logical Drive NOTICE: Migrate Operation Paused
LG:0 Logical Drive NOTICE: Add Drive to Logical Drive Completed
LG:0 Logical Drive NOTICE: Migrate to Logical Drive Completed
NOTICE: CHL:0 ID:132 Starting Media Scan
LG:0 NOTICE: CHL:0 ID:132 Starting Media Scan
NOTICE: Media Scan of CHL:0 ID:132 Completed
LG:0 NOTICE: Media Scan of CHL:0 ID:132 Completed
LG:0 Logical Drive NOTICE:CHL:0 ID:132 Continue Clone Operation
LG:0 Logical Drive NOTICE:CHL:0 ID:132 Starting Clone
LG:0 Logical Drive NOTICE:CHL:0 ID:132 Copy and Replace Completed
LG:0 Logical Drive NOTICE:CHL:0 ID:132 Clone Completed
LG:0 Logical Drive NOTIFY: Bad Data Block Recovered 000000084

C.2 Channel and Individual Drive Events

C.2.1 Critical:

CHL:0 ID:132 ALERT: Media Error Unrecoverable-0xD8001C7C
CHL:0 FATAL ERROR(0)

C.2.2 Warning:

CHL:1 ID:0 Target WARN: Unexpected Select Timeout
CHL:1 Drive Channel WARN: Unexpected Select Timeout
CHL:1 RCC Channel WARN: Gross Phase/Signal Error Detected
CHL:1 ID:0 Target WARN: Gross Phase/Signal Error Detected
CHL:1 Drive Channel WARN: Gross Phase/Signal Error Detected
CHL:1 RCC Channel WARN: Unexpected Disconnect Encountered
CHL:1 ID:0 Target WARN: Unexpected Disconnect Encountered
CHL:1 Drive Channel WARN: Unexpected Disconnect Encountered
CHL:1 RCC Channel WARN: Timeout Waiting for I/O to Complete
CHL:1 ID:0 Target WARN: Timeout Waiting for I/O to Complete
CHL:1 Drive Channel WARN: Timeout Waiting for I/O to Complete
CHL:1 RCC Channel WARN: Parity/CRC Error Detected
CHL:1 ID:1 Host Channel WARN: Parity/CRC Error Detected
CHL:1 Host Channel WARN: Parity/CRC Error Detected
CHL:0 ID:0 Drive WARN: Unexpected Drive Not Ready (00B)

CHL:0 ID:0 Drive WARN: Drive HW Error (00B)
CHL:0 RCC Channel WARN: Unit Attention Received
CHL:0 ID:0 Target WARN: Unit Attention Received (10B)
CHL:0 Drive Channel WARN: Unit Attention Received (10B)
CHL:0 ID:0 Drive WARN: Aborted Command (00B)
CHL:0 ID:0 Drive WARN: Unexpected Sense Received (00B)
CHL:0 ID:0 Drive WARN: Block Reassignment Failed –0 (10B)
CHL:0 RCC Channel WARN: Data Overrun/Underrun Detected
CHL:0 ID:0 Target WARN: Data Overrun/Underrun Detected
CHL:0 Drive Channel WARN: Data Overrun/Underrun Detected
CHL:0 RCC Channel WARN: Negotiation Error Detected
CHL:0 ID:0 Target WARN: Negotiation Error Detected
CHL:0 Drive Channel WARN: Negotiation Error Detected
CHL:0 RCC Channel WARN: Invalid Status/Sense Data Received
CHL:0 ID:0 Target WARN: Invalid Status/Sense Data Received
CHL:0 Drive Channel WARN: Invalid Status/Sense Data Received
CHL:0 WARN: Redundant Loop Connection Error Detected on ID:11
CHL:0 Host Channel WARN: Channel Failure
WARN:SMART-Slot 1 Predictable Failure Detected-Clone Failed
WARN:SMART-Slot 1 Predictable Failure Detected
WARN:SMART-Slot 1 Predictable Failure Detected-Starting Clone
WARN:SMART-Slot 1 Predictable Failure Detected(TEST)
CHL:0 WARN: Fibre Channel Loop Failure Detected
CHL:0 WARN: Redundant Loop for CHL:0 Failure Detected
CHL:0 WARN:Redundant Path for CHL:0 ID:0 Expected but Not Found
CHL:0 ID:0 WARN:Redundant Path for Chl:0 ID:0 Failure Detected

C.2.3 Notification:

CHL:0 ID:132 NOTICE: Media Error Recovered-0xD8001C7C
CHL:0 NOTICE: Fibre Channel Loop Connection Restored
CHL:0 ID:255 NOTICE: Redundant Path for CHL:0 Restored
CHL:0 NOTICE: Redundant Path for CHL:0 ID:0 Restored
CHL:0 ID:0 Drive NOTICE: Drive Recovered Error Reported
CHL:0 LIP(FF B) Detected
CHL:0 Host Channel Notification: Bus Reset Issued
CHL:0 ID:0 Drive NOTICE: Scan Drive Successful
CHL:0 ID:0 Drive NOTIFY: Block Successfully Reassigned - 0 (10B)
CHL:_ ID_ Drive #SEV#: Media Error Encountered

C.3 General Target Events

C.3.1 Critical:

SAF-TE Device(0)ALERT: Power Supply Failure Detected(Idx:132)
Peripheral Set 0 Device ALERT: Power Supply 0 Failure Detected
Peripheral Set 0 Device ALERT: Power Supply 0 Not Present
Peripheral Set 0 Device ALERT: Low Voltage Detected (5V)
SES(C0 I0)Power Supply 0:Device Not Supported !

SAF-TE Device(0)ALERT: Cooling Fan Not Installed(Idx:2)
SAF-TE Device(0)ALERT: Cooling Fan Failure Detected(Idx:2)
Peripheral Set 0 Device ALERT: Cooling Fan2 Failure Detected
Peripheral Set 0 Device ALERT: Cooling Fan2 Not Present
Peripheral Set 0 Device ALERT: Low Speed Detected (0 RPM)
SES(C0 I0)Cooling Fan 0:Device Not Supported!
SAF-TE Device(0)ALERT: Elevated Temperature Alert
Peripheral Device ALERT: CPU Cold Temperature Detected (0.0(C))
Peripheral Device ALERT: Elevated Temperature Alert
Peripheral Set 0 Device ALERT: Temperature Sensor 0 Failure Detected
Peripheral Set 0 Device ALERT: Temperature Sensor 0 Not Present
Peripheral Set 0 Device ALERT: Cold Detected (0(C))
SES(C0 I0)Temp Sensor 0:Device Not Supported!
UPS Connection Is Absent
SAF-TE Device(2) ALERT: UPS Power Failure Detected
Peripheral Set 0 Device ALERT: UPS 0 AC Power Failure Detected
Peripheral Set 0 Device ALERT: UPS 0 Battery Failure Detected
Warning: UPS AC Power-Loss Detected
Warning: UPS Battery Low - 0%
SES(C0 I0)UPS 2:Device Not Supported!
ALERT: Enclosure ID Conflict/Invalid

C.3.2 Notification:

SAF-TE Device(2) NOTICE: Fan Back On-Line(Idx:11)
NOTICE: FAN Module A Back On-Line(FAN0 0RPM)
NOTICE: Controller FAN 1 Back On-Line(0 RPM)
Peripheral Device NOTICE: Fan Back On-Line
Peripheral Set 1 Device NOTICE: FAN 1 Back On-Line
Peripheral Set 1 Device NOTICE: FAN 1 is Present
Peripheral Set 1 Device NOTICE: FAN 1 Back On-Line(0 RPM)
SES(C0 I0)Cooling Fan 1:Fan Back On-Line
SAF-TE Device(1) NOTICE: Temperature Back To Non-Critical LVs
Peripheral Set 1 Device NOTICE: Temperature 1 Back To Non-Critical LVs
Peripheral Set 1 Device NOTICE: Temperature 1 is Present
Peripheral Set 1 Device NOTICE: Temperature 1 Back To Non-Critical LVs(0(C))
SES(C0 I0)Temp Sensor 1:Temperature Back To Non-Critical LVs
SAF-TE Device(1) NOTICE: Power Supply Back On-Line(Idx:4)
Power Supply 0 Back-Online
Peripheral Set 0 Device NOTICE: PSU0 +12V Back On-Line
Peripheral Set 0 Device NOTICE: PSU0 +12V is Present
Peripheral Set 0 Device NOTICE: PSU0 +12V Back On-Line(5.0V)
SES(C0 I0)Power Supply 2:Power Supply Back On-Line
UPS Connection Detected
UPS AC Power Restored
UPS Battery Restored To Safe Level 0
SAF-TE Device(2) NOTICE: UPS Power Back On-Line
Peripheral Set 2 Device NOTICE: UPS 2 AC Power Back On-Line
Peripheral Set 2 Device NOTICE: UPS 2 Battery Back On-Line
SES(C0 I0)UPS 2:UPS Power Back On-Line

C.4 Controller Events

C.4.1 Critical:

ALERT: +3.3V Low Voltage Detected(205.9V)
ALERT: Controller FAN 0 Low Speed Detected(0 RPM)
Controller ALERT: Redundant Controller Failure Detected
Controller SDRAM ECC Multi-bits Error Detected
Controller SDRAM ECC Single-bit Error Detected
Controller SDRAM Parity Error Detected
Controller PCI Bus Parity Error Detected
BBU Absent or Failed! Correct It and Reset Ctlr to Take Effect
Controller BBU Absent or Failed!
Controller BBU Failure Detected!
Controller BBU Thermal Shutdown/Enter Sleep-Mode!

C.4.2 Warning:

Memory Not Sufficient to Fully Support Current Config.
ALERT: Cache volatilized
Force Controller Write-Through on Trigger Cause

C.4.3 Notification:

Board1 Cold Temperature Back To Non-Critical LVs(0(C))
+12V Upper Voltage Back within Acceptable Limits (0.0V)
+12V Lower Voltage Back within Acceptable Limits (0.0V)
Memory is Now Sufficient to Fully Support Current Config.
Controller NOTICE: NVRAM Factory Defaults Restored
Controller BBU is Charging!
NOTICE: Controller BBU Present!
NOTICE: Controller BBU Back On-Line!
NOTICE: Controller BBU Fully Charged!
Notify: Controller BBU Lifetime Exceeded - Replacement Recommended
Controller NOTICE: Redundant Controller Firmware Updated
NOTICE: Force Controller Shutdown Start
NOTICE: Force Controller Shutdown Complete

C.5 Data Service Events - Snapshot

C.5.1 Critical:

Snapshot ALERT - SC ID:3BB6514353CF3B85 Space Thresh Exceeded 80%

C.5.2 Warning:

Snapshot WARNING - SC ID:3BB6514353CF3B85 Creation Failed
Snapshot WARNING - SS ID:30A55D31700F2AAB Creation Failed

Snapshot WARNING - SI ID:1D3E09F11D48E27A Creation Failed
Snapshot WARNING - SC ID:3BB6514353CF3B85 Extend Failed
Snapshot WARNING - SC ID:3BB6514353CF3B85 Space Thresh Exceeded 60%

C.5.3 Notification:

Snapshot NOTICE - SC ID:3BB6514353CF3B85 Creation Complete
Snapshot NOTICE - SC ID:3BB6514353CF3B85 Extend Complete
Snapshot NOTICE - SC ID:3BB6514353CF3B85 Deleted
Snapshot NOTICE - SS ID:30A55D31700F2AAB Creation Complete
Snapshot NOTICE - SS ID:30A55D31700F2AAB Deleted
Snapshot NOTICE - SI ID:1D3E09F11D48E27A Creation Complete
Snapshot NOTICE - SI ID:1D3E09F11D48E27A Deleted
Snapshot NOTICE - SI ID:1D3E09F11D48E27A Activated
Snapshot NOTICE - SC ID:3BB6514353CF3B85 Space Thresh Exceeded 40%
Snapshot NOTICE – SC ID:3BB6514353CF3B85 Space Thresh Restored 80%
Snapshot NOTICE – SC ID:3BB6514353CF3B85 Purge Triggered- 80%
Snapshot NOTICE – SS ID:30A55D31700F2AAB Prune Triggered – 80%";
Snapshot NOTICE Non-optimal configuration may impact performance

C.1 Logical Drive Events

C.1.1 Critical:

Message	LG:0 Logical Drive ALERT: CHL:0 ID:132 Drive Missing
What Happened?	The LD member drive could have been accidentally removed or drive connection problems occurred before system boot-up.
What to Do?	The drive could have been accidentally removed or drive connection problems occurred in the boot time. Check drive connection and related enclosure status.

Message	LG: Logical Drive ALERT: CHL:0 ID:132 Drive Failure
What Happened?	The specified hard drive in the specified logical drive has failed
What to Do?	When a system is running, removing any LD member disk or member disk failure will be recognized as disk failure event. For example, RAID controller cannot access a member disk and decides this drive has failed. If a spare is available, the subsystem will proceed with an automatic rebuild. If there is no spare, replace the faulty drive using a drive that is known to be good and rebuild will be automatically initiated provided that a replacement drive can be recognized by the subsystem. For example, in a SCSI-based subsystem, a manual "Scan Drive" command should be executed. In subsystems using other types of disk drives, drive swap auto-detection should be enabled on the drive channels. If all disk drives are lost inside an enclosure, please check the related power supply status or backplane power connection.

Message	LG:0 Logical Drive ALERT: Rebuild Aborted/ Failed
What Happened?	Logical drive rebuild aborted. It could result from the one of the following reasons: 1. The rebuild has been manually canceled by a user. 2. The replacement drive used for a rebuild may have failed during the rebuild or the drive contains immanent defects. 3. System faults (for example, an LD fatally failed) occurred. Another member disk failed in a rebuild process.
What to Do?	Carefully identify and replace the faulty drive and perform logical drive rebuild again. It is best to stop host I/Os temporarily to this logical drive.

Message	LG:0 Logical Drive ALERT: Parity Regeneration Aborted/Failed
What Happened?	It could result from one of the following reasons: 1. The regenerate operation has been manually canceled by a user. 2. System faults or integrity problems forced system to abandon the operation. For example, any single member disk failed in a LD will force this operation to stop. 3. Irrecoverable parity inconsistency may also cause a halt to

	the operation by user's configuration/operation errors.
What to Do?	If the Check/Regenerate Parity Failure is caused by a drive failure, replace the faulty drive and rebuild the LD. Then execute the check parity operation again. If the Check/Regenerate Parity Failure results from the discovery of inconsistent parity, users have to change the firmware check parameters if users wish to complete the whole check parity operation.

Message	LG:0 Logical Drive ALERT: Creation Aborted
What Happened?	Logical drive creation process manually aborted or that some system integrity problems forced the subsystem to abort logical drive creation process.
What to Do?	Check proper system working conditions.

Message	LG:0 Logical Drive ALERT: Creation Failed
What Happened?	Logical drive creation process failed when one or more drive members failed or have been accidentally removed. Drive abnormality may also be the cause of the problems.
What to Do?	Check proper system working conditions and drive connection. Replace the faulty drives and repeat the creation process.

Message	LG:0 Logical Drive ALERT: Initialization Failed
What Happened?	Logical drive initialization failed. It could result from one of the following reasons: 1. Off-Line Initialization Failed Logical drive initialization failed. This event can result from one member drive failure during initialization for all types of RAID configuration (R0/1/3/5/6). 2. On-Line Initialization Failed This event can result from fatal failure of a logical drive, for example, two disk drives failed during the initialization process for a RAID 5 LD and three disk drives failed for a RAID 6 LD.
What to Do?	For off-line initialization, carefully identify and replace the faulty drive and perform create logical drive again. For on-line initialization, reboot system and check the LD 1. If LD can be restored to the degraded mode, then the initialization process will continue. 2. If the LD cannot be restored, then you have to replace the failed disks, recreate the LD, and restore data from a backup source.

Message	LG:0 Logical Drive ALERT: Expansion Aborted/Failed
What Happened?	Logical drive expansion failed. It could result from one of the following reasons: 1. The expansion has been canceled by a user. 2. For On-line expansion, an expansion operation is terminated when the target LD fatally fails. 3. For Off-line expansion, an expansion operation fails when any of its member drives fails.
What to Do?	Carefully identify and replace the faulty drive and perform logical

	drive rebuild first. For an online expansion, the expansion operation should continue with a single drive failure. When the LD rebuild process is completed, the LD will resume the expansion process. For an off-line expansion, the expansion will be discontinued by a single disk failure. The expansion operation will not continue after the rebuild is completed. Users have to start the expand process again.

Message	LG:0 ALERT: CHL:0 ID:132 Media Scan Failed
What Happened?	Media Scan failed to carry on with scanning drive(s) for certain reasons, e.g., a logical drive or a disk drive has fatally or physically failed.
What to Do?	Media Scan on a failed disk drive will stop, but the scan on healthy members of a target LD will continue. For some OEM version, the choice between proceeding or ending a Media Scan operation depends on the integrity of an LD. If the LD status is not good, e.g., a member drive fails and enters a degraded mode, then the Media Scan operation will pause for healthy members of the target LD. When the LD restores its good status, e.g., a faulty drive replaced and the LD is successfully rebuilt, then firmware will resume the Media Scan operation. Media Scan on the replacement drive will start from the beginning.

Message	LG:0 ALERT: CHL:0 ID:132 Media Scan Aborted
What Happened?	Media Scan operation is aborted for certain reasons, e.g., it has been manually canceled.
What to Do?	Check disk drive and system operation status.

Message	LG:0 Logical Drive ALERT: CHL:0 ID:132 Clone Failed/Aborted
What Happened?	Drive cloning failed. It could result from one of the following reasons: 1. Cloning has been canceled by a user. 2. The LD contains the source drive is fatal failed during the cloning process. 3. The source drive used for cloning has failed during the cloning process. 4. The target drive used for cloning has failed during the cloning process.
What to Do?	If the Logical Disk has fatally failed, then the cloning operation stops. If the source drive fails, then system firmware will proceed with the following: 1. If a spare drive is available, then LD will stop cloning and begins a rebuild using the available spare drive. 2. If no spare drive is available, then LD will stop cloning and start rebuild using the target drive (the target drive was originally a spare drive). If the target drive fails, then cloning will stop. Carefully identify

	and replace the faulty drive. Then restart the cloning operation.

Message	LG:0 Logical Drive ALERT: Logical Drive Bad Block Table FULL
What Happened?	There are too many damaged sectors and firmware used up the available bad block table entries. Bad blocks have been discovered and marked by "write" operations including host write commands, disk clone, Add Drive, etc.
What to Do?	The target disk drive may have become physically unreliable. The data in the target LD could have been corrupted, please restore the data from backup. Contact system vendor for help.

Message	LG:0 Logical Drive ALERT: Logical Drive Bad Block Table BAD
What Happened?	System failed to generate a bad block table. Logical drive may have generated serious integrity problems.
What to Do?	The data in the target LD could have been corrupted, please execute the Regenerate Parity function and restore data if instances of inconsistent parity are found. Contact system vendor for help.

Message	LG:0 Logical Drive ALERT: Logical Drive On-Line Init Table Bad
What Happened?	When an LD is still under an online initialization process and the system reboots for some reasons, the system will not be able to generate a correct On-line Init table at boot time. The logical drive may have encountered serious integrity problems.
What to Do?	Execute regenerate parity operation and ignore the inconsistent parity event if occurred. Contact system vendor for help.

Message	LG:0 Logical Drive ALERT: Bad Data Block Marked 000000084
What Happened?	During write operations (including rebuild, host write commands, disk cloning, and Add Drive), media errors may cause data inconsistency. System cannot regenerate the data by RAID parity calculation. The affected data blocks # are marked bad. The situation may result from the following: 1. Rebuild and host write command: For example, in a RAID 5 configuration, one member drive failed, and media errors are found on another member drive. 2. Disk clone and Add Drive: Concurrent occurrences of media errors are found on the same stripe across two member drives.
What to Do?	The data affected by data blocks failure will be lost. The host computer will return media error messages if host accesses fall within the affected blocks. Restoring data from a backup source is recommended.

Message	LG:0 Logical Drive ALERT: UNPROTECTED Block Marked 000000084
What Happened?	Media errors are found in an unprotected array, e.g., RAID0. The affected data blocks will be marked bad.
What to Do?	The data affected by data blocks failure will be lost. The host computer will return media error messages if host access falls within the affected blocks. Restoring data from a backup source is

	recommended.

Message	LG:0 Logical Drive ALERT: Bad Data Block Encountered 000000084
What Happened?	Data blocks previously marked Bad are re-encountered during host I/O access or certain operations such as Media Scan. The marked logical blocks will cause the firmware to return media error status to host.
What to Do?	The data affected by data blocks failure is lost. Restoring data from a backup source is recommended. The host computer will return media error messages if host access falls within the affected blocks.

Message	LG:0 ALERT: Inconsistent Parity Encountered Block 000000084
What Happened?	Users execute regenerate parity operation and a parity inconsistency is encountered. Problematic block address is shown at the end of the message.
What to Do?	Parity inconsistency means the data is corrupted or invalid across some data blocks. We recommend users to restore specific data from their routine backup.

Message	LG:0 Logical Drive ALERT: Cache Data Purged
What Happened?	There can be two causes for a RAID system to purge its cached data that belongs to a specific LD: 1. A fatal failure occurred on a specific LD, the related cached data will be purged. 2. Cached data is missing when system powers up and finds the battery is invalid to support the correct data in cache. System will check the cached data and issue an alert event by each LD if the cache is corrupt.
What to Do?	If this event occurs, it means the cached data for this LD could have been corrupted or invalid. The LD will be put into an off-line state. Users have to manually restart this LD and check the integrity of data before proceeding with any update operation. It is recommended to restore data from your routine backup.

Message	LG:0 Logical Drive ALERT: Fatal Fail/ Invalid Array/ Incomplete array
What Happened?	Fatal failure or incomplete array means that the LD has lost the protection by RAID configuration. If system cannot find some member disks for a specific LD at boot time, the LD will be considered as incomplete. If some member disks of a specific LD fail during operation, the LD will be considered as fatally failed. The possible causes include: 1. For RAID 5, two member disks have failed or are missing. 2. For RAID 6, three member disks have failed or are missing. 3. For RAID 1, two member disks have failed or are missing

	in the same mirror pair. 4. One disk failed or missing in a RAID 0/NRAID configuration.
What to Do?	The LD is not available at boot time or running time. If an LD has fatally failed, you may try to restart the RAID system and check if the system detects the failed disks. If one or two member disks are back then LD will be rebuilt. If the LD cannot be restored, then you have to replace the failed disks, recreate the LD, and restore data from a backup source.

Message	CHL:0 ID:0 ALERT: Media Error Unrecoverable-0x0
What Happened?	Media errors are found during read operations, and the errors can or cannot be recovered. The causes include: 1. Media errors are found in an unprotected array, e.g., a RAID0 and NRAID configuration. 2. Media errors are found in a degraded array, e.g., a RAID 5 array with a failed disk. Drive failure and the media errors on other members failed the recovery attempt. 3. Concurrent occurrences of media errors found on the same stripe across two member drives. Host receives a media error response for the specific logical block address related to the physical bad block.
What to Do?	The data affected by data blocks failure will be lost. The host computer will return media error messages if host access falls on the affected data blocks. Restoring data from a backup source is recommended.

C.1.2 Notification:

Message	LG:0 Logical Drive NOTICE: Continue Rebuild Operation
What Happened?	The target logical drive has been stored to its previous one-drive-failed status, and the rebuild operation is automatically resumed by firmware. This may occur when the system is powered off or reset during the rebuild process or that media errors have been discovered and corrected during the rebuild.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: Starting Rebuild
What Happened?	The rebuild process has begun.
What to Do?	This is the message displayed when a stand-by spare is available when a member drive fails or when a faulty drive is physically replaced. The subsystem should automatically detect a drive for rebuild if the hot-swap detection mechanism has been enabled.

Message	LG:0 Logical Drive NOTICE: Rebuild of Logical Drive Completed
What Happened?	The subsystem has successfully rebuilt a logical drive.

What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: Continue Parity Regeneration Operation
What Happened?	The parity regeneration process continues to execute.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: Starting Parity Regeneration
What Happened?	Start regenerating parity of a logical drive.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: Parity Regeneration of Logical Drive Completed
What Happened?	The parity regeneration process on logical drive _ is completed.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: Starting Creation
What Happened?	Logical drive creation process started. Note that the creation and initialization processes are separated. Creation of a logical drive only takes a second while the initialization may take hours with the arrays containing large capacity drives.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: Starting On-Line Initialization
What Happened?	Creation process is completed, initialization process has begun to generate parity sectors and readying member drives. The On-line initialization mode takes a longer time to complete and allows the logical drive to receive host I/Os immediately if appropriate host ID/LUN mapping has been applied.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: On-Line Initialization of Logical Drive Completed
What Happened?	Logical drive on-line initialization completed.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: Starting Off-Line Initialization
What Happened?	Creation process is completed, initialization process has begun to generate parity sectors and readying member drives. The Off-line initialization mode takes place immediately after the creation of a logical drive, and does not allow host I/Os to be distributed to the logical drive at the same time.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: Off-Line Initialization of Logical Drive Completed
What Happened?	Logical drive off-line initialization completed. The logical drive is now ready for host ID/LUN mapping, and then receiving host I/Os.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: Starting On-Line Expansion
What Happened?	Start expanding the logical drive's added or free capacity (previously not included) when system find appropriate time after selecting to expand its capacity. The On-Line mode allows the added capacity of a logical drive to appear immediately as a usable partition and ready for host I/Os. This partition may require being associated with a separate ID/LUN on the host bus. Response to host I/Os will be slower because the expansion process requires distributing data blocks and parity data across new drive sectors.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: On-Line Expansion of Logical Drive Completed
What Happened?	Logical drive on-line expansion completed. The added capacity will appear as a new partition, and this partition may require being associated with a separate ID/LUN on the host bus. Adjustments should also be made through the host OS volume management utility.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: Starting Off-Line Expansion
What Happened?	Start expanding the logical drive's added or free capacity (previously not included) when system find appropriate time after selecting to expand its capacity. The Off-Line mode does not allow the added capacity of a logical drive to be immediately available for host I/Os. This partition can only be associated with a separate ID/LUN on the host bus after the expansion process is completed. Access to the data already stored on the logical drive is not affected.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: Off-Line Expansion of Logical Drive Completed
What Happened?	Logical drive off-line expansion completed. The added capacity will appear as a new partition, and this partition may require being associated with a separate ID/LUN on the host bus.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: Continue Add Drive Operation
What Happened?	The target logical drive has been restored to its previous status, e.g., power-off or media errors, and the add drive operation can

	continue.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: Continue Migrate Operation
What Happened?	The target logical drive has been restored to its previous status, e.g., power-off or media errors, and the migration operation can continue.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: Starting Add Drive Operation
What Happened?	Add drive expansion process started
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: Starting Migrate Operation
What Happened?	Logical Drive Migrate process started
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: Add Drive Operation Paused
What Happened?	The add drive expansion process is halted by: 1. Logical drive expansion is canceled by a user. 2. One of the member drives failed during logical drive expansion. 3. Media errors are found on one or more of the logical drive members. 4. System faults have occurred.
What to Do?	Carefully identify and replace the faulty drive and correct system faults. The add drive expansion should continue once the faults are corrected.

Message	LG:0 Logical Drive NOTICE: Migrate Operation Paused
What Happened?	The migration process is halted by: 1. Logical drive migration is canceled by a user. 2. One of the member drives failed during logical drive migration. 3. Media errors are found on one or more of the logical drive members. 4. System faults have occurred.
What to Do?	Carefully identify and replace the faulty drive and correct system faults. The migration should continue once the faults are corrected.

Message	LG:0 Logical Drive NOTICE: Add Drive to Logical Drive Completed
What Happened?	The add drive expansion process is completed.
What to Do?	Press ESC to clear the message. Please remember to rearrange the added capacity which will appear as a new partition volume.

Message	LG:0 Logical Drive NOTICE: Migrate to Logical Drive Completed
What Happened?	The logical drive migration process is completed.
What to Do?	Press ESC to clear the message.

Message	NOTICE: CHL:0 ID:132 Starting Media Scan
What Happened?	Media Scan is manually or automatically started by the preset Task Schedules on a specific disk drive.
What to Do?	Press ESC to clear the message.

Message	LG:0 NOTICE: CHL:0 ID:132 Starting Media Scan
What Happened?	Media Scan is manually or automatically started by the preset Task Schedules on a specific member of Logical Drive 0.
What to Do?	Press ESC to clear the message.

Message	NOTICE: Media Scan of CHL:0 ID:132 Completed
What Happened?	Media Scan on a specific disk drive is completed.
What to Do?	Press ESC to clear the message.

Message	LG:0 NOTICE: Media Scan of CHL:0 ID:132 Completed
What Happened?	Media Scan on a specific member of Logical Drive 0 is completed.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: CHL:0 ID:132 Continue Clone Operation
What Happened?	The target logical drive has been restored to its previous status, e.g., a failed drive is restored, and the cloning process can continue.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: CHL:0 ID:132 Starting Clone
What Happened?	Cloning process started on one of the members of Logical Drive 0
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: CHL:0 ID:132 Copy and Replace Completed
What Happened?	Copy and replace is completed with a specific member of Logical Drive 0. The original member drive is replaced by another disk drive.
What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTICE: CHL:0 ID:132 Clone Completed
What Happened?	Cloning is completed with a specific member of Logical Drive 0, whether it is replaced (Copy and Replace) or cloned (Perpetual Clone) to a spare drive.

What to Do?	Press ESC to clear the message.

Message	LG:0 Logical Drive NOTIFY: Bad Data Block Recovered 000000084
What Happened?	Host writes fall onto the data blocks previously marked bad and overwrite the drive sectors.
What to Do?	Press ESC to clear the message.

C.2 Channel and Individual Drive Events

C.2.1 Critical:

Message	CHL:0 ID:132 ALERT: Media Error Unrecoverable-0xD8001C7C
What Happened?	Drive surface media errors found and after rewrites the drive sectors are determined as physically damaged and unusable.
What to Do?	Data will be reassigned to other drive sectors. Host writes should continue without problems.

Message	CHL:0 Fatal Error (0)
What Happened?	Channel 0 has fatally failed.
What to Do?	Contact your system vendor for help.

C.2.2 Warning:

Message	CHL:_ ID:_ Target WARN: Unexpected Select Timeout
What Happened?	Drive target select timeout. The specified hard drive cannot be selected by the RAID controller/subsystem. Whether the drive has been removed, or the cabling/termination/canister is out of order.
What to Do?	This could occur if a disk drive is not properly installed or accidentally removed, or when cable links/backplane traces/connection become unreliable. Check cabling and drive installation, and contact your RAID system supplier.

Message	CHL:_ Drive Channel WARN: Unexpected Select Timeout
What Happened?	Drive Channel target select timeout. The specified drive channel cannot be selected by the RAID controller/subsystem. Faults might have occurred with the cabling/termination/internal data paths.
What to Do?	Check cabling and drive installation, and contact your RAID system supplier.

Message	CHL:_ RCC Channel WARN: Gross Phase/Signal Error Detected
What Happened?	Phase/signal abnormality detected on the specific RCC channel.
What to Do?	Rare occurrence of phase/signal error could be recovered by firmware retry. Contact your RAID system supplier.

Message	CHL:_ ID:0 Target WARN: Gross Phase/Signal Error Detected
What Happened?	Phase/signal abnormality detected with a specific target disk drive.
What to Do?	Check proper installation of disk drives and contact your RAID system supplier.

Message	CHL:_ Drive Channel WARN: Gross Phase/Signal Error Detected
What Happened?	Phase/signal abnormality detected on the specific drive channel.
What to Do?	Contact your RAID system supplier.

Message	CHL:_ RCC Channel WARN: Unexpected Disconnect Encountered
What Happened?	Unexpected disconnect detected on the specific RCC channel.
What to Do?	Check cabling/termination and canister connections, and contact your RAID system supplier.

Message	CHL:_ ID:_ Target WARN: Unexpected Disconnect Encountered
What Happened?	Unexpected disconnect detected on the specific target disk drive.
What to Do?	Check cabling/termination and canister installation, and contact your RAID system supplier.

Message	CHL:_ Drive Channel WARN: Unexpected Disconnect Encountered
What Happened?	Unexpected disconnect detected on the specific drive channel.
What to Do?	Check cabling/termination and canister connections, and contact your RAID system supplier.

Message	CHL:_ RCC Channel WARN: Timeout Waiting for I/O to Complete
What Happened?	I/O timeout on specific RCC channel.
What to Do?	Contact your RAID system supplier.

Message	CHL:_ ID:_ Target WARN: Timeout Waiting for I/O to Complete
What Happened?	Drive-side target I/O timeout. Possible drive-side cabling/termination and canister connection abnormal or drive malfunctioning.
What to Do?	Check drive-side cabling/termination/canister/disk drive connections and contact your RAID system supplier.

Message	CHL:_ Drive Channel WARN: Timeout Waiting for I/O to Complete
What Happened?	I/O timeout on specific drive channel.
What to Do?	Contact your RAID system supplier.

Message	CHL:_ RCC Channel WARN: Parity/CRC Error Detected
What Happened?	RCC channel parity or CRC errors detected.

What to Do?	Rare occurrence of these errors can be managed by firmware. Contact your RAID system supplier.

Message	CHL:_ ID:_ Host Channel WARN: Parity/CRC Error Detected
What Happened?	Parity or CRC errors detected through host ID, CHL_ ID_.
What to Do?	Occurrences of these faults usually can be negotiated between RAID subsystems and application servers. If occurrences become frequent, contact your RAID system supplier.

Message	CHL:_ Host Channel WARN: Parity/CRC Error Detected
What Happened?	Host channel parity or CRC errors detected.
What to Do?	Parity and CRC faults can often be negotiated between application servers and RAID subsystems. Contact your RAID system supplier if this occurs too frequently.

Message	CHL:_ ID:_ Drive WARN: Unexpected Drive Not Ready (00B)
What Happened?	Unexpected drive not ready returned from a disk drive on CHL:_ ID:_.
What to Do?	Check proper installation of disk drives or cabling connection. This event often occurs when hot-swapping disk drives and may be caused by signal glitches. If the disk drive fails to respond, the subsystem should consider the disk drive as a failed drive when the timeout threshold is reached. Contact your RAID system supplier.

Message	CHL:_ ID:_ Drive WARN: Drive HW Error (00B)
What Happened?	Drive hardware error returned from a disk drive on CHL:_ ID:_.
What to Do?	The target disk drive may have failed. Contact your RAID system supplier for a replacement.

Message	CHL:_ RCC Channel WARN: Unit Attention Received
What Happened?	Unit attention received from RCC channel CHL:_.
What to Do?	Rare occurrences of these errors can be managed by firmware. Contact your RAID system supplier.

Message	CHL:_ ID:_ Target WARN: Unit Attention Received (10B)
What Happened?	Drive-side target unit attention received on a disk drive CHL:_ ID:_.
What to Do?	Rare occurrence of these errors can be managed by firmware. Check disk drive connection and contact your RAID system supplier.

Message	CHL:_ Drive Channel WARN: Unit Attention Received (10B)
What Happened?	Drive-side target unit attention received on a drive channel CHL:_ ID:_.
What to Do?	Rare occurrence of these errors can be managed by firmware. Check proper installation and contact your RAID system supplier.

Message	CHL:_ ID:_ Drive WARN: Aborted Command (00B)
What Happened?	Aborted command reported from a disk drive CHL:_ ID:_
What to Do?	Rare occurrence of these errors can be managed by firmware. Contact your RAID system supplier.

Message	CHL:_ ID:_ Drive WARN: Unexpected Sense Received (00B)
What Happened?	Drive-side target unexpected sense received on a disk drive CHL:_ ID:_.
What to Do?	Check drive-side cabling/installation/working conditions. This may result from bad signal quality or poor connection. Contact your RAID system supplier.

Message	CHL:_ ID:_ Drive WARN: Block Reassignment Failed – 0(10B)
What Happened?	Data regeneration and reassignment failed. The Bad Block Table is full. Too many bad blocks have been found on the disk drive. The adjacent data blocks on other member drives might have also been damaged.
What to Do?	Rare occurrence of these errors can be managed by firmware. Contact your RAID system supplier.

Message	CHL:_ RCC Channel WARN: Data Overrun/Underrun Detected
What Happened?	Data overrun/underrun errors detected on the RCC channel CHL:_.
What to Do?	Rare occurrence of these errors can be managed by firmware. Contact your RAID system supplier.

Message	CHL:_ ID:_ Target WARN: Data Overrun/Underrun Detected
What Happened?	Data overrun/underrun errors detected on a disk drive CHL:_ ID:_.
What to Do?	Rare occurrence of these errors can be managed by firmware. Contact your RAID system supplier.

Message	CHL:_ Drive Channel WARN: Data Overrun/Underrun Detected
What Happened?	Data overrun/underrun errors detected on the drive channel CHL:_.
What to Do?	Rare occurrence of these errors can be managed by firmware. Contact your RAID system supplier.

Message	CHL:_ RCC Channel WARN: Negotiation Error Detected
What Happened?	Negotiation errors occurred on the RCC channel CHL:_.
What to Do?	Rare occurrence of these errors can be managed by firmware. Contact your RAID system supplier.

Message	CHL:_ ID:_ Target WARN: Negotiation Error Detected
What Happened?	Negotiation errors occurred with the communications with a disk drive CHL:_ ID:_. The event could also occur with drive target data bus sync/wide negotiation abnormality.
What to Do?	Rare occurrence of these errors can be managed by firmware. Contact your RAID system supplier.

Message	CHL:_ Drive Channel WARN: Negotiation Error Detected
What Happened?	Negotiation errors occurred with the communications over a drive channel CHL:_. The event could also occur with drive target data bus sync/wide negotiation abnormality.
What to Do?	Rare occurrence of these errors can be managed by firmware. Contact your RAID system supplier.

Message	CHL:_ RCC Channel WARN: Invalid Status/Sense Data Received
What Happened?	Invalid status/sense data received on the RCC channel CHL:_
What to Do?	Rare occurrence of these errors can be managed by firmware. Contact your RAID system supplier.

Message	CHL:_ ID:_ Target WARN: Invalid Status/Sense Data Received (10B)
What Happened?	Invalid status/sense data received by a disk drive CHL:_ ID:_.
What to Do?	Rare occurrence of these errors can be managed by firmware. Contact your RAID system supplier.

Message	CHL:_ Drive Channel WARN: Invalid Status/Sense Data Received
What Happened?	Invalid status/sense data received on the drive channel CHL:_
What to Do?	Rare occurrence of these errors can be managed by firmware. Contact your RAID system supplier.

Message	CHL:_ WARN: Redundant Loop Connection Error Detected on ID:_
What Happened?	One of the dual loop members might have failed or been disconnected. Make sure all cabling are properly connected and topological configurations properly set.
What to Do?	Check cabling and channel bus working conditions. If the problem does not result from cabling problems, contact your RAID system supplier.

Message	CHL:_ Host Channel WARN: Channel Failure
What Happened?	Host channel CHL:_ link failure or disconnection occurred.
What to Do?	Check proper cabling connection and host HBA, link connection devices, etc. Contact your RAID system supplier.

Message	WARN:SMART-Slot _ Predictable Failure Detected-Clone Failed
What Happened?	Hard drive(s) report SMART-detected defects. A spare drive is

	commenced to clone or replace the member drive showing defects. The cloning operation failed. The member drive suspected of errors might have failed, or the clone process has been interrupted, e.g., yet another member has failed or the spare drive used for cloning has shown immanent faults.
What to Do?	Carefully check drive status and replace the unstable/failed drive. Contact your RAID system supplier.

Message	WARN:SMART-Slot _ Predictable Failure Detected
What Happened?	Hard drive slot:_ reported SMART-detected errors, e.g., abnormal rotation speed, seek time delay, etc.
What to Do?	Carefully check drive status and replace the unstable/failed drive. Contact your RAID system supplier.

Message	WARN:SMART-Slot _ Predictable Failure Detected-Starting Clone
What Happened?	One or more hard drive(s) reports SMART-detected defects. Starts using a pre-defined spare drive to clone the member drive suspected of inherent errors. This happens when SMART detection is enabled with the "Copy & Replace" or "Perpetual Clone" reaction schemes.
What to Do?	Press ESC to clear the message. After the cloning process, you may consider replacing the defective drive with a good one and configure it as a spare drive in case drive failure might happen in the future.

Message	WARN:SMART-Slot _ Predictable Failure Detected (TEST)
What Happened?	This results from a manually executed test command. This ensures your disk drive supports SMART detection and report functionality.
What to Do?	Press ESC to clear the message.

Message	CHL:_ WARN: Fibre Channel Loop Failure Detected
What Happened?	Fibre Channel loop integrity warning. Lack of bypass or nodes on a loop caused a Fibre loop to fail.
What to Do?	Check host- or drive-side connection. Contact your RAID system supplier.

Message	CHL:_ WARN: Redundant Loop for CHL:_ Failure Detected
What Happened?	One of the dual loop members may have failed or been disconnected. Make sure all channels are properly connected and topological configurations properly set.
What to Do?	Check cabling and system installation. Contact your supplier if no connection errors could be found.

Message	CHL:_ WARN: Redundant Path for CHL:_ ID:_ Expected but Not Found
What Happened?	One of the dual loop members may have failed or been disconnected. This event is regenerated in the event when one of the dual loop connections to an expansion enclosure can not be found after power-up. Make sure all channels are properly connected and topological configurations properly set.

What to Do?	Check cabling and system installation. A manual "Issue LIP" command may restore the connection. Contact your supplier if no connection errors could be found.

Message	CHL:_ ID:_ WARN: Redundant Path for CHL:_ ID:_ Failure Detected
What Happened?	One of the dual loop members connecting device CHL:_ ID:_ may have failed or been disconnected. Make sure all channels are properly connected and topological configurations properly set.
What to Do?	Check cabling and system installation. Contact your supplier if no connection errors could be found.

C.2.3 Notification:

Message	CHL:_ ID:_ NOTICE: Media Error Recovered-0xD8001C7C
What Happened?	Data once stored on bad drive sectors has been successfully re-conducted from adjacent data blocks/stripes onto the original drive blocks or distributed onto other healthy sectors. ** Infortrend firmware translates SATA error codes into SCSI standard sense keys and sense codes. That's the error codes at the end of message.
What to Do?	Press ESC to clear the message.

Message	CHL:_ NOTICE: Fibre Channel Loop Connection Restored
What Happened?	Specific Fibre Channel loop connection restored.
What to Do?	Check proper host- or drive-side activities. Press ESC to clear the message.

Message	CHL:_ ID:_ NOTICE: Redundant Path for CHL:_ Restored
What Happened?	Redundant path, one of the drive-side loop members, for CHL:_ ID:_ is restored.
What to Do?	Check proper host- or drive-side activities. Press ESC to clear the message.

Message	CHL:_ NOTICE: Redundant Path for CHL:_ ID:_ Restored
What Happened?	The redundant path (CHL:_) connecting drive (CHL:_ ID:_) is restored.
What to Do?	Check proper host- or drive-side activities. Press ESC to clear the message.

Message	CHL:_ ID:_ Drive NOTICE: Drive Recovered Error Reported
What Happened?	Hard drive unrecoverable media error reported. A bad block is encountered in the specified hard drive. The RAID controller will ask the hard drive to retry.
What to Do?	Press ESC to clear the message.

Message	CHL:_ LIP(FF B) Detected

What Happened?	A LIP is issued on CHL:_.
What to Do?	Check proper host- or drive-side activities. Press ESC to clear the message.

Message	CHL:_ Host Channel Notification: Bus Reset Issued.
What Happened?	Bus reset on the CHL:_ host channel.
What to Do?	Check proper host- or drive-side activities. Press ESC to clear the message.

Message	CHL:_ ID:_ Drive NOTICE: Scan Drive Successful
What Happened?	A new drive or a replacement drive has been successfully scanned/recognized by the subsystem.
What to Do?	Press ESC to clear the message.

Message	CHL:_ ID:_ Drive NOTIFY: Block Successfully Reassigned – 0 (10B)
What Happened?	Writes conducted, retried, and media errors were discovered. Writes will then be re-directed to other drive sectors.
What to Do?	Press ESC to clear the message. If this message becomes frequent, the target disk drive might have developed serious media errors.

Message	CHL:_ ID_ Drive #SEV#: Media Error Encountered
What Happened?	Media errors are encountered in a hard drive.
What to Do?	Press ESC to clear the message. If this message becomes frequent, the target disk drive might have developed serious media errors.

C.3 General Target Events

IMPORTANT!

Firmware revision 3.63 added event strings for I2C slave devices. The I2C device error is indicated by a bracketed message such as "**(XXX Status Failure Detected!)**" by the end of a peripheral device event. The cause of an error can be an I2C bus arbitration error, bus timeout, or a slave device content error.

C.3.1 Critical:

Message	SAF-TE Device(0) ALERT: Power Supply Failure Detected(Idx:132)
What Happened?	Power supply failure detected through SAF-TE monitoring device.
What to Do?	Check the Power Supply Status LED on the front panel of system or check the PSU Status LED on the power supply unit on the rear panel. If the Amber LED is lit constantly, that indicates that a power

	supply has failed. If the LED is off, check the power cord is connected correctly or the corresponding AC is ok. Contact your RAID supplier for a replacement module.

Message	Peripheral Set 0 Device ALERT: Power Supply 0 Failure Detected
What Happened?	Power supply failure detected through an I2C serial bus.
What to Do?	Check the Power Supply Status LED on the front panel of system or check the PSU Status LED on the power supply unit on the rear panel. If the Amber LED is lit constantly, that indicates that a power supply has failed. If the LED is off, check the power cord is connected correctly or the corresponding AC is ok. Contact your RAID supplier for a replacement module.

Message	Peripheral Set 0 Device ALERT: Power Supply 0 Not Present
What Happened?	No power supply module is installed in the expected module slot.
What to Do?	Check proper module installation. If an installed module can not be detected, contact your RAID supplier.

Message	Peripheral Set 0 Device ALERT: Low Voltage Detected (5V)
What Happened?	Low voltage detected from the power supply. Power supply may have become unstable.
What to Do?	Check proper module installation and contact your RAID supplier for a replacement module.

Message	SES(C0 I0)Power Supply 0:Device Not Supported!
What Happened?	Unrecognizable device type. This event may result from an incorrect configuration with the SES remote device monitoring.
What to Do?	Check proper module installation and contact your RAID system supplier.

Message	SAF-TE Device(0) ALERT: Cooling Fan Not Installed(Idx:2)
What Happened?	No cooling fan module is installed in the expected module slot.
What to Do?	Check proper module installation. If an installed module can not be detected, contact your RAID supplier.

Message	SAF-TE Device(0) ALERT: Cooling Fan Failure Detected(Idx:2)
What Happened?	Cooling fan failure detected through SAF-TE monitoring device.
What to Do?	Check the Cooling Fan Status LED on the front panel of system. If the Amber LED is lit constantly, that indicates at least one cooling fan module has failed. Contact your RAID supplier for a replacement module.

Message	Peripheral Set 0 Device ALERT: Cooling Fan2 Failure Detected
What Happened?	Cooling fan failure detected through an I2C serial bus.

What to Do?	Check the Cooling Fan Status LED on the front panel of RAID/SBOD system. If the Amber LED is lit constantly, that indicates at least one cooling fan module has failed. Contact your RAID supplier for a replacement module.

Message	Peripheral Set 0 Device ALERT: Cooling Fan2 Note Present
What Happened?	No cooling fan module is installed in the expected module slot.
What to Do?	Check proper module installation or contact your RAID supplier for a replacement module.

Message	Peripheral Set 0 Device ALERT: Low Speed Detected (0 RPM)
What Happened?	Low rotation speed detected on cooling fan. This is a fault detected through an I2C serial bus.
What to Do?	Check proper module installation or contact your RAID supplier for a replacement module.

Message	SES(C0 I0)Cooling Fan 0:Device Not Supported!
What Happened?	Unrecognizable device type. This event may result from an incorrect configuration with the SES remote device monitoring.
What to Do?	Check proper module installation and contact your RAID system supplier.

Message	SAF-TE Device(0) ALERT: Elevated Temperature Alert
What Happened?	Elevated temperature detected through SAF-TE monitoring device. Improper installation site condition or cooling fan failure might have caused the system temperature to rise.
What to Do?	Check proper cooling and cooling fan working condition. Contact your RAID system supplier.

Message	Peripheral Device ALERT: CPU Cold Temperature Detected (0.0I)
What Happened?	Temperature below the lower CPU threshold detected. Improper installation site condition might have caused the situation.
What to Do?	Correct your installation site condition. Contact your RAID system supplier.

Message	Peripheral Device ALERT: Elevated Temperature Alert
What Happened?	High temperature threshold violated and detected through an I2C serial bus. Improper installation site condition or cooling fan failure might have caused the situation.
What to Do?	Correct your installation site condition and cooling fan failure. Contact your RAID system supplier if a replacement is necessary.

Message	Peripheral Set 0 Device ALERT: Temperature Sensor 0 Failure Detected
What Happened?	Temperature sensor failure reported through an I2C serial bus.

What to Do?	Contact your RAID system supplier.

Message	Peripheral Set 0 Device ALERT: Temperature Sensor 0 Not Present
What Happened?	Temperature sensor failed or serial bus configuration/connection faults occurred.
What to Do?	Contact your RAID system supplier.

Message	Peripheral Set 0 Device ALERT: Cold Detected
What Happened?	Temperature dropped below the lower system threshold detected. Improper installation site condition might have caused the situation.
What to Do?	Correct your installation site condition. Contact your RAID system supplier.

Message	SES(C0 I0)Temp Sensor 0: Device Not Supported!
What Happened?	Unrecognizable device type. This event may result from an incorrect configuration with the SES remote device monitoring.
What to Do?	Check proper module installation and contact your RAID system supplier.

Message	UPS Connection Is Absent
What Happened?	The COM2 link to a UPS device is absent.
What to Do?	Check cabling to the UPS device. If hardware faults occurred, contact your RAID system supplier.

Message	SAF-TE Device(2) ALERT: UPS Power Failure Detected
What Happened?	UPS power failure detected through SAF-TE monitoring device.
What to Do?	Contact your UPS supplier.

Message	Peripheral Set 0 Device ALERT: UPS 0 AC Power Failure Detected
What Happened?	AC power source failure reported by UPS and to the subsystem firmware through an I2C serial bus.
What to Do?	Check your power source connection or contact your RAID system supplier.

Message	Peripheral Set 0 Device ALERT: UPS 0 Battery Failure Detected
What Happened?	UPS battery failure reported by UPS and to the subsystem firmware through an I2C serial bus.
What to Do?	Contact your UPS supplier.

Message	Warning: UPS AC Power-Loss Detected
What Happened?	UPS AC power loss reported through the COM2 serial port connection.
What to Do?	Contact your UPS supplier.

Message	Warning: UPS Battery Low - __%
What Happened?	UPS battery charge low, may not be able to support subsystem during a power outage.
What to Do?	Wait for the UPS to recharge or you may use the Event Triggered mechanisms implemented with firmware. These mechanisms use conservative operation modes to reduce the chance of data loss in the event of power outage. Contact your UPS supplier.

Message	SES(C0 I0)UPS 2: Device Not Supported!
What Happened?	Unrecognizable device type. This event may result from an incorrect configuration with the SES remote device monitoring.
What to Do?	Check proper module installation and contact your RAID system supplier.

Message	ALERT: Enclosure ID Conflict/Invalid
What Happened?	Invalid drive configuration of Enclosure ID.
What to Do?	Check ID rotary switch at the front of RAID or SBOD enclosures. Please refer to documentation that came with your array.

C.3.2 Notification:

Message	SAF-TE Device(2) NOTICE: Fan Back On-Line(Idx:11)
What Happened?	A once missing or failed cooling fan is restored. This message is reported through the SAF_TE monitoring device.
What to Do?	Press ESC to clear this message.

Message	NOTICE: Fan Module A Back On-Line(FAN0 _RPM)
What Happened?	A once missing or failed cooling fan is restored. This message is reported through an I2C serial bus with RPM reading.
What to Do?	Press ESC to clear this message.

Message	NOTICE: Controller FAN 1 Back On-Line(0 RPM)
What Happened?	A once missing or failed cooling fan is restored. This message is reported through an I2C serial bus with RPM reading.
What to Do?	Press ESC to clear this message.

Message	Peripheral Device NOTICE: Fan Back On-Line
What Happened?	An I2C serial bus message indicating a once missing or failed cooling fan is restored.
What to Do?	Press ESC to clear this message.

Message	Peripheral Set 1 Device NOTICE: FAN 1 Back On-Line
What Happened?	An individual cooling fan once missing or failed is restored.
What to Do?	Press ESC to clear this message.

Message	Peripheral Set 1 Device NOTICE: FAN 1 is Present

What Happened?	An individual cooling fan once missing or failed is restored.
What to Do?	Press ESC to clear this message.

Message	Peripheral Set 1 Device NOTICE: FAN 1 Back On-Line(_RPM)
What Happened?	An individual cooling fan once missing or failed is restored, and the PRM reading displays. This applies to enclosure implementations using cooling fans capable of reporting rotation speed.
What to Do?	Press ESC to clear this message.

Message	SES(C0 I0) Cooling Fan 1:Fan Back On-Line
What Happened?	A cooling fan once missing or failed is restored. This message is reported through an SES monitoring device.
What to Do?	Press ESC to clear this message.

Message	SAF-TE Device(1) NOTICE: Temperature Back To Non-Critical LVs
What Happened?	Operating temperature back within normal temperature range.
What to Do?	Press ESC to clear this message.

Message	Peripheral Set 1 Device NOTICE: Temperature 1 Back To Non-Critical LVs
What Happened?	Temperature sensor 1 reports operating temperature is now within normal temperature range.
What to Do?	Press ESC to clear this message.

Message	Peripheral Set 1 Device NOTICE: Temperature 1 is Present
What Happened?	Temperature sensor 1 once missing and is now present on the serial bus.
What to Do?	Press ESC to clear this message.

Message	Peripheral Set 1 Device NOTICE: Temperature 1 Back To Non-Critical LVs (_I
What Happened?	Temperature sensor 1 reports operating temperature is now within normal temperature range and the Celsius reading displays.
What to Do?	Press ESC to clear this message.

Message	SES(C0 I0) Temp Sensor 1: Temperature Back to Non-Critical LVs
What Happened?	Temperature sensor 1 reports operating temperature is now within normal temperature range.
What to Do?	Press ESC to clear this message.

Message	SAF-TE Device(1) NOTICE: Power Supply Back On-Line (Idx:4)
What Happened?	A power supply once missing or failed is now restored. This event is reported through the SAF-TE monitoring interface.
What to Do?	Press ESC to clear this message.

Message	Power Supply 0 Back-Online
What Happened?	Power supply module 0 is back online
What to Do?	Press ESC to clear this message.

Message	Peripheral Set 0 Device NOTICE: PSU0 +12V Back On-Line
What Happened?	Power supply module 0 restored normal +12V voltage range
What to Do?	Press ESC to clear this message.

Message	Peripheral Set 0 Device NOTICE: PSU0 +12V is Present
What Happened?	Power supply module 0 restored normal +12V voltage range
What to Do?	Press ESC to clear this message.

Message	Peripheral Set 0 Device NOTICE: PSU0 +12V Back On-Line(5.0V)
What Happened?	Power supply module 0 restored normal +12V voltage range
What to Do?	Press ESC to clear this message.

Message	SES(C0 I0)Power Supply 2: Power Supply Back On-Line
What Happened?	Power supply module 2 once missing or failed is now restored.
What to Do?	Press ESC to clear this message.

Message	UPS Connection Detected
What Happened?	The COM2 serial link to UPS device is now valid.
What to Do?	Press ESC to clear this message.

Message	UPS AC Power Restored
What Happened?	UPS reports AC power source is now restored.
What to Do?	Press ESC to clear this message.

Message	UPS Battery Restored to Safe Level 0
What Happened?	UPS battery is charged to a safe level, able to protect system operation.
What to Do?	Press ESC to clear this message.

Message	SAF-TE Device(2) NOTICE: UPS Power Back On-Line
What Happened?	Power supplied by UPS is now restored.
What to Do?	Press ESC to clear this message.

Message	Peripheral Set 2 Device NOTICE: UPS 2 AC Power Back On-Line
What Happened?	UPS 2 AC power source restored.
What to Do?	Press ESC to clear this message.

Message	Peripheral Set 2 Device NOTICE: UPS 2 Battery Back On-Line
What Happened?	UPS 2 Battery charge or battery failure restored.
What to Do?	Press ESC to clear this message.

Message	SES(C0 I0)UPS 2: UPS Power Back On-Line
What Happened?	UPS 2 connected through SES interface reports power back online.
What to Do?	Press ESC to clear this message.

C.4 Controller Events

C.4.1 Critical:

Message	ALERT: +3.3V Low Voltage Detected(_._V)
What Happened?	The detected +3.3V voltage source is lower than the preset voltage threshold.
What to Do?	1. Check power supply working condition 2. Check voltage threshold settings, if you do not know the appropriate threshold value for voltage monitoring, please use the default setting for voltage threshold. 3. If the alert still occurs then it can result from controller hardware problems. Please contact your RAID system supplier for a replacement controller.

Message	ALERT: Controller FAN 0 Low Speed Detected (_RPM)
What Happened?	This only applies to controller fans in the SentinelRAID series. Low rotation speed detected.
What to Do?	Contact your RAID system supplier for a replacement.

Message	Controller ALERT: Redundant Controller Failure Detected
What Happened?	One RAID controller in a controller pair failed. This message indicates that its partner has taken the workload and managed the failover operation. The fault can result from the following: 1. A RAID controller failure 2. The partner controllers cannot communicate with each other. The other major reason can be the failure of the inter-controller communication channels. It could result from the defects of one or both RAID controllers, or that of the backplane.
What to Do?	Contact your RAID system supplier for a replacement controller and restore the dual-active configuration.

Message	Controller SDRAM ECC Multi-bits Errors Detected
What Happened?	Memory multi-bits errors.
What to Do?	Contact your RAID system supplier to return the subsystem/RAID controller for repair.

Message	Controller SDRAM ECC Single-bit Errors Detected
What Happened?	Memory single-bit errors.
What to Do?	Subsystem firmware should be able to manage the single-bit

	errors. Contact your RAID system supplier if this message becomes frequent.

Message	Controller SDRAM Parity Errors Detected
What Happened?	Memory parity errors.
What to Do?	Subsystem firmware should be able to manage the parity errors in memory. Contact your RAID system supplier if this message becomes frequent.

Message	Controller PCI Bus Parity Errors Detected
What Happened?	Serious system faults have occurred.
What to Do?	Stop host I/Os immediately to prevent further data integrity errors and contact your RAID system supplier.

Message	BBU Absent or Failed! Correct It and Reset Ctlr to Take Effect
What Happened?	BBU (battery cell pack) has failed or is accidentally removed. When a new module is installed, reset the subsystem for the configuration to take effect.
What to Do?	Check proper installation of the module or contact your RAID system supplier for a replacement module.

Message	Controller BBU Absent or Failed!
What Happened?	BBU (battery cell pack) has failed or is accidentally removed. When a new module is installed, reset the subsystem for the configuration to take effect.
What to Do?	Check proper installation of the module or contact your RAID system supplier for a replacement module.

Message	Controller BBU Failure Detected!
What Happened?	BBU (battery cell pack) has failed or is accidentally removed. When a new module is installed, reset the subsystem for the configuration to take effect.
What to Do?	Check proper installation of the module or contact your RAID system supplier for a replacement module.

Message	Controller BBU Thermal Shutdown/Enter Sleep-Mode!
What Happened?	The following conditions will cause the charger circuits to enter a low-power and self-protection state: 1. The temperature sensor on the charger circuit reports elevated temperature reading (>= 45 degree Celsius). 2. The BBU (battery cell pack) has been charged for over 14.3 hours. The BBU charger will enter a timer fault state.
What to Do?	1. Check proper ventilation within the subsystem. You may also check the readings from other sensors within the enclosure. Airflow might have been disrupted by the absence of one or several major modules or the failure of a cooling fan. Once the thermal condition is improved, charging will resume automatically.

	2. If a new battery module has been charged for over 14.3 hours and this event is issued, you may remove and re-install the battery module. An empty battery module may take more than 14.3 hours to be fully charged. There is a timer embedded with the charger, doing so can reset the timer. Charging will resume automatically.

C.4.2 Warning:

Message	Memory Not Sufficient to Fully Support Current Config.
What Happened?	The installed memory size does not support current configuration. Try using a DIMM module of a larger size.
What to Do?	Check proper installation of the module or contact your RAID system supplier for a certified module.

Message	ALERT: Cache volatilized.
What Happened?	Cache exists and does not belong to any specific LD. It could result from the one of the following reasons: 1. The cable broken between RAID system and JBOD if the specific LD was located at the JBOD. 2. The all member disk drives are missing of the specific LD. 3. Other reasons that RAID system cannot find the specific LD.
What to Do?	Check proper installation of the module, for example, the connection cable to JBOD. We recommend you to check the data and restore backup if necessary.

Message	Force Controller Write-Through on Trigger Cause
What Happened?	Preset conditions, e.g., elevated temperature or component failure, has forced the subsystem to adopt a more conservative operating mode by disabling the Write-Back caching.
What to Do?	Check proper working condition and correct system faults. Press ESC to clear the message.

C.4.3 Notification:

Message	Board1 Cold Temperature Back to Non-Critical LVs(_(C))
What Happened?	Board1 (usually the RAID controller main board) ventilation condition is restored to normal.
What to Do?	Check proper working condition and press ESC to clear the message.

Message	+12V Upper Voltage Back within Acceptable Limits (_._V)
What Happened?	+12V voltage dropped back within preset thresholds.

What to Do?	Check proper working condition and press ESC to clear the message.

Message	+12V Lower Voltage Back within Acceptable Limits (_._V)
What Happened?	+12V voltage dropped back within preset thresholds.
What to Do?	Check proper working condition and press ESC to clear the message.

Message	Memory is Now Sufficient to Fully Support Current Config.
What Happened?	Memory of a sufficient capacity is installed and the subsystem should work properly.
What to Do?	Check proper working condition and press ESC to clear the message.

Message	Controller NOTICE: NVRAM Factory Defaults Restored
What Happened?	The factory defaults for configuration profile is restored.
What to Do?	Press ESC to clear the message.

Message	Controller BBU is Charging!
What Happened?	Once a new BBU is installed or BBU charge drops below a preconfigured level, the subsystem will start charging its batteries.
What to Do?	Press ESC to clear the message. All batteries have a limited life expectancy. If this message appears too frequently, the BBU may have gradually lost its ability to hold charge.

Message	NOTICE: Controller BBU Present!
What Happened?	A BBU once missing or removed is now electrically connected to the subsystem.
What to Do?	Check proper working condition and press ESC to clear the message.

Message	NOTICE: Controller BBU Back On-Line!
What Happened?	A BBU once missing or removed is now restored to its normal working condition.
What to Do?	Check proper working condition and press ESC to clear the message.

Message	NOTICE: Controller BBU Fully Charged!
What Happened?	A BBU is fully charged.
What to Do?	Check proper working condition and press ESC to clear the message.

Message	Notify: Controller BBU Lifetime Exceeded - Replacement Recommended
What Happened?	The battery cell pack has reached its life expectancy.

What to Do?	Contact your local dealer for a replacement module.

Message	Controller NOTICE: Redundant Controller Firmware Updated
What Happened?	Firmware updated on a redundant controller RAID subsystem. Resetting the subsystem may be required for the new firmware to take effect.
What to Do?	Updating firmware on a complex system configuration is recommended for experienced engineers only. Unless for bug fixing, updating firmware is not always necessary.

Message	NOTICE: Force Controller Shutdown Start
What Happened?	System start to shutdown triggered by user or other configured event action.
What to Do?	Check proper working condition and press ESC to clear the message.

Message	NOTICE: Force Controller Shutdown Complete
What Happened?	System shutdown complete triggered by user or other configured event action.
What to Do?	Check proper working condition and press ESC to clear the message.

C.5 Data Service Events

C.5.1 Critical:

Message	Snapshot ALERT - SC ID:3BB6514353CF3B85 Space Thresh Exceeded 80%
What Happened?	Snapshot took up disk space in a Snapshot Container that exceeds the critical capacity threshold configured by user.
What to Do?	You may need to examine SC capacity threshold settings (purge policy) and configure the appropriate threshold value for your system. The related purge operation will be triggered by user's configuration and the old snapshot images will be purged to free more Snapshot Container space. User can setup multiple thresholds with different alerting levels that trigger event notification. For example, when snapshot 60% of container space is used, system will send a notification-level event. Then system will send a warning-level event and purge the old snapshot images when space usage reaches 80%.

C.5.2 Warning:

Message	Snapshot WARNING - SC ID:3BB6514353CF3B85 Creation Failed
What Happened?	Creation of a Snapshot Container failed. It could result from system internal error or hardware fatal failures.
What to Do?	Choose or recreate another valid Logical Drive and execute the snapshot container creation command again.

Message	Snapshot WARNING - SS ID:30A55D31700F2AAB Creation Failed
What Happened?	Creation of a Snapshot Set failed. It could result from insufficient disk space, system internal errors, or hardware fatal failures. If the original source volume has been mapped to host, user should to un-map the source volume before configuring the snapshot function.
What to Do?	Check disk space usage in the existing snapshot containers and purge the old snapshot images to free more space. We

	recommend users to configure appropriate purge policy, the capacity threshold, and the event notification setting.

Message	Snapshot WARNING - SI ID:1D3E09F11D48E27A Creation Failed
What Happened?	Creation of a Snapshot Image failed. It could result from insufficient disk space, the total number of snapshot images exceeds the configured limit, system internal errors, or hardware fatal failures.
What to Do?	Check usable space in snapshot containers and purge the unused snapshot images to free more space. Or check the upper threshold for the max. number of snapshot images for a specific snapshot set.

Message	Snapshot WARNING - SC ID:3BB6514353CF3B85 Expand Failed.
What Happened?	A Snapshot Container expansion failed. It could result from system errors.
What to Do?	Find another valid LD/partition to expand the snapshot container again. If it still fails, please check the integrity of snapshot container and the related logical drive status.

Message	Snapshot WARNING - SC ID:3BB6514353CF3B85 Space Thresh Exceeded 60%
What Happened?	Disk space in a Snapshot Container has been used up exceeding the capacity threshold previously configured by a user.
What to Do?	Check capacity threshold and configure an appropriate threshold value. Once the upper threshold is exceeded, the related purge action will be triggered and the old snapshot images will be purged to free more Snapshot Container space.

C.5.3 Notification:

Message	Snapshot NOTICE - SC ID:3BB6514353CF3B85 Creation Complete
What Happened?	Creation of a Snapshot Container is completed.
What to Do?	Clear the message.

Message	Snapshot NOTICE - SC ID:3BB6514353CF3B85 Extend Complete
What Happened?	The expansion of a Snapshot Container is completed.
What to Do?	Clear the message.

Message	Snapshot NOTICE - SC ID:3BB6514353CF3B85 Deleted
What Happened?	The specific Snapshot Container is deleted.
What to Do?	Clear the message.

Message	Snapshot NOTICE - SS ID:30A55D31700F2AAB Creation Complete
What Happened?	The creation of a Snapshot Set is completed.
What to Do?	Clear the message.

Message	Snapshot NOTICE - SS ID:30A55D31700F2AAB Deleted
What Happened?	The specific Snapshot Set is deleted. User have to remap the source volume to host side ID/LUN again to continue I/Os.
What to Do?	Clear the message.

Message	Snapshot NOTICE - SI ID:1D3E09F11D48E27A Creation Complete
What Happened?	A Snapshot Image is created.
What to Do?	Clear the message.

Message	Snapshot NOTICE - SI ID:1D3E09F11D48E27A Deleted
What Happened?	The specific Snapshot Image is deleted. The necessary data will be merged to other snapshot images to keep the point-in-time image.
What to Do?	Clear the message.

Message	Snapshot NOTICE - SI ID:1D3E09F11D48E27A Activated
What Happened?	The specific Snapshot Image is activated. The updated data at source volume will be recorded and stored in the activated snapshot image.
What to Do?	Clear the message.

Message	Snapshot NOTICE - SC ID:3BB6514353CF3B85 Space Thresh Exceeded 40%
What Happened?	40% of the disk space in a Snapshot Container has been used up exceeding the capacity notification threshold.
What to Do?	Check capacity threshold settings and configure the appropriate threshold value for your system. The related purge operation will be triggered by the exceeded capacity threshold and the old snapshot images will be purged to free more Snapshot Container space.

Message	Snapshot NOTICE – SC ID:3BB6514353CF3B85 Space Thresh Restored 50%
What Happened?	Disk space in the Snapshot Container has been released, used space drops below the preset threshold. System has successfully purged old snapshot images to regain disk space.
What to Do?	Clear the message.

Message	Snapshot NOTICE – SC ID:3BB6514353CF3B85 Purge Triggered-80%
What Happened?	The purge action is triggered when the capacity threshold is violated.
What to Do?	Clear the message.

Message	Snapshot NOTICE – SS ID:30A55D31700F2AAB Prune Triggered – 80%"
What Happened?	The Snapshot scheduler prune action is triggered when system detects the number or life span of the snapshot images have exceeded the preset threshold.
What to Do?	Clear the message.

Message	Snapshot NOTICE Non-optimal configuration may impact performance
What Happened?	System will check the related settings at boot time for the optimal configuration to execute snapshot function, including: 1. RAID Cache write-back policy 2. Disk buffer disabled 3. BBU protection enabled 4. BBU Event –triggered operation enabled System will automatically default to the above configuration when snapshot function is activated.
What to Do?	Restore the suggested configuration value for snapshot function. System will prompt user if the configuration is not configured as the default values for the optimal snapshot function.