Infortrend

External RAID Subsystems & Controllers

Generic Operation Manual

Revision 1.6.9

Firmware Revision: 3.42x

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@infortrend.com.tw
support@infortrend.com.tw
http://www.infortrend.com.tw

Americas
Infortrend Corporation
3150 Coronado Dr., Unit C
Santa Clara, CA 95054
USA
Tel: +1-408-988-5088
Fax: +1-408-988-6288
sales@infortrend.com
tsd@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@infortrend.com.cn
support@infortrend.com.cn
http://www.infortrend.com.cn

Europe (EMEA)
Infortrend Europe Limited
5 Elmwood Chineham Business Park
Basingstoke, Hampshire
RG24 8WG, UK
Tel: +44-1256-70-77-00
Fax: +44-1256-70-78-89
sales@infortrend-europe.com
support@infortrend-europe.com
http://www.infortrend-europe.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@infortrend.co.jp
support@infortrend.co.jp
http://www.infortrend.co.jp

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, the Infortrend logo, RAIDWatch, and EonStor® are registered trademarks of Infortrend Technology, Inc. Other names prefixed with "IFT" are trademarks of Infortrend Technology, Inc.

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

DEC and Alpha are registered trademarks of Compaq Computer Corp. (formerly of Digital Equipment Corporation).
Microsoft, Windows, Windows NT and MS-DOS are registered trademarks of Microsoft Corporation in the U.S. and other countries.

Novell and NetWare are registered trademarks of Novell, Inc. in the U.S. and other countries.
SCO, OpenServer, and UnixWare are trademarks or registered trademarks of The Santa Cruz Operation, Inc. in the U.S. and other countries.

Solaris is a trademark of SUN Microsystems, Inc.
UNIX is a registered trademark of The Open Group in the U.S. and other countries. All other names, brands, products or services are trademarks or registered trademarks of their respective companies.

RMA Policy

Please visit our websites (www.infortrend.com, www.infortrend.com.tw, www.infortrend.com.cn, or www.infortrend-europe.com) where our RMA policy is given a detailed explanation.

Supported Models

This manual supports the following Infortrend controllers/subsystems:

- EonStor: subsystems that come with SCSI, iSCSI, or Fibre host channels and SCSI, Fibre, or SATA drive interfaces.
- SentinelRAID: SCSI-based external RAID controllers (including the 5.25" full-height and 1U canister configuration)
- EonRAID: Fibre-based external RAID controllers (including the 1U canister controller head configurations)

Table of Contents

Chapter 6 Terminal Screen Messages

Chapter 7 Terminal Operation

Chapter 8 Fibre Operation

Chapter 9 Host-side and Drive-side Parameters

Chapter 10 Enclosure Monitoring

Chapter 11 Data Integrity

Chapter 12 iSCSI Options

Chapter 13 Redundant Controller 13-1

Chapter 14 Record of Settings

Chapter 15 Array Expansion

Chapter 16 S.M.A.R.T. Configuration

Chapter 17 Implementations for AV Applications

Appendix A Firmware Functionality Specifications

Appendix B System Functions: Upgrading Firmware

Functional Table of Contents

This functional table of contents helps you to quickly locate the descriptions of firmware functions.

List of Tables

List of Figures

Chapter 1

Chapter 2

Chapter 5

Chapter 7

Chapter 8

Chapter 9

Chapter 13

Chapter 15

About This Manual

This manual provides all of the necessary information that a system administrator needs to configure and maintain one of Infortrend's external RAID controllers or subsystems. For hardware-related information, please refer to the ***Installation and Hardware Reference Manual*** that came with your RAID subsystem/controller. Also available is the ***User's Manual*** for the Java-based GUI RAID manager for remote and concurrent management of RAID systems.

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

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

Chapter 1	Introduces basic RAID concepts and configurations, including RAID levels, logical drives, spare drives, and the use of logical volumes. It is recommended that users unfamiliar with RAID technologies should read this chapter before creating a configuration.
Chapter 2	Tells the user how to begin with a RAID. At the beginning of this chapter, we raise some basic questions to which the user should know the answers prior to creating a RAID.
Chapter 3	Teaches the user how to configure the RS-232C terminal emulation interface and the connection through a LAN port.
Chapter 4	Helps the user to understand screen messages on the LCD display.
Chapter 5	Gives step-by-step instructions on creating a RAID using the LCD keypad panel.
Chapter 6	Teaches the user how to interpret the information found on the RS-232 terminal emulation.
Chapter 7	Gives step-by-step instructions on how to create a RAID via the RS-232 session.
Chapter 8	Includes all the Fibre Channel-specific functions implemented since the firmware release 3.12.
Chapter 9	Provides the advanced options for the host and drive channel configurations including I/O queuing, number of LUNs, host loop, in-band, and disk access characteristics settings.
Chapter 10	Shows how to supervise system operating and component status.
Chapter 11	Details various methods, either manual or automatic, for maintaining array data integrity.
Chapter 12	Provides the configuration options implemented

	with iSCSI RAID subsystems
Chapter 13	Addresses the concerns regarding the redundant controller configuration and the configuration process.
Chapter 14	Provides the recording forms with which a system administrator can keep a record of his configuration.
Chapter 15	Shows how to expand a configured array or logical volume.
Chapter 16	Shows how to use the SMART-related functionality.
Chapter 17	Illustrates the Maximum Response Time setup
Appendix A	Lists the important firmware features supported with the firmware version, arranged in accordance with the latest firmware version as of press date.
Appendix B	Teaches the user how to upgrade firmware and boot record.
Appendix C	Lists all of the controller event messages.

Firmware Version & Other Information

Firmware Version: 3.42x

Date: May/30/05

Revision History:

Version 1.0:	• Initial release
Version 1.1:	• Added redundant controller configuration
Version 1.2:	• Added host-side and drive-side SCSI parameters • Added S.M.A.R.T. with implemented fault-prevention methods. • Added system functions • Added Fault-bus configuration to be compatible with 3101 and 3102 series • Added Host-side interface installation details • Added Event Messages for error message identification • Added all advanced functions available since 2.23K and 3.11F upward • Added a functional table of contents for quick searching functions • Moved SCSI/Fibre Cable Specifications to Hardware Manual
Version 1.3:	• Added Chapter 8 "Fibre Operation" for the new functions available since firmware release 3.12
Version 1.4:	• Added firmware features available with firmware revisions 3.14, 3.15, and 3.21 • Revised details about redundant controllers, host LUN mapping, etc. • Modified string definitions in Chapter 14 "In-band SCSI Drives and Utilities" section • Corrected descriptions of "Controller Unique Identifier" • Added the configuration process for out-of-band configuration via LAN port
Version 1.5:	• Removed Chapter 14 • Revised the descriptions for some functional items • Added firmware features available from revision 3.25

Version 1.61:

- Added features available by revision 3.31
- Removed Appendix E
- Moved array expansion to Chapter 12
- Added variable stripe size, write policy per array
- Added Media Scan
- Added controller immediate array availability, time zone, date and time setting
- Added IO channel diagnostics
- Added controller auto-shutdown and cache-flush mechanisms
- Added system monitoring via enclosure modules
- Added disabling cache coherency using write-through mode
- Added descriptions about new firmware utility items
- Added details about enabling RAIDWatch and its sub modules via Ethernet port

Version 1.62:

- Added functions previously included in controller Hardware Manual
- Added remote redundant functionality
- Updated firmware event messages
- Updated LCD panel navigation map
- Added Appendix E - Dial-Out Modem

Version 1.63:

- Added Task Scheduler
- Added Periodic Cache Flush
- Added parity regeneration options
- Added Event Triggered Operations
- Updated firmware event messages
- Added new chapters for specific functionality groups
- Removed Appendix A - LCD Navigation Road Map. The Road Map will be available as print-out copies.

Version 1.64: Added description or configuration procedure for the following:

- Maximum Drive Response Time
- Host-side 64-bit LBA support
- Host-side Ordered Tag support
- Host-side Re-routing support
- Adaptive Write-back/Write-through switching
- SSH connection support
- Issue LIP command
- Updated firmware event messages

Version 1.65: Corrected functional details

Version 1.66:	Added Chapter 12 “iSCSI Options”
Version 1.67:	Updated event messages
Version 1.68:	Updated event messages and added the Jumbo frame configuration option
Version 1.69	• Removed the Jumbo frame configuration option • Updated redundant controller details in Chapter 13 • Added description for NCQ support • Updated Communication Parameters • Modified in-band connection settings and Ethernet connection requirements. • Updated supported features in Appendix A • Removed Rolling Firmware Upgrade • Added Add Host – ID/WWN Label Declaration to Chapter 8 • Updated Solaris 10 cylinder/head/sector mapping

Chapter

1

RAID Functions: An Introduction

Redundant Arrays of Independent Disks, or RAID, offers the following advantages: availability, capacity, and performance. Choosing the right RAID level and drive failure management can increase capacity and performance, subsequently increasing availability. Infortrend's external RAID controllers and subsystems provide complete RAID functionality and enhanced drive failure management.

1.1 Logical Drive

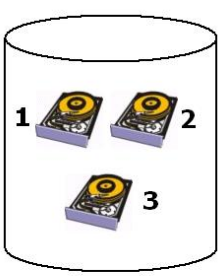

Logical Drive

The advantages mentioned above are achieved by creating "logical drives." A logical drive is an array of independent physical drives. The logical drive appears to the host as a contiguous volume, the same as a local hard disk drive does.

Figure 1 - 1 Logical Drive

The following section describes the different methods to create logical arrays of disk drives, such as spanning, mirroring and data parity. These methods are referred to as "RAID levels."

1.2 Logical Volume

What is a logical volume?

The concept of a logical volume is very similar to that of a logical drive. A logical volume is the combination of one or several logical drives. These logical drives are combined into a larger capacity

using the RAID 0 method (striping). When data is written to a logical volume, it is first broken into data segments and then striped across different logical drives in a logical volume. Each logical drive then distributes data segments to its member drives according to the specific RAID level it is composed of.

The member logical drives can be composed of the same RAID level or each of a different RAID level. A logical volume can be divided into a maximum of 64 partitions. During operation, the host sees a non-partitioned logical volume or a partition of a logical volume as one single physical drive.

1.3 RAID Levels

Using a RAID storage subsystem has the following advantages:

- Provides disk spanning by weaving all connected drives into one single volume.
- Increases disk access speed by breaking data into several blocks when reading/writing to several drives in parallel. With RAID, storage speed increases as more drives are added as the channel bus allows.
- Provides fault-tolerance by mirroring or parity operation.

What Are the RAID Levels?

Table 1 - 1 RAID Levels

RAID Level	Description	Capacity	Data Availability
NRAID	Non-RAID	N	
RAID 0	Disk Striping	N	==NRAID
RAID 1 (0+1)	Mirroring Plus Striping (if N>1)	N/2	>>NRAID ==RAID 5
RAID 3	Striping with Parity on dedicated disk	N-1	>>NRAID ==RAID 5
RAID 5	Striping with interspersed parity	N-1	>>NRAID ==RAID 5
RAID 10 **(Logical Volume)**	Striping with RAID 1 logical drives	/	>>NRAID >>RAID 5
RAID 30 **(Logical Volume)**	Striping with RAID 3 logical drives	/	>>NRAID >>RAID 5
RAID 50 **(Logical Volume)**	Striping with RAID 5 logical drives	/	>>NRAID >>RAID 5

NOTE: *Drives on different channels can be included in a logical drive and logical drives of different RAID levels can be used to compose a logical volume. There are more combinations than RAID 10, 30, and 50.*

RAID Level	Performance Sequential	Performance Random
NRAID	Drive	Drive
RAID 0	R: Highest W: Highest	R: High W: Highest
RAID 1 (0+1)	R: High W: Medium	R: Medium W: Low
RAID 3	R: High W: Medium	R: Medium W: Low
RAID 5	R: High W: Medium	R: High W: Low

NRAID

Disk Spanning

NRAID	
Minimum Disks Required	1
Capacity	N
Redundancy	No

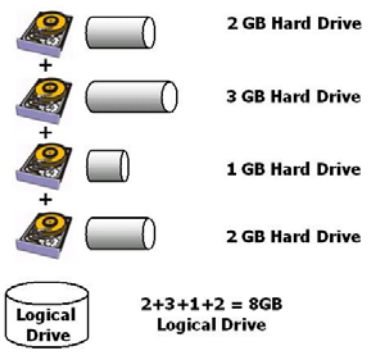

NRAID stands for Non-RAID. The capacity of all drives is combined to become one logical drive (no block striping). In other words, the capacity of the logical drive is the total capacity of the physical member drives. NRAID does not provide data redundancy.

Figure 1 - 2 NRAID

JBOD

Single Drive Control

JBOD	
Minimum Disks Required	1
Capacity	1
Redundancy	No

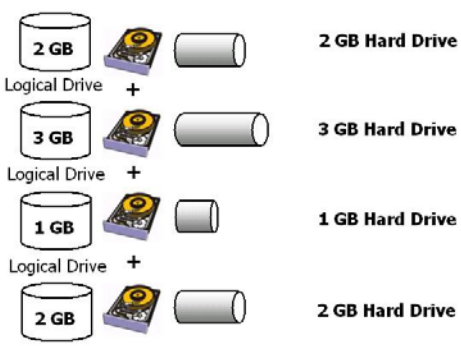

JBOD stands for Just a Bunch of Drives. The controller treats each drive as a stand-alone disk; therefore, each drive is an independent logical drive. JBOD does not provide data redundancy.

Figure 1 - 3 JBOD

RAID 0

Disk Striping

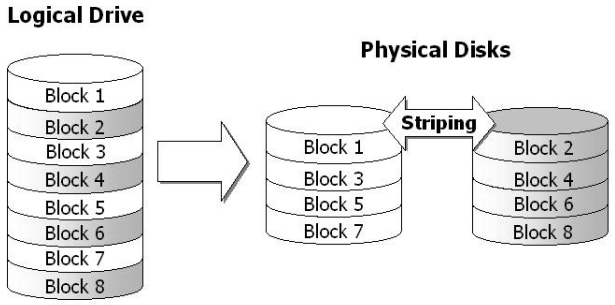

RAID 0	
Minimum Disks Required	2
Capacity	N
Redundancy	No

RAID 0 provides the highest performance but no redundancy. Data in the logical drive is striped (distributed) across several physical drives.

Figure 1 - 4 RAID 0

RAID 1

Disk Mirroring

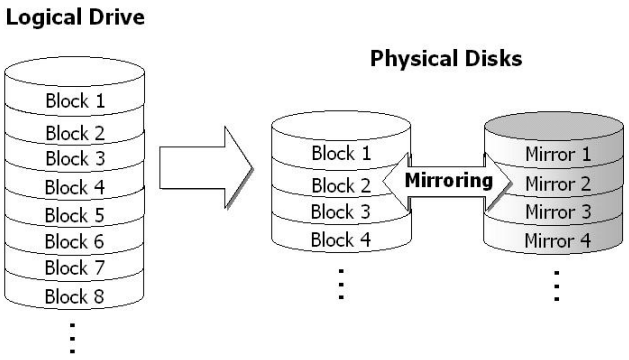

RAID 1	
Minimum Disks Required	2
Capacity	N/2
Redundancy	Yes

RAID 1 mirrors the data stored in one hard drive to another. RAID 1 can only be performed with two hard drives. If there are more than two hard drives, RAID (0+1) will be automatically applied.

Figure 1 - 5 RAID 1

RAID (0+1)

Disk Striping with Mirroring

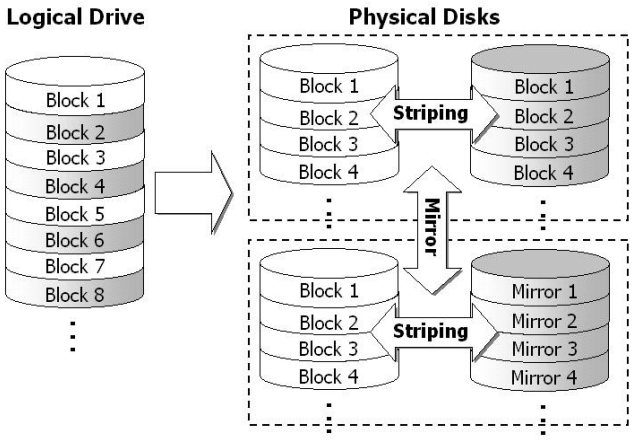

RAID (0+1)	
Minimum Disks Required	4
Capacity	N/2
Redundancy	Yes

RAID (0+1) combines RAID 0 and RAID 1 - Mirroring and Striping. RAID (0+1) allows multiple drive failure because of the full redundancy of the hard drives. If there are more than two hard drives assigned to perform RAID 1, RAID (0+1) will be automatically applied.

Figure 1 - 6 RAID (0+1)

IMPORTANT!

"RAID (0+1)" will not appear in the list of RAID levels supported by the controller. If you wish to perform RAID 1, the controller will determine whether to perform RAID 1 or RAID (0+1). This will depend on the number of drives that have been selected for the logical drive.

RAID 3

Disk Striping
with Dedicated Parity Disk

RAID (0+1)	
Minimum Disks Required	3
Capacity	N-1
Redundancy	Yes

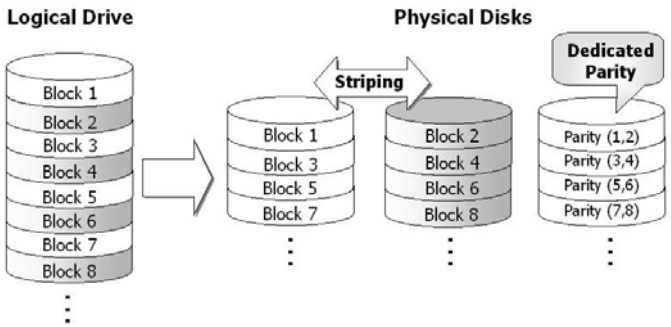

RAID 3 performs Block Striping with Dedicated Parity. One drive member is dedicated to storing the parity data. When a drive member fails, the controller can recover/regenerate the lost data of the failed drive by comparing and re-calculating data on the remaining drives.

Figure 1 - 7 RAID 3

RAID 5

Striping with Interspersed Parity

RAID 5	
Minimum Disks Required	3
Capacity	N-1
Redundancy	Yes

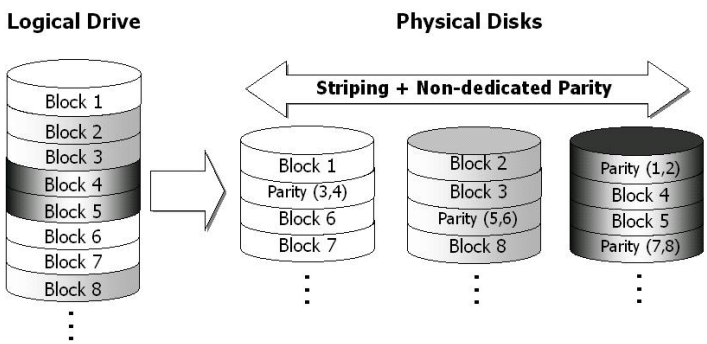

Figure 1 - 8 RAID 5

RAID 5 is similar to RAID 3 but the parity data is not stored in a dedicated hard drive. Parity information is interspersed across the

drive array. In the event of a drive failure, the controller can recover/regenerate the lost data of the failed drive by comparing and re-calculating data on the remaining drives.

RAID 10, RAID 30, and RAID 50 are implemented as logical volumes; please refer to the proceeding discussions for details.

1.4 Spare Drives

Global and Local Spare Drives

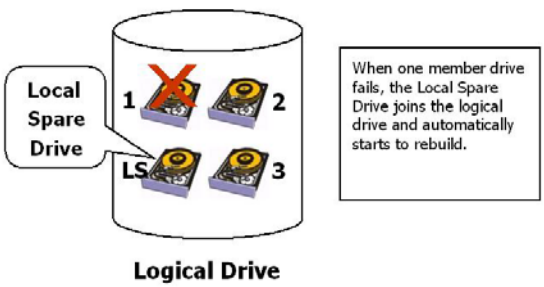

A Local Spare Drive is a standby drive assigned to serve one specified logical drive. When a member drive of this specified logical drive fails, the Local Spare Drive becomes a member drive and automatically starts to rebuild.

Figure 1 - 9 Local (Dedicated) Spare

Sp

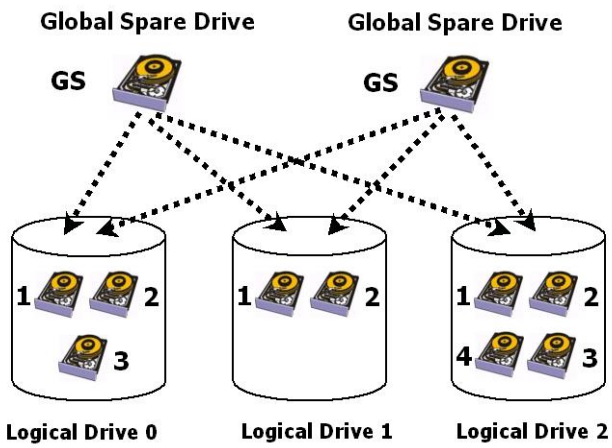

A Global Spare Drive serves more than one specified logical drive. When a member drive from any of the logical drives fails, the Global Spare Drive will join that logical drive and automatically start to rebuild.

Figure 1 - 10 Global Spare

Global Spare Drives serve any logical drive.

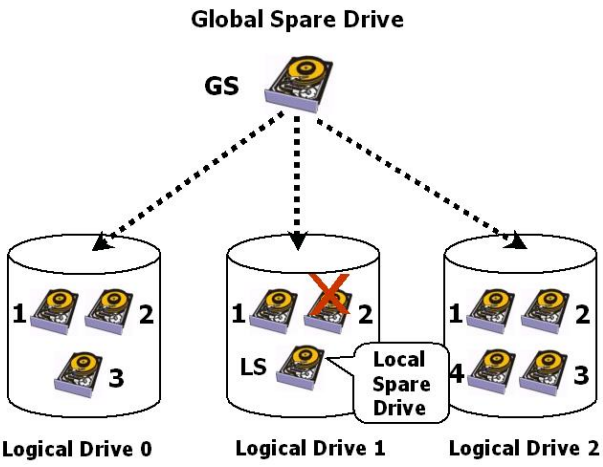

The example on the left provides both Local Spare Drive and Global Spare Drive functions. On certain occasions, applying these two functions together will better fit various needs. Take note that the **Local Spare Drive always has higher priority than the Global Spare Drive.**

Figure 1 - 11 Global Spare Rebuild

When a member drive from any logical drive fails, the Global Spare Drive joins that logical drive and automatically starts to rebuild.

In the example shown below, the members of Logical Drive 0 are 9GB drives, and the members in Logical Drives 1 and 2 are 4GB drives.

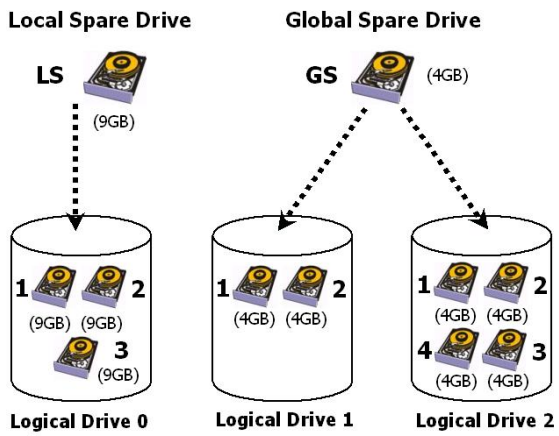

A Local Spare always has higher priority than a Global Spare.

It is not possible for the 4GB Global Spare Drive to join Logical Drive 0 because of its insufficient capacity. However, using a 9GB drive as the Global Spare Drive for a failed drive that comes from Logical Drive 1 or 2 will bring a huge amount of excess capacity since these logical drives require 4GB only. As shown on the left, the 9GB Local Spare Drive will aid Logical Drive 0 once a drive in this logical drive fails. If the failed drive is in Logical Drive 1 or 2, the 4GB Global Spare drive will immediately give aid to the failed drive.

Figure 1 - 12 Mixing Local and Global Spares

1.5 Identifying Drives

Whenever there is a failed drive in the RAID 5 logical drive, make it a point to replace the failed drive with a new, healthy drive to keep the logical drive working.

If you replaced the wrong drive when trying to replace a failed drive you will no longer be able to access the logical drive because you have inadvertently failed another drive.

To prevent this from happening, the controller provides an easy way to identify the faulty drive. By forcing certain drive LEDs to light for a configurable period of time, the faulty drive can be identified, and thus reduce the chance of removing the wrong drive. This function can be especially helpful in an installation site operating with hundreds of drives.

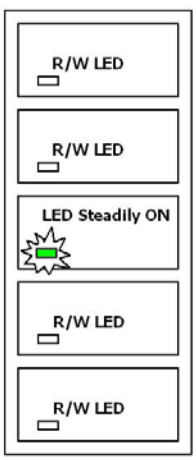

Flash Selected Drive

The Read/Write LED of the drive you selected will light steadily for a configurable period of time, from 1 to 999 seconds.

Figure 1 - 13 Flash Selected Drive

Flash All Drives

The Read/Write LEDs of all connected drives will light for a configurable period of time. If the LED of the defective drive did not light on the "Flash Selected SCSI Drive" function, use "Flash All SCSI Drives" to verify the fault. If the "Flash All SCSI Drives" function is executed, and the defective drive's LED still does not respond, it can be a drive tray problem or the drive is dead.

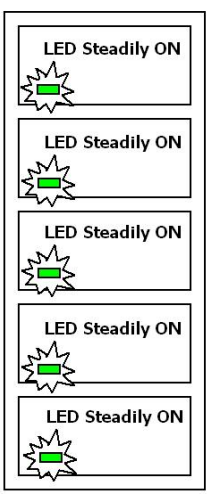

Figure 1 - 14 Flash All Drives

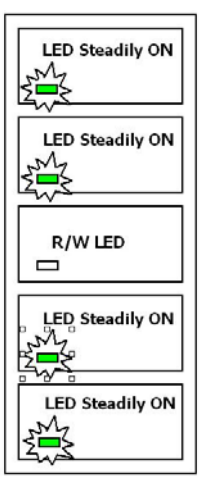

Flash All but Selected Drives

Except the selected drive, the Read/Write LEDs of all connected drives will light for a configurable period of time ranging from 1 to 999 seconds. If an administrator cannot be sure of the exact location of a specific drive, this function will help to indicate where it is. This can prevent removal of the wrong drive when a drive fails and is about to be replaced.

Figure 1 - 15 Flash All but Selected Drives

The drive-identifying function can be selected from "Main Menu"/"View and Edit SCSI Drives"/"Identify SCSI Drives."

1.6 Rebuild

Automatic Rebuild and Manual Rebuild

1. Automatic Rebuild

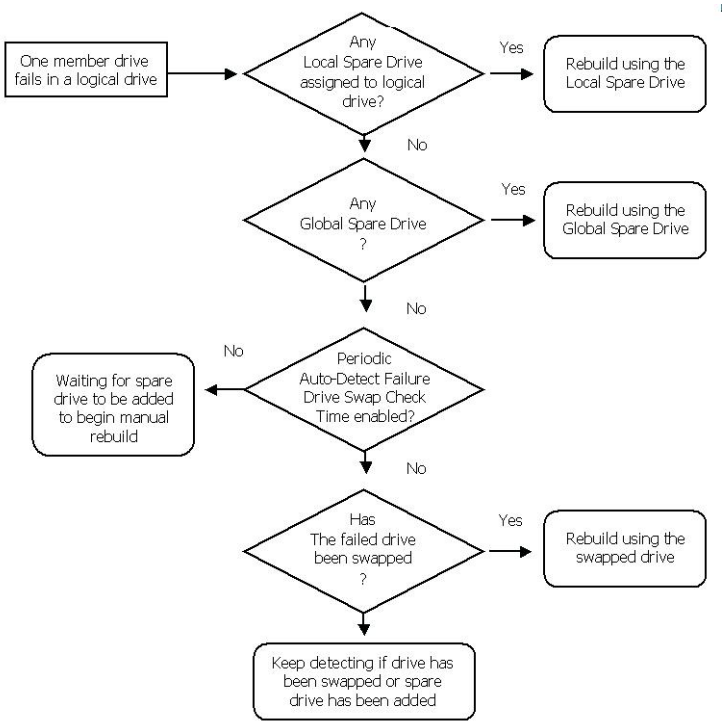

Figure 1 - 16 Automatic Rebuild

Rebuild with Spare: When a member drive in a logical drive fails, the controller will first examine whether there is a Local Spare Drive assigned to this logical drive. If yes, rebuild is automatically started.

If there is no Local Spare available, the controller will search for a Global Spare. If there is a Global Spare, rebuild automatically begins using the Global Spare.

Failed Drive Swap Detected: If neither a Local Spare Drive nor Global Spare Drive is available, and the "***Periodic Auto-Detect Failure Drive Swap Check Time***" is "Disabled," the controller will not attempt to rebuild unless the user applies a forced-manual rebuild.

When the "***Periodic Auto-Detect Failure Drive Swap Check Time***" is "Enabled" (i.e., a check time interval has been selected), the controller will detect whether a faulty drive has been swapped (by checking the failed drive's channel/ID). Once the failed drive has been replaced with a healthy drive, the rebuild will begin immediately.

If the failed drive is not swapped but a local spare is added to the logical drive, rebuild will begin with the spare.

If the S.M.A.R.T. function is enabled on drives and the reaction scheme is selected for securing data on a failing drive, the spare will also be used for restoring data. Please refer to ***Chapter 9***, Advanced Functions, for more details.

2. Manual Rebuild

When a user applies a forced-manual rebuild, the controller will first examine whether there is any Local Spare assigned to the logical drive. If yes, it will automatically start to rebuild.

If there is no Local Spare available, the controller will search for a Global Spare. If there is a Global Spare, logical drive rebuild will be automatically conducted.

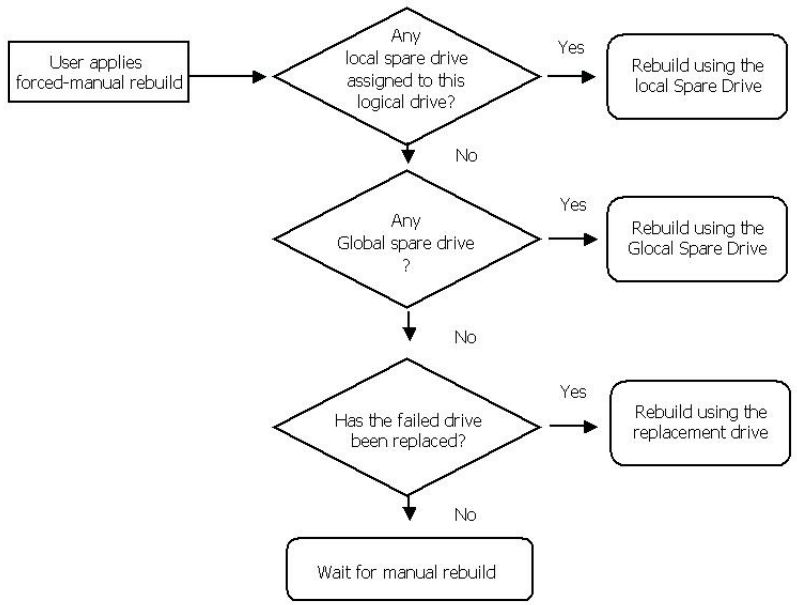

Figure 1 - 17 Manual Rebuild

If none of the spares are available, the controller will examine the SCSI channel and ID of the failed drive. Once the failed drive has been replaced by a healthy one, it starts to rebuild using the new drive. If there is no available drive for rebuilding, the controller will

not attempt to rebuild until the user applies another forced-manual rebuild.

3. Concurrent Rebuild in RAID (0+1)

RAID (0+1) allows multiple drive failures and rebuilds to be concurrently conducted on more than one of its members. Drives newly swapped must be scanned and set as Local Spares. These drives will be used for rebuilding at the same time - you do not need to repeat the rebuild process for each member drive.

1.7 Logical Volume (Multi-Level RAID)

What Is a Logical Volume?

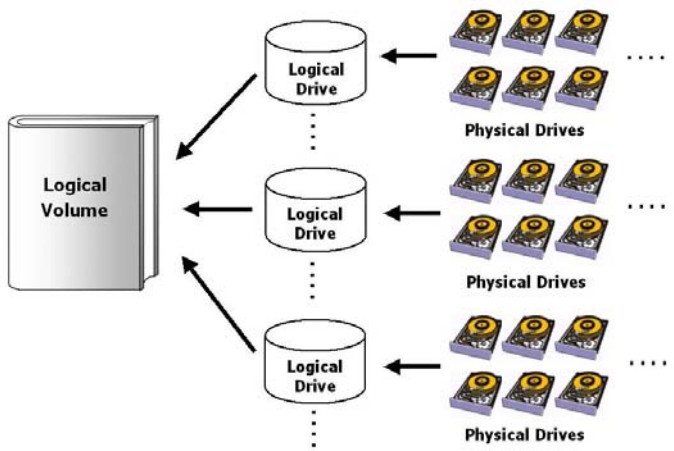

Figure 1 - 18 Logical Volume

A logical volume is a combination of RAID 0 (Striping) and other RAID levels. Data written to a logical volume is first broken into smaller data segments and striped across different logical drives in a logical volume. Each logical drive then distributes data segments to its member drives according to its mirroring, parity, or striping scheme. A logical volume can be divided into a maximum of eight partitions. During normal operation, the host sees a non-partitioned logical volume or a partition of a partitioned logical volume as one single physical drive.

The benefits of using a logical volume are achieved by:

1. Extending the MTBF (mean time between failure) by using more redundancy drives (spare drives).

2. Decreasing the time to rebuild and reducing the chance of data loss caused by simultaneous drive failures because drives are included in different drive groups using a multi-level logical structure.

3. Avoiding the chance of data loss caused by channel bus failure with flexible drive deployment.

As diagramed below, numerous drives can be included in a logical drive, and one of them is used for redundancy. By grouping these drives into several logical drives, and then into a logical volume, the chance of two drives failing in a logical unit is greatly reduced. Each logical drive can have one or more local spares. A failed drive can be immediately replaced by a local spare, reducing the risk of losing data if another drive fails soon afterwards.

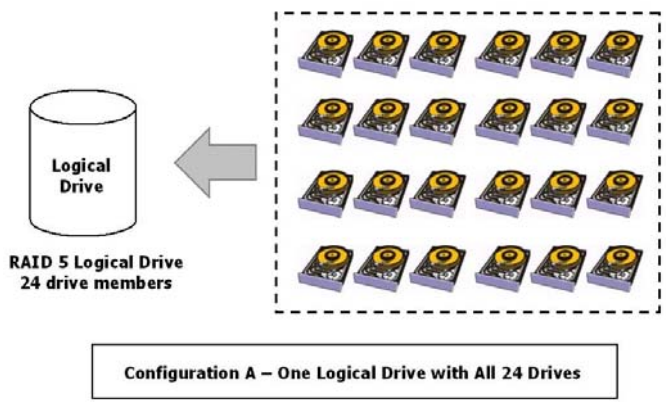

Figure 1 - 19 Logical Drive Composed of 24 Drives

As illustrated above, Configuration A is a RAID 5 logical drive consisting of 24 physical drives. Configuration B is a logical volume made of four RAID 5 logical drives.

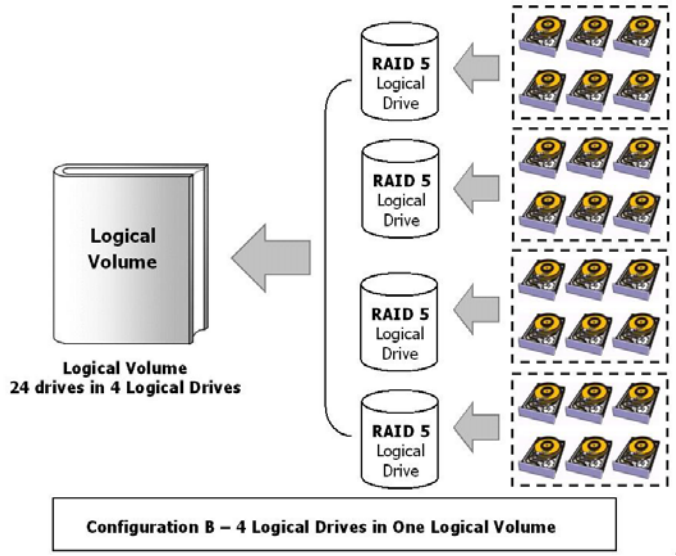

Figure 1 - 20 Logical Volume Composed of 4 Logical Drives and 24 Drives

Configuration B can help reduce the chance of encountering points of failure:

a) Higher Redundancy: Configuration A has one dedicated spare, while Configuration B allows the configuration of four spares. In Configuration B, the risk of simultaneous drive failure in a logical drive is significantly reduced compared to Configuration A. The total array capacity is comparatively smaller by the use of spares.

b) Less Rebuild Time: The time during rebuild is a time of potential hazard. For example, a RAID 5 logical drive can only withstand single drive failure - if another drive fails during the rebuild process, data will be lost. The time span for rebuilding a faulty drive should be minimized to reduce the possibility of having two drives fail at the same time.

Configuration A is a large logical drive and takes a long time to rebuild. All members will be involved during the rebuild process. In Configuration B, the time span is shorter because only six members will participate when rebuilding any of the logical drives.

c) Channel Failure Protection: Channel failure may sometimes result from unlikely matters such as a cable failure. A channel failure will cause multiple drives to fail at the same time and inevitably lead to a fatal failure. Using a logical volume with drives coming from different drive channels can avoid this point of failure.

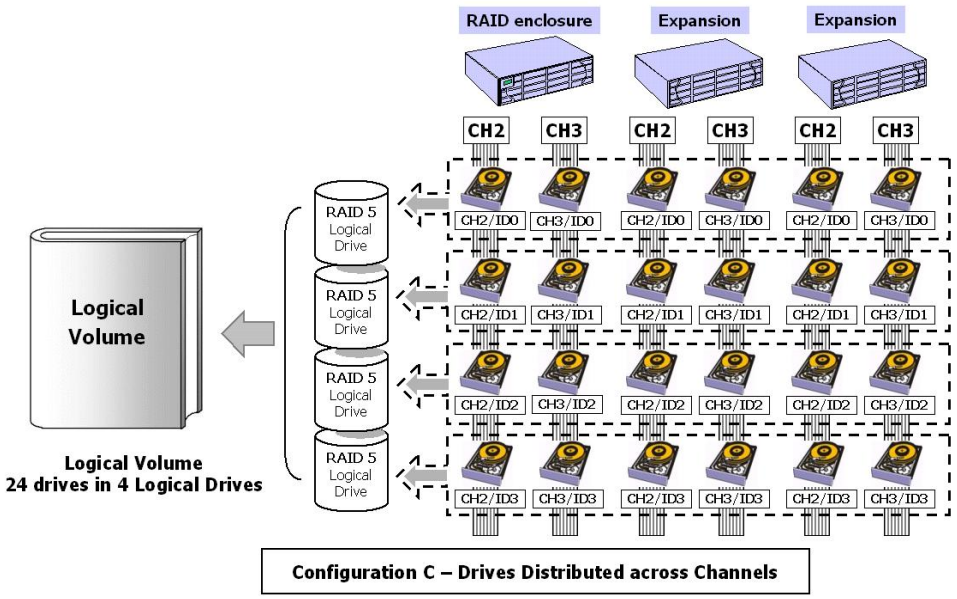

Figure 1 - 21 Logical Volume with Drives on Different Channels

As illustrated above, if one of the drive channels fails, each logical drive loses one of its members. Logical drives still have the chance to rebuild its members. Data remains intact and the rebuild can be performed after the failed channel is recovered. No access interruptions to the logical volume will be experienced from the host side.

Spare Drives Assigned to a Logical Volume?

A Local Spare cannot be assigned to a Logical Volume. If a drive fails, it fails as a member of a logical drive; therefore, the controller allows a Local Spare's assignment to logical drives rather than logical volumes.

Limitations:

The logical volume can not have any logical drive stated as "fatal failed." If there is any failed drive in any of its member logical drives, the controller will start to rebuild that logical drive. If any of the member logical drives fail fatally, the logical volume fails fatally and data will not be accessible.

To avoid a logical volume failure:

1. Logical drives as members of a logical volume should be configured in RAID levels that provide redundancy, i.e., - RAID levels 1 (0+1), 3, or 5.

2. Rebuild the logical drive as soon as possible whenever a drive failure occurs. Use of local spares is recommended.

3. A logical drive should be composed of physical drives from different drive channels. Compose the logical drive with drives from different drive channels to avoid the fatal loss of data caused by bus failure.

Partitioning the Logical Drive or Partitioning the Logical Volume?

Once a logical drive has been divided into partitions, the logical drive can no longer be used as a member of a logical volume. The members of a logical volume should have one partition only with the entire capacity.

If you want to use a partitioned logical drive for a logical volume, delete the other partitions in this logical drive until there remains one partition only with the entire capacity. Note that deleting the partition of the logical drive will also destroy all data. Data should be backed up before making partition configurations.

When a logical drive is used as a member of a logical volume, this logical drive can no longer be partitioned in "View and Edit Logical Drives." Instead, the Logical Volume can be divided into eight partitions in "View and Edit Logical Volume."

The procedure for partitioning a logical volume is the same as that for partitioning a logical drive. After the logical volume has been partitioned, map each partition to a host ID/LUN to make the partitions available as individual drives.

As members of a logical volume, all logical drives will be forced to adopt a consistent write policy. Whenever the write policy of a logical volume is changed, for example, the corresponding setting in its members will also be changed.

RAID Expansion with Logical Volume?

The Logical Volume can also be expanded using the RAID expansion function. The concept of expanding a logical volume is similar to that of expanding a logical drive. To perform RAID expansion on a logical drive, replace each member physical drive with a drive of larger capacity or add a new drive, then perform

logical drive expansion to utilize the newly added capacity. For information about RAID expansion, please refer to ***Chapter 9*** "Advanced Configurations."

To perform RAID expansion on a logical volume, expand each member logical drive, then perform "RAID Expansion" on the logical volume.

Steps to expand a Logical Volume:

1. Expand each member logical drive.
2. Expand the logical volume.
3. Map the newly added capacity (in the form of a new partition) to a host LUN.

IMPORTANT!

- *If a logical unit has already been partitioned, and you wish to expand its capacity, the added capacity will be appended to the last partition. You will not be able to proceed with expansion using a firmware version earlier than 3.27 when the unit already has eight partitions.*

- *Unless you move your data and merge two of the partitions, you will not be allowed to expand your logical volume. This is a precautionary limitation on logical unit expansion.*

Different Configuration Settings Using Logical Volume?

Redundant Controller

Without logical volume - logical drives can be assigned to the primary or the secondary controller. The host I/Os directed to a logical drive will be managed by the controller which owns the logical drive. If a controller fails, the host I/Os originally assigned to the failed controller will be taken over by the existing controller. When the controller fails back (a failed controller is replaced by a new one), logical drives will be returned to the replacement controller in its original configuration.

With logical volume - logical volumes can also be assigned to different controllers. The only difference is logical volumes will be used as base units when shifting control during controller failure.

A Logical Volume with Logical Drives of Different Levels?

Multi-level RAID Systems

RAID (0+1) - this is a standard feature of Infortrend RAID controllers. It brings the benefits of RAID 1 (high availability) and RAID 0 (enhanced I/O performance through striping). Simply choose multiple drives (more than two) to compose a RAID 1 logical drive, and RAID (0+1) will be automatically implemented.

RAID (1+0) or RAID 10 - a logical volume is a multi-level RAID implementation A logical volume is a logical composition which stripes data across several logical drives (the RAID 0 method). A logical volume with several RAID 1 members can be considered as a RAID (1+0), or RAID 10 volume.

RAID (3+0) or RAID 30 - a logical volume with several RAID 3 members. A logical volume with several RAID 3 members can be considered as a RAID (3+0), or RAID 53 as defined in "The *RAID* Book" (from The RAID Advisory Board).

RAID (5+0) or RAID 50 - a logical volume with several RAID 5 members.

Chapter

2

RAID Planning

This chapter summarizes the RAID configuration procedures and provides some useful tools for first-time configuration:

2.1 Considerations	Things you should know before setting up the array
2.2 Configuring the Array	The basic configuration procedure
2.3 Operation Theory	A brief introduction to data bus and system drive mapping
2.4 Tunable Parameters	A useful tool that gives you a glimpse of important parameters

2.1 Considerations

After you understand the basic ideas behind RAID levels, you may still be wondering how to begin. Here are the answers to some questions that may help you through the decision making process.

1. How many physical drives do you have?

When initially creating the drive groups, you should know how many drives you have in your RAID system and in the JBOD attached to the RAID controlling unit.

2. Limitations? Does my OS or HBA support 64-bit LBA?

Firmware 3.31 and above support 64-bit LBA addressing. A maximum of 64TB capacity can be included in a single logical drive.

If you have a picture of large arrays in mind, make sure both your host OS and HBA support 64-bit LBA. Not all OSes and HBA cards support 64-bit LBA. Some OSes, such as SuSE Linux 9, Windows 2003, Linux kernel 2.6X, support the addressability. Your HBA must also support this feature. Please refer to the documentation that came with your OS and HBA card to ensure the support for this feature.

Up to 128 members can be included in each logical drive.

Extreme array sizes can cause operational problems with system backup and should be avoided.

3. How many drives on each drive channel?

Optimal system planning is always a compromise between pros and cons. As a general rule, the number of drives you should connect on each channel equals the data bus bandwidth divided by the maximum transfer rate you can get from each of your hard drives. Knowing the mechanical performance of your hard drives can also help to determine how many drives to connect over a drive channel.

Always use fast and large drives of the same capacity for your disk array. A logical drive composed of an adequate number of larger drives can be more efficient than a logical drive comprised of many but smaller drives.

4. How many drives would you like to appear to the host computer?

You must decide what capacity will be included in a logical configuration of drives, be it a logical drive or a logical volume. A logical configuration of drives will appear to the host as a single capacity volume.

You may compose a large logical volume consisting of drives on different drive channels/enclosures, and have it divided into smaller partitions. Each partition will appear as an independent capacity volume. In a performance-oriented configuration, you may configure the same number of drives into several RAID 0 logical drives to maximize the array performance.

5. What kind of host application?

The frequency of read/write activities can vary from one host application to another. The application can be a SQL server,

Oracle server, Informix, or other database server of a transaction-based nature. Applications like video playback and video post-production editing require read/write activities of larger files coming in a sequential order.

Choose an appropriate RAID level for what is the most important for a given application – capacity, availability, or performance. Before creating your RAID, you need to choose an optimization scheme and optimize each array/controller for your application. Stripe size and write policy can be adjusted on a per logical drive basis.

6. Dual loop, hub, or switch?

Unpredictable situations like a cable coming loose can cause system down time. Fibre Channel dual loop or redundant data paths using a flexible LUN mapping method can guarantee there is no single point of failure. The use of a Fibre Channel switch makes cabling and topology more flexible. Change the channel mode, connection type, and other associated settings to adjust the RAID array to your demands.

7. Optimization mode

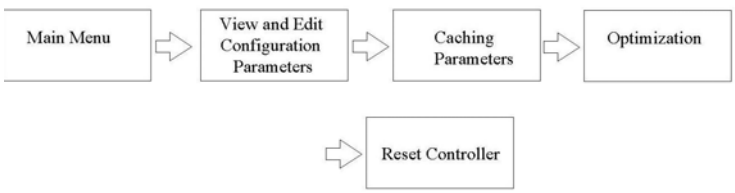

Figure 2 - 1 Optimization Setting

You should select an optimization scheme best suited to your applications before configuring a RAID array. Once the optimization mode is selected, it will be applied to all arrays in the system.

Two options are available: Sequential I/Os and Random I/Os. You may refer to the "Caching Parameters" section in ***Chapter 5*** and ***Chapter 7*** for the stripe size variables and their relation to RAID levels.

Numerous controller parameters are tuned for each optimization mode. Although stripe size can be adjusted on a per logical drive basis, users are not encouraged to make a change to the default values.

For example, smaller stripe sizes are ideal for I/Os that are transaction-based and randomly accessed. However, using the wrong stripe size can cause problems. When an array of the

4KB stripe size receives files of 128KB size, each drive will have to write many more times to store data fragments of the size of 4KB.

Unlike the previous firmware versions, controller optimization mode can be changed without changing the array stripe size.

The default values in optimization modes guarantee the optimal performance for most applications. Consult Table 2-2 for all the controller parameters that are related to system performance and fault- tolerance.

8. What RAID level?

Different RAID levels provide varying levels of performance and fault tolerance.

Table 2 - 1 RAID Levels

RAID Level	Description	Capacity	Data Availability
NRAID	Non-RAID	N	N/A
RAID 0	Disk Striping	N	==NRAID
RAID 1 (0+1)	Mirroring Plus Striping (if N>1)	N/2	>>NRAID ==RAID 5
RAID 3	Striping with Parity on dedicated disk	N-1	>>NRAID ==RAID 5
RAID 5	Striping with interspersed parity	N-1	>>NRAID ==RAID 5
Logical Volume	Striping one or more logical drives of different RAID levels	*	Higher; depends on its members

RAID Level	Performance Sequential	Performance Random
NRAID	Drive	Drive
RAID 0	R: Highest W: Highest	R: High W: Highest
RAID 1 (0+1)	R: High W: Medium	R: Medium W: Low
RAID 3	R: High W: Medium	R: Medium W: Low
RAID 5	R: High W: Medium	R: High W: Low
Logical Volume	Depends on its members; see above	Depends on its members

9. Any spare drives?

(Swap Drive Rebuild / Spare Drive Rebuild)
Spare drives allow for the unattended rebuilding of a failed drive, heightening the degree of fault tolerance. If there is no spare drive, data rebuild has to be manually initiated by replacing a failed drive with a healthy one.

As is often ignored, a spare drive (whether dedicated or global) must have a capacity no smaller than the members of a logical drive.

2.2 Configuring the Array

2.2.1 Starting a RAID System

Here is a flowchart illustrating basic steps to be taken when configuring a RAID system. Hardware installation should be completed before any configuration takes place.

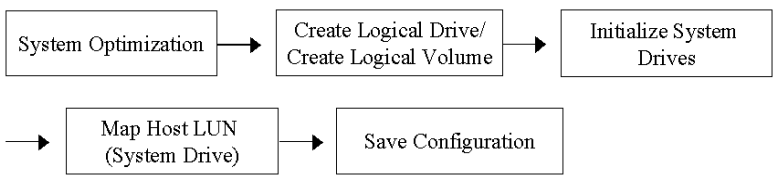

Figure 2 - 2 Array Configuration Process

Drives must be configured and the controller properly initialized before a host computer can access the storage capacity.

1. Use the LCD panel, terminal program, or the RAIDWatch manager to start configuring your array.

2. When powered on, the controller scans all the hard drives that are connected through the drive channels. If a hard drive is connected after the controller completes initialization, use the "Scan SCSI Drive" function (only available on SCSI-based controllers or subsystems) to let the controller recognize its presence.

3. Optimize the controller's parameters for your applications.

4. Configure one or more logical drives to contain your hard drives based on the desired RAID level, and/or partition the logical drive or logical volume into one or several partitions.

NOTE:

- *A "logical drive" is a set of drives grouped together to operate under a given RAID level and it appears as a single contiguous volume. The controller is capable of grouping drives into as many as 128 logical drives, configured in the same or different RAID levels.*

- *A total of 32 "logical volumes" can be created each from one or several logical drives. A logical drive or logical volume can be divided into a maximum of 64 "Partitions."*

5. The next step is to make logical drives or storage partitions available through the host ports. When associated with a host ID or LUN number, each capacity volume appears as one system drive. The host SCSI or Fibre adapter will recognize the system drives after the host bus is re-initialized.

6. The last step is to save your configuration profile in the host system drive (the use of RAIDWatch manager is necessary for this purpose) or to the logical drives you created.

The controller is independent from the host operating system. The host operating system will not be able to tell whether the attached storage is a physical hard drive or the virtual system drives created by the RAID controller.

2.3 Operation Theory

2.3.1 I/O Channel, SCSI ID, and LUN

Depending on the interface used by a RAID system, a SCSI drive channel (SCSI bus) can connect up to 15 drives (excluding the RAID controller itself). A Fibre Channel has up to 125 drives in a loop. Each device occupies one unique ID.

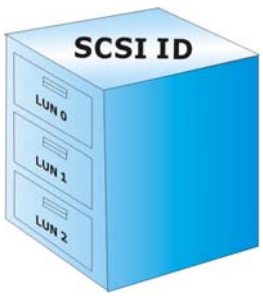

The figure on the left illustrates the idea of mapping a system drive to host ID/LUN combinations. The host ID is like a cabinet, and the drawers are the LUNs (LUN is short for Logical Unit Number). Each cabinet (host ID) contains up to 32 drawers (LUNs). Data can be made available through one of the LUNs of a host ID. Most host adapters treat a LUN like another device.

Figure 2 - 3 SCSI ID/LUNs

2.3.2 Grouping Drives into an Array

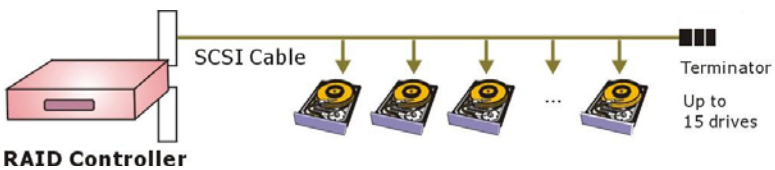

Figure 2 - 4 Connecting Drives

The physical connection of a RAID controller should be similar to the one shown above. Drives are connected through I/O paths that have been designated as drive channels.

The next diagram shows two logical configurations of drives and the physical locations of its members. Using drives from different channels can lower the risk of fatal failure if one of the drive channels fails. There is no limitation on the locations of spares.

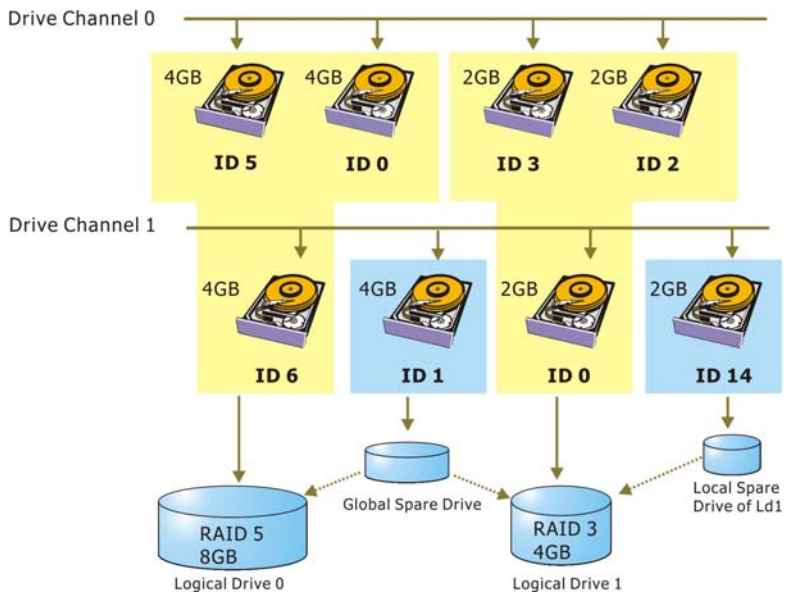

Figure 2 - 5 Physical Locations of Drive Members

A drive can be assigned as the Local Spare Drive that serves one specific logical drive, or as a Global Spare Drive that participates in the rebuild of any logical drive. Spares automatically joins a logical drive when a drive fails. Spares are not applicable to logical drives that have no data redundancy (NRAID and RAID 0).

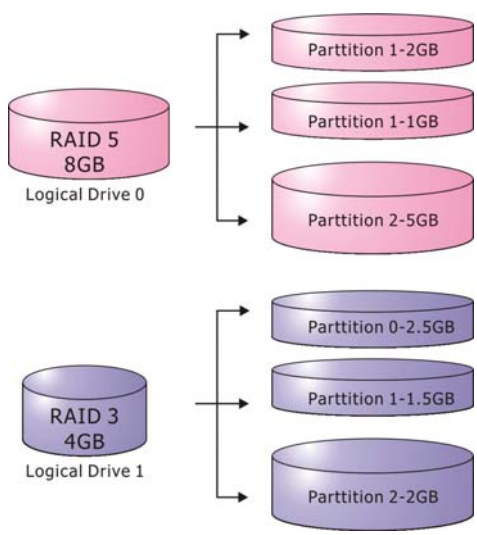

Figure 2 - 6 Partitions in Logical Configurations

You may divide a logical drive or logical volume into partitions of desired capacity, or use the entire capacity as a single volume.

1. It is not a requirement to partition any logical configuration. Partitioning helps to manage a massive capacity.

2. Note that a logical drive can not be included in a logical volume if it has already been partitioned.

2.3.3 Making Arrays Available to Hosts

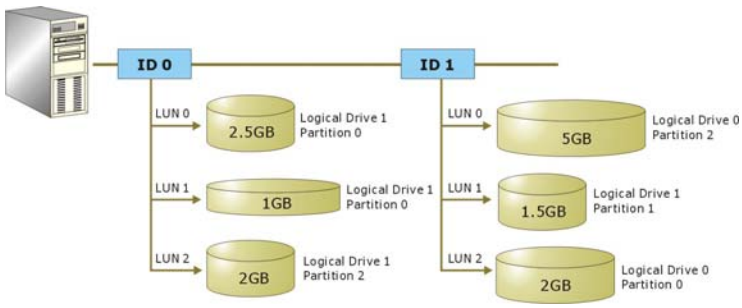

Figure 2 - 7 Mapping Partitions to Host ID/LUNs

Host ID mapping is a process that associates a logical configuration of drives with a host channel ID/LUN. To avail logical partitions on host channel(s), map each partition to a host ID or one of the LUNs under host IDs. Each ID or LUN will appear to the host adapter as one virtual hard drive.

There are alternatives in mapping for different purposes:

1. Mapping a logical configuration to IDs/LUNs on different host channels allows two host computers to access the same array. This method is applicable when the array is shared in a clustering backup.

2. Mapping partitions of an array to IDs/LUNs across separate host channels can distribute workload over multiple data paths.

3. Mapping across separate host channels also helps to make use of all bandwidth in a multi-path configuration. Firmware automatically manages the process when one data path fails and the workload on the failed data path has to be shifted to the existing data paths.

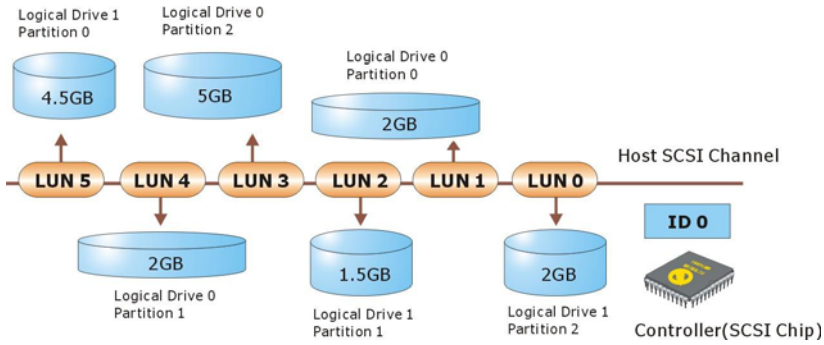

Figure 2 - 8 Mapping Partitions to LUNs under ID

2.4 Tunable Parameters

Fine-tune the controller and the array parameters for your host applications. Although the factory defaults guarantee the optimized controller operation, you may refer to the table below to facilitate tuning of your array. Some of the performance and fault-tolerance settings may also be changed later during the preparation process of your disk array.

Use this table as a checklist and make sure you have each item set to an appropriate value.

Table 2 - 2 Controller Parameter Settings

(1)	Parameters that should be configured at the initial stage of system configuration
(2)	Parameters that can be changed later
(3)	Non-critical

User-Defined Parameters	Default	Alternate Settings
Fault Management:		
(1) Automatic Logical Drive Rebuild - Spare Drive	Enabled when Spare Drive is available	RAID 1 + Local Spare RAID 3 + Local Spare RAID 5 + Local Spare Global Spare
(1) S.M.A.R.T.	Disabled	Detect Only Perpetual Clone Clone + Replace
(3) Clone Failing Drive	Manual function	Replace After Clone Perpetual Clone
(1) Rebuild Priority	Low (higher priority requires more system resources)	Low Normal Improved High
(1) Verification on Write	Disabled	On LD Initialization On LD Rebuild On Normal Drive Writes

(3) SDRAM ECC	Disabled	Enabled
(2) Periodic Cache Flush	Disabled	Continuous to 10 minutes
(1) Event Notification	Reports to user interface and onboard alarm	Over SNMP Traps Over Java-Based Management Software
(2) Periodic Auto-Detect Failure Drive Swap Check Time	Disabled	Disabled, 5~60 seconds
(2) Periodic Drive Check Time	Disabled	Disabled, 0.5~30 seconds
(2) Rebuild Priority	low	Low, normal, improved, high

Controller:		
(1) Channel Mode	*	Host, Drive, RCCOM, Drive + RCCOM (selectable on the EonRAID 2510FS models)
(1) Host and Drive Channel IDs	*	*
(1) Controller Unique Identifier	Preset on some models	Hex number from 0 to FFFFF (FW 3.25 and above)
(2) Data Rate	Auto	Depends on problems solving
(1) Date and Time	N/A	
(1) Time Zone	+ 8 hrs	

Optimization Mode:			
(1) Write-back Cache	Enabled	Disabled	
(1) Array Stripe Size	Related to controller general setting		4KB to 256KB
(2) Adaptive Write Policy	Disabled	Enabled	
(1) Optimization for Random/Sequential	Sequential	Either	
(2) Array Write Policy	Related to controller general setting		W/B or W/T

SCSI Parameters:		
(1) Data Transfer Rate	*	Async. To 160.0MHz
(1) Maximum Tag Count	32	1-128
(1) Maximum Queued I/O Count	32	32 to 1024
(2) LUN's per SCSI ID	8	Up to 32
(1) Periodic Drive Check Time	Disabled	Enabled
(1) Periodic SAF-TE and SES Device Check Time	5	Disabled to 60 seconds
(1) Periodic Auto-Detect Failure Drive Swap Check Time	Disabled	5 to 60 seconds
(1) Number of Host-LUN Connection	32	1 to 1024

(1) Tag per Host-LUN Connection	32	1 to 256
(1) Wide Transfer	*	Enabled/Disabled
(1) Parity Check	Disabled	Enabled
Spin-Up Parameters: (available on SCSI- & Fibre-based models)		
(1) Motor Spin-Up	Disabled	Enabled
(1) Reset at Power-UP	Enabled	Disabled
(1) Initial Disk Access Delay	*	None to 75 seconds

Data Integrity:		
(3) Task Scheduler	N/A	Execute on initialization Start time and date Execution period Media scan mode Media scan priority Logical drive selection

Fibre Channel Parameters:		
(1) Fibre Connection Options	*	Loop Only Point-to-Point Only Loop Preferred Point-to-Point Preferred
(1) Fibre Channel Dual-Loop	Enabled	Enabled by cabling connection
(1) Host ID/WWN Name List	*	User configurable
(1) RCC through Fibre Channel	*	Dedicated or sharing drive channel(s)

Array Configuration:		
(1) Disk Reserved Space	256MB	
(1) Guaranteed Latency Loss (AV option) - Max Drive Response Timeout	Disabled	160, 250, or 500ms
(2) Array Assignment	Primary controller	Secondary controller
(1) Array Partitioning	1	Up to 64
(1) Auto-assign Global Spare	disabled	enabled

Enclosure Monitoring:		
(2) Periodic SAF-TE/SES Device Check Time	30 seconds	Disabled, 50ms~60 seconds
(2) Event Triggered Operation	N/A	Controller, fan, PSU, BBU, UPS, and elevated temperature Auto-shutdown: 2 mins~1 hour
(1) Thresholds for Voltage and	CPU temp: 0~90˚C	User-Defined; do not change parameters unless necessary

Voltage and Temperature Self-Monitoring	Board temp: 0~80˚C 3.3V: 2.9~3.6V 5V: 4.5~5.5V 12V: 10.8~13.2V	

Others:		
(3) Password	N/A	User-Defined; Password Validation Timeout: 1 to Always Check Configurable
(3) LCD Display Controller Name	N/A	User-defined
(1) Network Protocol Support	All enabled except SSH	No configuration options

Chapter

3 Accessing the Array: Serial Port, Ethernet, and Access Security

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

3.1 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 ports. Simply use an RS-232C cable to connect between the controller/enclosure's RS-232C port and the PC serial (COM) port.

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 default is set at 38400 baud, 8 bit, 1 stop bit and no parity. Use the COM1 serial port of the controller.

- In most cases, connecting RD, TD, and SG 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 as presented with the "(Hyper) Terminal" icon in the Accessories window.

- For other details on connecting to the serial port, please refer to the ***Installation and Hardware Reference Manual*** that came with your controller.

3.1.1 Configuring RS-232C Connection via LCD Keypad Panel

The subsystem's baud rate default is set to 38400. If necessary, follow the steps below to change the baud rate using the front panel keypad:

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

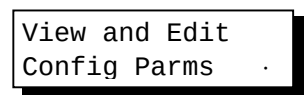

Select "Communication Parameters ..", then press **ENT**.

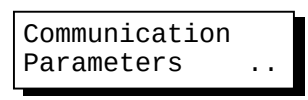

Select "RS-232 Configuration ..", then press **ENT**.

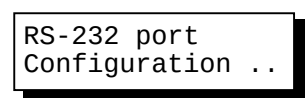

Select "COM1 Configuration ..", then press **ENT**.

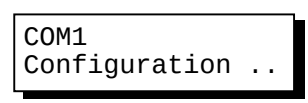

Select "Baud-rate XXXXX..", then press **ENT**.

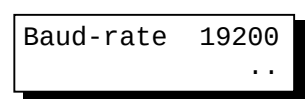

The baud rate default is 38400. If another baud rate is preferred, press the up or down arrow keys to select the baud rate, then press **ENT** for 2 seconds to confirm the selected baud rate. Set the identical baud rate to your RAID array and your terminal computer.

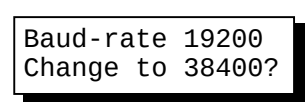

- The following baud rates are available: 2400, 4800, 9600, 19200 and 38400.

- Terminal connection should work properly using the above setting. You may check the following option in your COM port configuration if you encounter problems:

  1. "Comm Route Dir ..": The communication route should be configured as "direct to port" instead of "through PPP."

     NOTE: The communication route option has been removed from later firmware revisions.

  2. "Term Emul. Enab ..": Make sure the terminal function has not been accidentally disabled.

3.1.2 Starting RS-232C Terminal Emulation

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

←→↑↓	To select options
[Enter]	To go to a submenu or to execute a selected option
[Esc]	To escape and go back to the previous menu
[Ctrl] [L]	To refresh the screen information

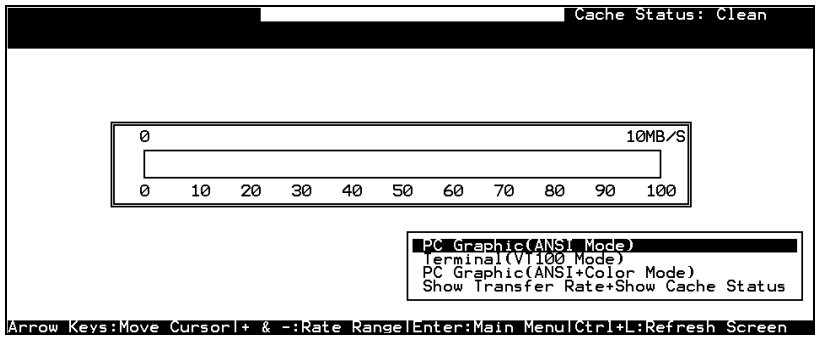

IMPORTANT!

- *If the RS-232C cable is connected while the controller is powered on, press [Ctrl] + [L] to refresh the screen information.*

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.

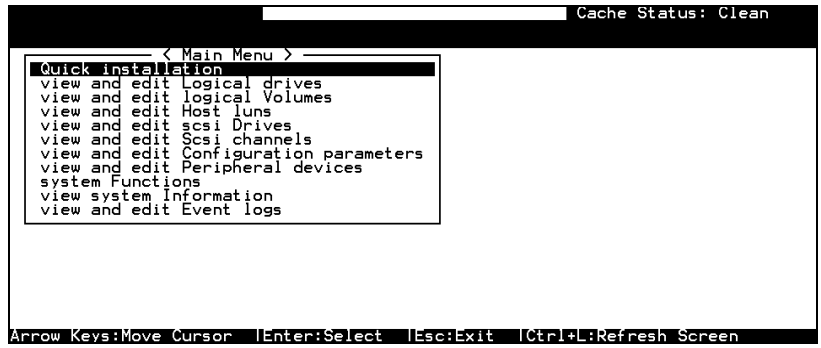

Choose a functional item from the main menu to begin configuring your RAID.

3.2 Communication Parameters

The Communication Parameters is the first functional sub-menu you access once the physical RS-232C link is connected. In addition to the baud rate and terminal emulation option which have been setup in section 3.1, the sub-menu contains several 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>."

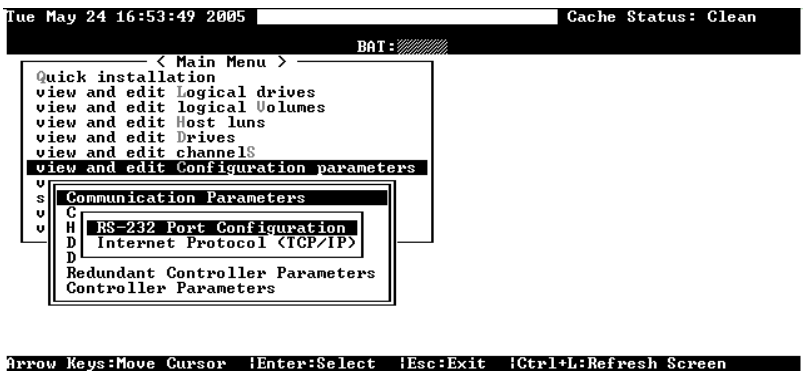

RS-232 Port Configuration

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

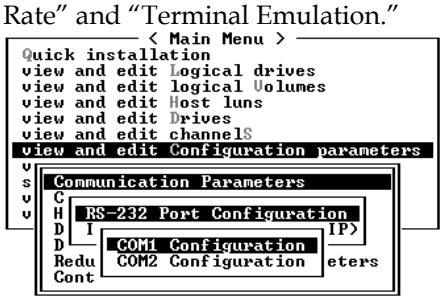

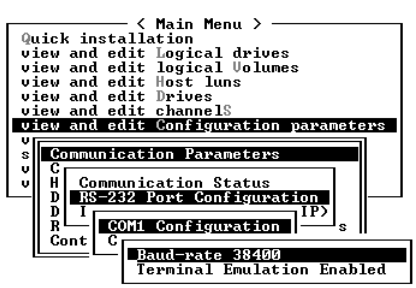

NOTE:

COM2 is mostly used for UPS connection on the EonStor subsystems. The COM2 port on the EonStor can no longer be used for management access.

Terminal Emulation:

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

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.

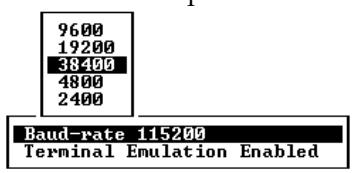

Internet Protocol <TCP/IP>

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

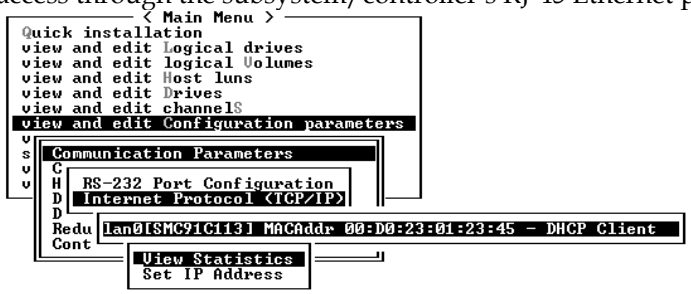

To access the configuration options, press Enter on "Internet Protocol <TCP/IP>" to display the information of Ethernet port. Press Enter on the chip information to display the "View Statistics" and "Set IP Address" options.

View Statistics

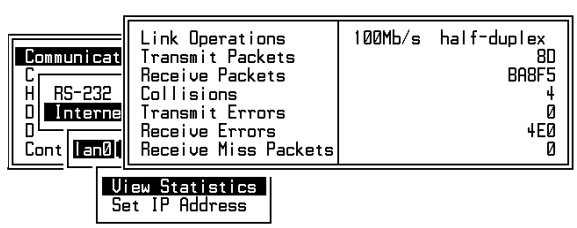

This windows displays the current Ethernet link status.

Set IP Address

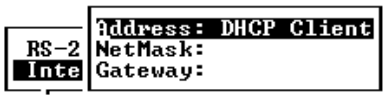

Provide a valid IP address for your controller/subsystem's Ethernet port. Consult your network administrator for an IP address and the associated NetMask and Gateway values. You may also key in "DHCP" if your local network supports automatic IP configuration.

NOTE:

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. RAIDWatch and its sub-module, the Configuration Client, recognize a configured array through its Ethernet port IP.

3.3 Out-of-Band via Ethernet

The RAIDWatch manager software provides a graphical interface to RAID subsystems. RAIDWatch comes with an installer program that facilitates the installation of software modules for local/remote access to the array through standard TCP/IP.

Below is a tree structure of RAIDWatch's installation options:

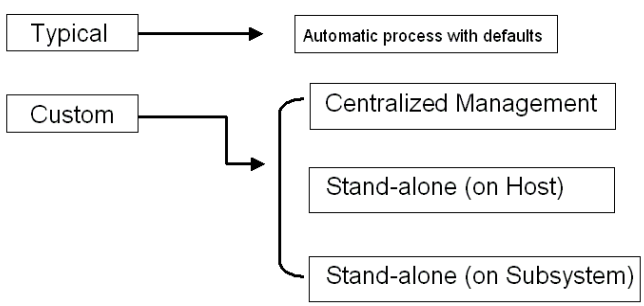

Figure 3 - 1 Installation Options

If you install RAIDWatch using the "***Typical***" option or the first two custom modes "***Centralized Management***" and "***Stand-alone (on Host),***" prepare the following:

1. TCP/IP for the network connection to the controller/subsystem's Ethernet port.
2. Configure a valid IP and appropriate network settings for the RAID controller/subsystem's Ethernet port.

If you install RAIDWatch using the "***Stand-alone (on Subsystem)***" mode, prepare the following:

1. TCP/IP for the network connection to the controller/subsystem's Ethernet port.
2. Configure a valid IP and appropriate network settings for the RAID controller/subsystem's Ethernet port.
3. Array reserved space: If you select the "Custom" installation mode, "Stand-alone (on Subsystem)," the RAIDWatch installation program will automatically create a reserved space on the first disk drive and distribute necessary program files to that space.
4. Unless the in-band connection over SCSI or Fibre host busses is preferred, there is no need to configure the Peripheral Device setting with the later versions of RAIDWatch manager using Ethernet connection.

NOTE:

- *Only the connections via Ethernet are discussed here. For more details about the connection over In-band method (the host-to-subsystem connection), please refer to RAIDWatch's User's Manual.*
- *Unless you want to start using RAIDWatch under the following conditions, it is unnecessary to configure the Peripheral Device settings:*
  1. *Using RAIDWatch to access a brand new array that has no configured drive groups*
  2. *Using the In-band method (over the existing host data paths) for array management. The Peripheral Device settings present the RAID array on the host data paths as if it is an external device.*

  If you use Ethernet link and configure the Peripheral Device settings, the application servers will mistakenly identify a RAID array as an external device and require for a driver. Normal RAID array operation will require no supporting drivers.

What is the "Disk Reserved Space?"

RAIDWatch and Reserved Space:

- If you install RAIDWatch using the third installation mode, "Stand-alone (on Subsystem)," user's configuration data and the manager's main programs will be kept in a small section of disk space on a data drive.

  After you create a logical drive, all its member drives will have a reserved space. The segregated disk space is called a "Disk Reserved Space." When configuring a logical drive, the firmware automatically segregates 256MB of disk space from each of the member drives.

- In the event of single controller failure, the surviving controller can continue the management access to the array because the RAIDWatch manager's main program is kept in the reserved space on the hard drives. Operators' access to the system will not be interrupted.

- The reserved space also allows "Drive Roaming." Arrays (hard drives included in logical configurations) can be readily recognized by another RAID subsystem/RAID controller. This is useful when hardware faults occur, and you have to replace a faulty subsystem/controller.

Other Concerns

Availability Concern

For safety reasons, it is better to create a reserved space on more than one logical drive. The reserved space information can be seen during the array creation process or in the "View and Edit Drives" menu.

Whatever data is put into the reserved space, the firmware will automatically duplicate and distribute it to the reserved section on every data drive. Even if one hard drive or one logical drive fails, an exact replica still resides on other member drives.

Web-based Management

The controller firmware has embedded http server. Once properly configured, the controller/subsystem's Ethernet port behaves like an HTTP server. This applies to RAIDWatch's third installation mode, "Stand-alone (on Subsystem)."

3.3.1 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 controller/subsystem's Ethernet port and an Ethernet port from your local network.

3.3.2 Configuring the Controller

To prepare the controller for Ethernet connection, do the following:

1. Use a Terminal Emulator to Begin Configuration

Connect the subsystem's serial port to a PC running a VT-100 terminal emulation program or a VT-100-compatible terminal.

Make sure the included null modem is already attached to enclosure serial port or the host computer's COM port. The null modem converts the serial signals for connecting to a standard PC serial interface. For more details, please refer to the descriptions above in ***Section 3.1*** and ***Section 3.2***.

2. Assign an IP Address to the Ethernet Port

Assign an IP address to the controller Ethernet port and specify the Net Mask and Gateway values. Power off your system and then power on again for the configuration to take effect.

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."

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

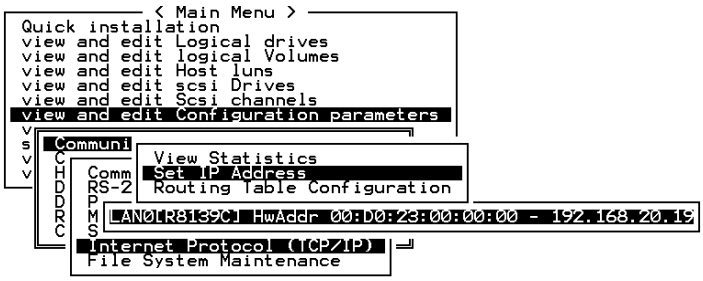

Provide the IP address, NetMask, and Gateway values accordingly.

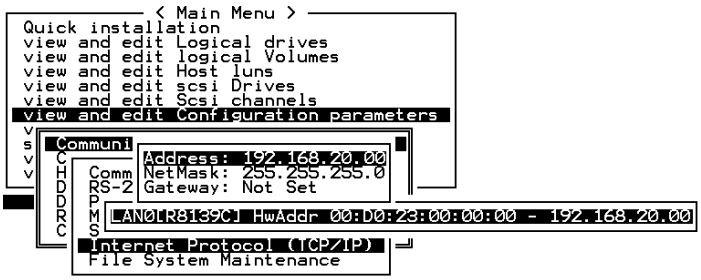

PING the IP address from your management computer to make sure the link is up and running.

3.4 Telnet Connection

1. Use an Ethernet cable with RJ-45 phone jacks to connect the Ethernet port on the subsystem/controller module.
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.
3. Consult your network administrator for an IP address that will be assigned to the controller Ethernet port.
4. Assign an IP address to the controller Ethernet port and specify the NetMask and Gateway values. Power off your system and then power on again for the configuration to take effect.
5. 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."

6. Provide the IP address, NetMask, and Gateway values accordingly.
7. PING the IP address from your management computer to make sure the link is up and running.
8. Open a command prompt and key in "telnet <IP address>" to access the embedded firmware utility.

3.5 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

To access the firmware embedded configuration utility over the SSH network, do the following:

1. Open a web browser.
2. Key in the line below in the browser's URL field.

   http://<controller IP>/ssh.htm

3. Enter controller name (if there is any) as user name and the password set for the RAID subsystem. If neither the controller name or password has been set, press Enter at each command prompt to proceed.

Chapter

4 LCD Screen Messages

4.1 The Initial Screen

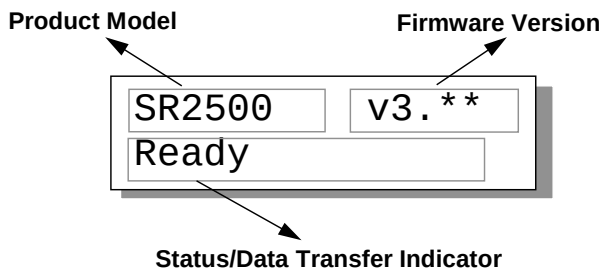

Status/Data Transfer Indicator:

Ready	There is at least one logical drive or logical volume mapped to a host ID/LUN.
No Host LUN	No logical drive created or the logical drive has not yet been mapped to any host ID/LUN.
▮▮▮▮▮▮▮	Indicates data transfer. Each block indicates 256Kbytes of data throughput.

4.2 Quick Installation Screen

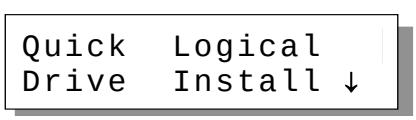

Press **[ENT]** to create a logical drive. The controller/subsystem will start initialization of one logical drive with all the connected SCSI drives and automatically map the logical drive to LUN 0 of the first host channel. The "Quick Installation" can only be performed when there is no logical drive.

4.3 Logical Drive Status

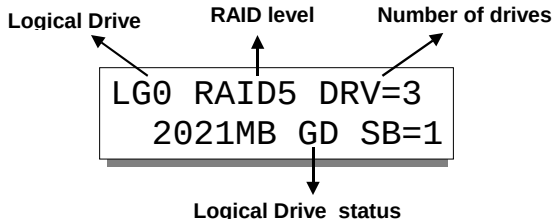

Logical Drive:	The Logical Drive number.
RAID level:	The RAID level used in this logical drive.
Number of Drives:	The number of physical drives included in this configuration.

Logical Drive status:

XxxxMB	The capacity of this logical drive.
SB=x	Standby drives available to this logical drive. Except the spares dedicated to other logical configurations, all spare drive(s) will be counted in this field, including Global and Local Spares.
xxxxMB INITING	The logical drive is now initializing.
xxxxMB INVALID	For firmware version before 3.31: The logical drive has been created with "Optimization for Sequential I/O", but the current setting is "Optimization for Random I/O." -OR- The logical drive has been created with "Optimization for Random I/O," but the current setting is "Optimization for Sequential I/O." Firmware version 3.31 has separate settings for array optimization and array stripe size. This message will not appear when the optimization mode is changed.
xxxxMB GD SB=x	The logical drive is in good condition.
xxxxMB FL SB=x	One drive failed in this logical drive.
xxxxMB RB SB=x	Logical Drive is rebuilding.
xxxxMB DRVMISS	One of the drives is missing.
INCOMPLETE ARRAY	Two or more drives failed in this logical drive.

4.4 Logical Volume Status

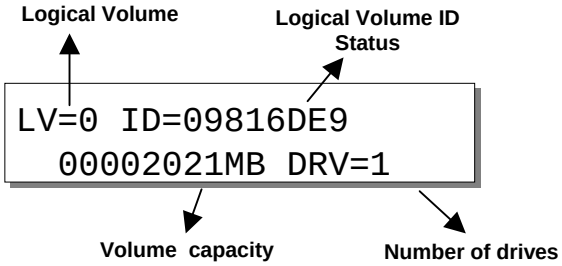

Logical Volume:	The Logical Volume number.
DRV=x:	The number of logical drive(s) contained in this logical volume.
Logical Volume ID Status:	The unique ID number of the logical volume (controller random generated).
Logical Volume Status:	
xxxMB	The capacity of this logical volume.

4.5 SCSI Drive Status

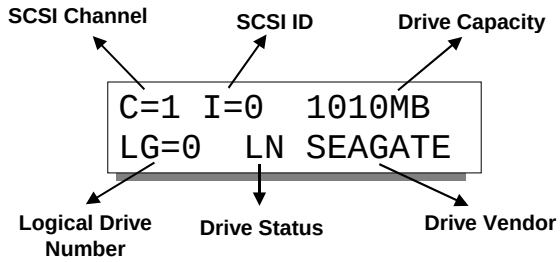

Drive Status:

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
GlobalSB	Global Spare Drive
NEW DRV	New drive
BAD DRV	Failed drive
ABSENT	Drive does not exist; an empty drive tray
MISSING	Drive missing (drive was once there)
SB-MISS	Spare drive missing

4.6 Channel Status

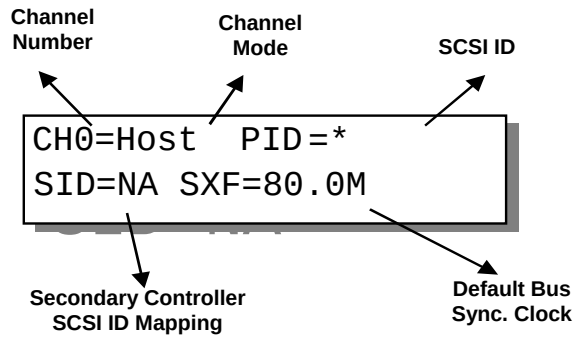

SATA Drive Channel (Single Controller Configuration)

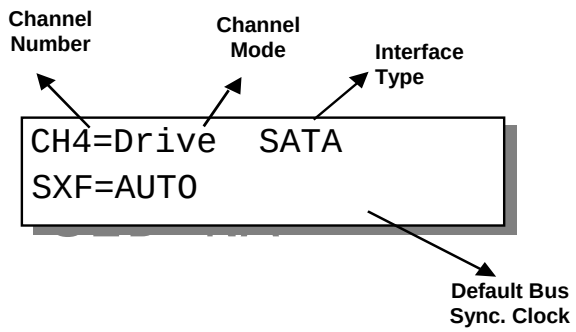

Channel Mode:

Host	Host Channel mode
Drive	Drive Channel mode

Default SCSI Bus Sync Clock:

80.0M	The default setting of this channel is 80.0MHz in Synchronous mode
Async	The default setting of this SCSI channel is in Asynchronous mode

Primary Controller SCSI ID Mapping:

*	Multiple SCSI ID's applied (Host Channel mode only)
(ID number)	Primary Controller is using this SCSI ID for host LUN mapping.
NA	No SCSI ID applied (Drive Channel mode only)

Secondary Controller SCSI ID Mapping:

*	Multiple SCSI ID's applied (Host Channel mode only)
(ID number)	Secondary Controller is using this SCSI ID for host LUN mapping.
NA	No SCSI ID applied (Drive Channel mode only)

4.7 Controller Voltage and Temperature

Press **ENT** for two seconds to enter the Main Menu. Press τ or σ 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** and then choose "View Ctlr Periph Device Status..", then press **ENT**.

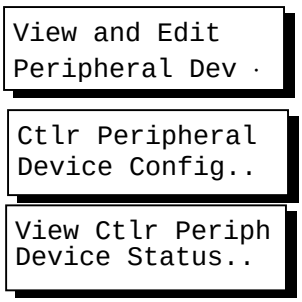

Press the up or down arrow keys to choose either "Voltage Monitor", or "Temperature Monitor".

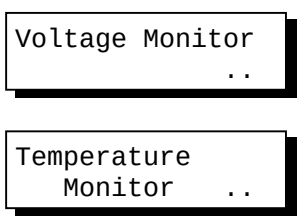

Select "Temperature and Voltage Monitor" by pressing **Enter**. Press the up or down arrow keys to browse through the various voltage and temperature statuses.

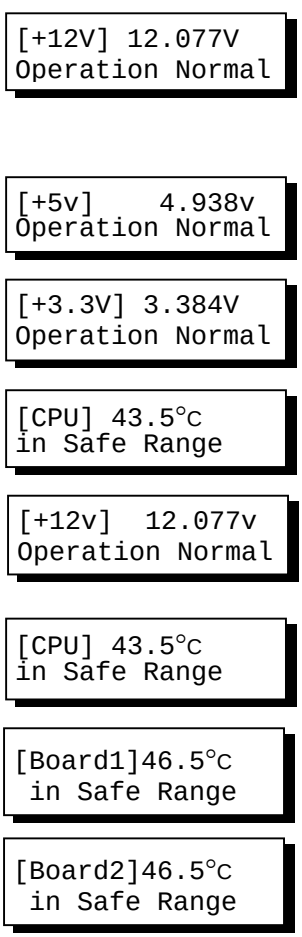

4.9 View and Edit Event Logs

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 delete a specified item and all events prior to this event, press **ENT** for 2 seconds.

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

IMPORTANT!

- *The event log will be cleared after the controller is powered off or reset.*

Chapter

5 LCD Keypad Operation

A navigation roadmap for the configuration menu options of LCD keypad is separately available as a PDF file. You may check your Product Utility CD or visit Infortrend's FTP site for the latest update.

5.1 Power On the 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 enclosure.

NOTE:

- *The RAID controller units of the EonStor subsystems come in separate packages. Chances are that a secondary controller may be installed to the controller slot at the upper position. As the result, your subsystem's LCD panel provides access to the Secondary controller. However, in the dual-redundant 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 "P" or "S." If the LCD displays "S," press the Up and Down arrow keys for one second to switch around the access to different controllers.

5.2 Caching Parameters

Optimization Modes

Mass storage applications can be divided into two categories according to its read/write characteristics: database and video/imaging. To optimize the controller for these two categories, the controller has two embedded optimization modes with controller behaviors adjusted to different read/write parameters. They are the Optimization for Random I/O and the Optimization for Sequential I/O.

Limitations: There are limitations on the use of optimization modes.

1. You can select the stripe size of each array (logical drive) during the initial configuration. However, changing stripe size is only recommended for experienced engineers who have tested the effects of tuning stripe sizes for different applications.

2. The array stripe size can only be changed during the initial configuration process.

3. Once the controller optimization mode is applied, access to different logical drives in a RAID system will follow the same optimized pattern. You can change the optimization mode later without having to re-organize your array.

Database and Transaction-based Applications:

These kinds of applications usually include SQL server, Oracle server, Informix, or other database services. These applications keep the size of each transaction down to the minimum, so that I/Os can be rapidly processed. Due to their transaction-based nature, these applications do not read or write a bunch of data in sequential order - access to data occurs randomly. The transaction size usually ranges from 2K to 4K. Transaction performance is measured in "I/Os per second" or "IOPS."

Video Recording/Playback and Imaging Applications:

These kinds of applications usually includes video playback, video post-production editing, or other similar applications. These applications have the tendency to read or write large files from and into storage in sequential order. The size of each I/O can be 128K, 256K, 512K, or up to 1MB. The efficiency of these applications is measured in "MB/sec."

When an array works with applications such as video or image oriented applications, the application reads/writes from the drive as large-block, sequential threads instead of small and randomly accessed files.

The controller optimization modes have read-ahead buffer and other R/W characteristics tuned to obtain the best performance for these two major application categories.

Optimization Mode and Stripe Size

Each controller optimization mode has preset values for the stripe size of arrays created in different RAID levels. If you want a different value for your array, you may change the controller optimization mode, reset the controller, and then go back to create the array. Once the array is created, stripe size cannot be changed.

Using the default value should be sufficient for most applications.

	Opt. For Sequential I/O	Opt. for Random I/O
RAID0	128	32
RAID1	128	32
RAID3	16	4
RAID5	128	32
NRAID	128	32

Optimization for Random or Sequential I/O

Select from Main Menu "View and Edit Config Parms," "Caching Parameters," and press **ENT.** Choose "Optimization for Random I/O" or "Optimization for Sequential I/O," then press **ENT** for two seconds to confirm. Press **ESC** to leave and the setting will take effect after the controller is restarted.

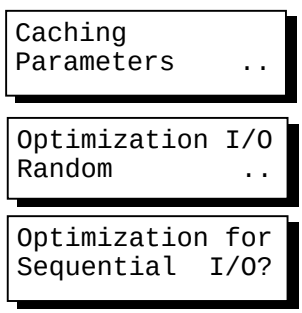

IMPORTANT!

- *The original 512GB threshold on array optimization mode is canceled. If the size of an array is larger than 16TB, only the optimization for sequential I/O can be applied. Logical drives of this size are not practical; therefore, there is actually no limitation on the optimization mode and array capacity.*

Write-back/Write-through Cache Enable/Disable

As one of the submenus in "Caching Parameters," this option controls the cached write function. Press **ENT** to enable or disable "Write-back Cache." Press **ENT** for two seconds to confirm.

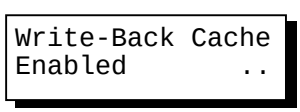

The current status will be displayed on the LCD.

The Write-through mode is safer if your controller is not configured in a redundant pair and there is no battery backup.

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

Write-back caching can dramatically improve write performance by caching the unfinished writes in memory and letting them be committed to drives latter in a more efficient manner. In the event of power failure, a battery module can hold cached data for days. In the event of controller failure, data cached in the failed controller has an exact replica on its counterpart controller and therefore remains intact.

Periodic Cache Flush

If for a reason Write-Back caching is preferred for better performance and yet data integrity is of the concern, e.g., lack of the battery protection, the system can be configured to flush the cached writes at every preset interval.

Note that the "Continuous Sync" option holds data in cache for as long as necessary to complete a write operation and immediately commits it to hard drives if it does not come in a series of sequential write requests.

Select from Main Menu "View and Edit Config Parms," "Caching Parameters," and press **ENT.** Use the arrow keys to scroll through the options and select "Periodic CachFlush Time", and then press **ENT** to proceed. The "Set Cache Flush Time – Disable" appears. The default is "Disable." Use your arrow keys to select an option from "ConSync," "1/2min," to "10min." "ConSync" stands for "continuously synchronized." Press **ENT** to select and press **ESC** to leave and the setting will take effect immediately.

```
Period CachFlush
Time - Disable
```

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

IMPORTANT!

- *Every time you change the caching parameters, you must reset the controller for the changes to take effect.*
- *In the redundant controller configuration, write-back will only be applicable when there is a synchronized cache channel between partner*

controllers.

5.3 View Connected Drives:

A RAID system consists of many physical drives that can be modified and configured as the members of one or several logical drives.

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 SCSI Drives," then press **ENT**.

```
View and Edit
Drives
```

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

```
C=2 I=0 1010MB
New DRV SEAGATE
```

```
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.

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   ..
```

The Revision Number of the selected SCSI drive will be shown. Press τ to see other information.

```
Revision Number:
0274
```

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

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 a 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 Local/Global Spare Drive to a logical drive whose*

members have a block number equal to or smaller than the Local/Global Spare Drive but you should not do the reverse.

5.4 Creating a Logical Drive

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**.

```
View and Edit
Logical Drives
```

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

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

Choosing a RAID Level

Press the up or down arrow keys to choose the desired RAID level, then press **ENT** for two seconds. "TDRV" (Total Drives) refers to the number of available SCSI drives.

```
TDRV=4 Create
LG Level=RAID5 ?
```

Choosing Member Drives

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

```
RAID X  Selected
To Select drives
```

Press **ENT**, then use the up or down arrow keys to browse through the available drives. Press **ENT** again to select/deselect the 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. "C=1 I=0" refers to "Channel 1, SCSI ID 0".

```
C=1 I=0   1010MB
NEW DRV  SEAGATE
```

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

After all the desired drives have been selected, press **ENT** for two seconds to continue. Press the up or down arrow keys to choose "Create Logical Drive," then press **ENT** for two seconds to start initializing the logical drive.

```
Create Logical
Drive          ?
```

Logical Drive Preferences

You may also choose "Change Logical Drive Parameter," then press **ENT** to change related parameters before initializing the logical drive.

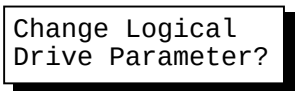

Maximum Drive Capacity

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

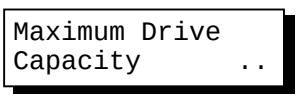

Use the up and down arrow keys to change the maximum size that will be used on each drive.

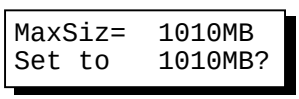

Spare Drive Assignments

The Local Spare Drive can also be assigned here. Press the up or down arrow keys to choose "Spare Drive Assignments," then press **ENT**.

Available 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 Spare Drive. Press **ENT** again for two seconds.

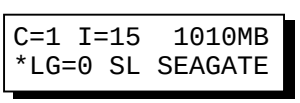

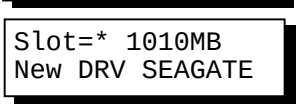

Disk Reserved Space

This menu allows you to change the size of disk reserved space. Default is 256MB. We recommended using the default value.

Choices are 256MB and 64KB. With 64KB, logical drives are backward compatible to RAID controllers running earlier firmware versions. Press **ENT** and use the up or down keys to choose the size you prefer.

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 controller's caching mode setting. Other choices are "Write-back" and "Write-through."

```
Write Policy
Default      ..
```

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. Users may proceed with array configuration, e.g., including this array in a logical volume.

```
Initialization
Mode   Online..
```

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

Stripe Size

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 be sure of the performance gains it might bring you.

```
Stripe size
Default      ?
```

Listed below are the default values for an array. The default value for the stripe size is determined by the controller Optimization Mode and the RAID level chosen for an array.

Table 5 - 1 RAID Level, Optimization Modes, and Stripe Size

	Opt. for Sequential I/O	Opt. for Random I/O
RAID0	128	32
RAID1	128	32
RAID3	16	4
RAID5	128	32
NRAID	128	32

When you are done setting logical drive preferences, press **ESC** and use your arrow keys to select "Create Logical Drive?". Press **ENT** for two seconds to proceed.

Beginning Initialization

Press **ESC** to return to the previous menu. Use the up or down arrow keys to choose "Create Logical Drive," then press **ENT** for two seconds to start initializing the logical drive.

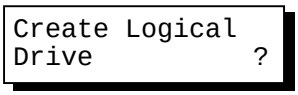

The Online Mode:

If the online initialization method is applied, the array will be available for use immediately. The array initialization runs in the background while data can be written onto it and users can continue configuring the RAID system.

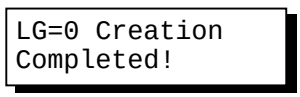

The Offline Mode:

The controller will start to initialize the array parity if using the "offline" mode. Note that if NRAID or RAID 0 is selected, initialization time is short and completes almost immediately.

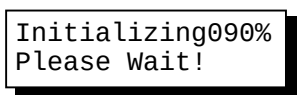

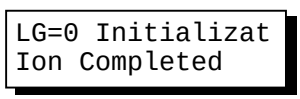

The logical drive's information displays when the initialization process is completed. If the "online" mode is adopted, array information will be displayed immediately.

5.5 Creating a Logical Volume

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 .
```

Press the up or down arrow keys to select an undefined entry for logical volume, then press **ENT** for two seconds to proceed. "LV" is short for Logical Volume.

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

Proceed to select one or more logical drives. Press **ENT** to proceed. "LD" is short for Logical Drive.

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

Use the up or down arrow keys to browse through the logical drives. Press **ENT** again to select/deselect the drives. An asterisk (*) mark will appear when the logical drive is selected. After all the desired logical drive(s) have been selected, press **ENT** for two seconds to continue.

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

Two submenus will appear.

Initialization Mode

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

```
Initialization
Mode    Online..
```

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 caching mode setting. Other choices are "Write-back" and "Write-through."

```
Write Policy
Default       ..
```

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 ?
```

The logical volume has been successfully created.

```
Lv=0 Creation
Completed
```

Press ESC to clear the message. Another message will prompt; press ESC to clear it.

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

Logical volume information will be displayed below.

```
Create Logical
Volume Successed
```

Logical Volume Assignment

If you have two controllers, you may choose to assign this logical volume to the secondary controller. The assignment can be done during or after the initial configuration.

```
Change Logical
Volume Params  ?
```

```
Logical Volume
Assignments   ..
```

If the redundant controller function has been enabled, and the secondary controller IDs have been assigned to I/O channels, the assignment menus should appear as listed on the right.

```
Red Ctlr Assign
to Sec. Ctlr   ?
```

If settings related to redundant controllers have not been completed, you may set them after the volume is successfully created.

```
Logical Volume
Assignment    ..
```

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.

```
Red Ctlr Assign
to Sec. Ctlr   ?
```

Press ESC, and the LCD will display the logical volume’s information when initialization is finished.

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

5.6 Partitioning a Logical Drive/Logical Volume

Partitioning and the creation of a logical volume, are not requirements for creating a RAID system. The configuration processes for partitioning a logical drive are the same as those for partitioning a logical volume.

IMPORTANT!

- *If operated 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.*

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 .
```

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

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

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

```
Partition
Logical Volume..
```

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=1:
    2021MB     ?
```

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

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

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

The rest of the drive space will be automatically allocated as the last partition. You may go on to create up to 32 partitions using the same method as described above.

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

Press **ESC** several times to go back to the Main Menu.

5.7 Mapping a Logical Volume/Logical Drive to Host LUN

The process of mapping a logical drive is identical to that of mapping a logical volume. The process of mapping a logical volume is used as an example.

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**.

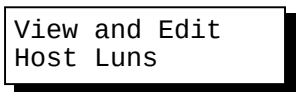

Note some details before proceeding:

1. **Primary/Secondary IDs**: A logical group of drives (logical drive/logical volume) previously assigned to the primary controller cannot be mapped to a secondary ID. Neither can those be assigned to the secondary controller be mapped to a primary ID.

2. **Reserved IDs**: For a SCSI-based controller, ID 7 is reserved for the controller itself. If there are two controllers, controllers might occupy ID 6 and ID 7. Please check your system *hardware manual* for details on preserved IDs.

Press the up or down arrow keys to select a configured host ID, and then press **ENT** for two seconds to confirm. IDs are available as Primary or Secondary Controller IDs.

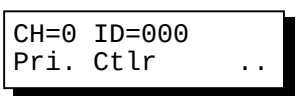

Press the up or down arrow keys to select the type of logical configuration. Available choices are “Map to Logical Volume,” “Map to Logical Drive,” or “Map to Physical Drive.” Confirm your choice by pressing **ENT.**

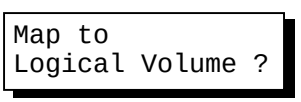

Press the up or down arrow keys to select a LUN number, then press **ENT** to proceed.

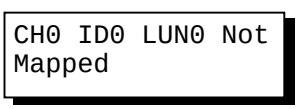

Press **ENT** for two seconds to confirm the selected LUN mapping.

Press the up or down arrow keys to select a partition from the logical volume. Press **ENT** for two seconds to

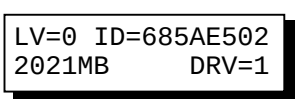

map the selected partition to this LUN. If the logical configuration has not been partitioned, you can map the whole capacity to a host LUN.

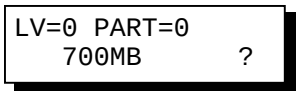

Mapping information will be displayed on the subsequent screen. Press **ENT** for two seconds to confirm the LUN mapping.

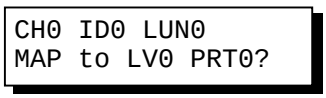

With any of the host ID/LUN successfully associated with a logical capacity, the "No Host LUN" message in the main menu will change to "Ready."

If you want to create more host IDs, please move to ***Section 5.12 Viewing and Editing SCSI*** Channels for more details on channel mode and channel IDs setting.

5.8 Assigning Spare Drive and Rebuild Settings

Adding a Local Spare Drive

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

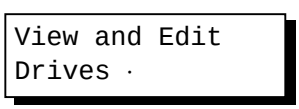

SCSI drive information will be displayed on the LCD. Press the up or down arrow keys to select a drive that is stated as "NEW DRV" or "USED DRV" that has not been assigned to any logical drive, nor as spare drive or failed drive, then press **ENT** to select it.

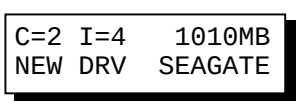

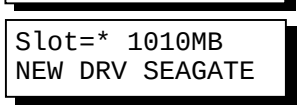

Press the up or down arrow keys to select "Add Local Spare Drive," then press **ENT**.

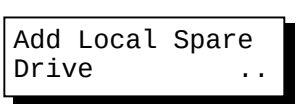

Press the up or down arrow keys to select the logical drive where the Local Spare Drive will be assigned, then press **ENT** for two seconds to confirm.

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

Adding a Global Spare Drive

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

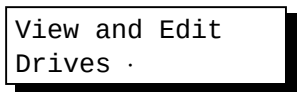

SCSI drive information will be displayed on the LCD. Press the up or down arrow keys to select a SCSI drive that has not been assigned to any logical drive yet, then press **ENT**.

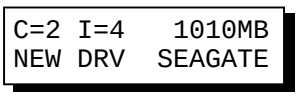

Press the up or down arrow keys to select "Add Global Spare Drive," then press **ENT**.

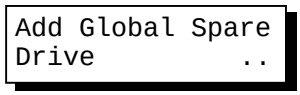

Press **ENT** again for two seconds to add the spare drive. The message "Add Global Spare Drive Successful" will be displayed on the LCD.

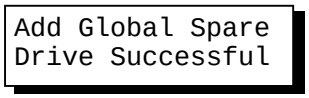

Rebuild Settings

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

Press the up or down arrow keys to select "Disk Array Parameters," then press **ENT**.

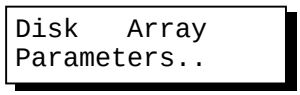

Press the up or down arrow keys to select "Rebuild Priority Low," then press **ENT**. "Low" refers to the temporary setting.

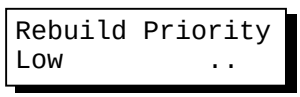

Press **ENT** again and the abbreviation mark ".." will change to a question mark "?". Press the up or down arrow keys to select priority "Low," "Normal," "Improved," or "High".

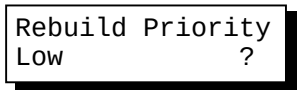

Press **ENT** to confirm and the question mark "?" will turn into "..".

```
Rebuild Priority
High           ..
```

NOTE:

- *The rebuild priority determines how much of the system resources are used when rebuilding a logical drive. The default setting of the rebuild priority is "LOW." Rebuild will have less impact on host I/O access, but will take a longer time to complete. Changing the priority to a higher level will achieve a faster rebuild, but will significantly increase the host I/O response time. The default setting "LOW" is recommended.*

5.9 Viewing and Editing Logical Drives and Drive Members

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 ·
```

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

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

Press the up or down arrow keys to select "View SCSI Drives..", then press **ENT** .

```
View Drives
               ..
```

Press the up or down arrow keys to scroll through the list of member drives.

```
C=1 I=0   1010MB
 LG=0 LN SEAGATE
```

```
Slot=* 1010MB
LG=0 LN SEAGATE
```

Deleting a Logical Drive

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 ·
```

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

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

```
Delete Logical
Drive          ..
```

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

Press **ENT** for two seconds to delete. The selected logical drive has now been deleted.

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

Deleting a Partition of a Logical Drive

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

```
View and Edit
Logical Drives .
```

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

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

```
Partition
Logical Drive ..
```

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       ?
```

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 rest of the drive space will be automatically allocated to the last partition as diagrammed below.

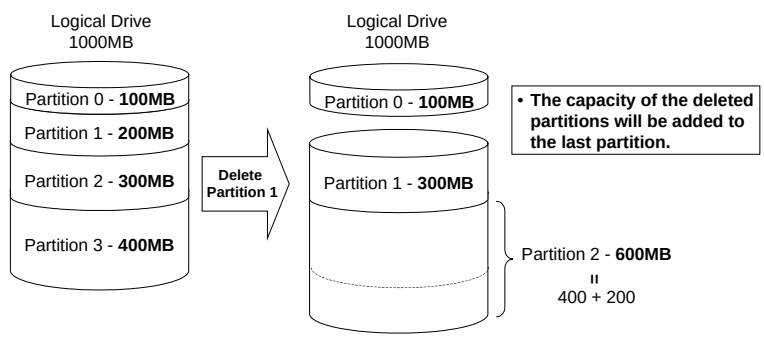

Figure 5 - 1 Drive Space Allocated to the Last Partition

WARNING!

- *Whenever there is a partition change, data will be erased, and all host LUN mappings will be removed. Therefore, every time the size of a partition has been changed, it is necessary to reconfigure all host LUN mappings of the associated partitions.*

Assigning a Name to a Logical Drive

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 .
```

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

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

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

```
Logical Drive
Name          ..
```

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 15.

```
Enter LD Name:
_
```

Rebuilding a Logical Drive

If you want the controller to auto-detect a replacement drive, make sure you have the following items set to enabled:

1. Periodic Drive Check Time
2. Periodic Auto-Detect Failure Drive Swap Check Time

These two configuration options can be found under "View and Edit Configuration Parameters" -> " Drive-Side SCSI Parameters".

Press **ENT** for two seconds to enter the Main Menu. Press the up or down arrow

```
View and Edit
Logical Drives .
```

keys to select "View and Edit Logical Drives..", then press **ENT**.

Press the up or down arrow keys to select the logical drive that has a failed member, then press **ENT**.

```
LG0 RAID5 DRV=3
  2012MB FL SB=0
```

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

```
Rebuild Logical
Drive          ..
```

Press **ENT** for two seconds to start rebuilding the logical drive.

```
Rebuild Logical
Drive          ?
```

The rebuilding progress will be displayed (as a percentage) on the LCD. When rebuilding is already started or the logical drive is being rebuilt by a Local Spare Drive or Global Spare Drive, choose "Rebuild Progress" to see the rebuild progress.

```
LG0 RAID5 DRV=3
  2012MB RB SB=0
```

```
Rebuilding   25%
Please Wait!
```

```
Rebuild Progress
                ..
```

IMPORTANT!

- *The Rebuild function will appear only if a logical drive (with RAID level 1, 3 or 5) has a failed member.*
- *Use the "Identify Drive" function to check the exact location of a failed drive. Removing the wrong drive may cause a logical drive to fail and data loss is unrecoverable.*

Regenerating Logical Drive Parity

If no verifying method is applied to data writes, this function can be manually performed to ensure that parity errors can be mended.

From the Main Menu, press the up or down arrow keys to select "View and Edit Logical Drives."

```
View and Edit
Logical Drives
```

If you have more than one logical drive, use the up or down arrow keys to select the logical drive you would like to check the parity for, and then press **ENT.**

```
LG0 RAID5 DRV=3
  4095MB GD SB=0
```

```
Regenerate
Parity         ..
```

Press the up or down arrow keys to select "Regenerate Parity" and then press **ENT.**

To stop the regeneration process, press ESC and enter the submenu to select "Abort Regenerate Parity".

```
Abort Regenerate
Parity          ..
```

IMPORTANT!

- *If the Parity Regenerating process is stopped by a drive failure, the process cannot restart until the logical drive is rebuilt.*

Media Scan

Media Scan is used to examine drives and detect the presence of bad blocks. If any data blocks have not been properly committed, data from those blocks are automatically recalculated, retrieved, and stored onto undamaged sectors. If bad blocks are encountered on yet another drive during the rebuild process, the block LBA (Logical Block Address) of those bad blocks will be shown. If rebuild is carried out in this situation, rebuild will continue with the unaffected sectors, salvaging a majority of the stored data.

From the Main Menu, press the up or down arrow keys to select "View and Edit Logical Drives".

```
View and Edit
Logical Drives
```

The first logical drive displays. If you have more than one logical drive, use the up or down keys to select the logical drive you want to scan, and then press **ENT.**

```
LG0 RAID5 DRV=3
  4095MB GD SB=0
```

Press the up or down arrow keys to select "Media Scan" and then press **ENT.**

```
Media Scan
                ..
```

Press **ENT** again to display the first configuration option, "Priority." Press **ENT** on it and use arrow keys to select an option. Press **ENT** to confirm the change on priority level.

```
Priority
Normal          ..
```

```
Priority
To High         ?
```

Use arrow keys to move one level down to another option, "Iteration Count". This option determines how many times

```
Iteration Count
Single          ..
```

```
Iteration Count
to Continuous   ?
```

the scan is performed on the logical drive. If set to "Continuous," the scan will run in the background continuously until it is stopped by user.

If Media Scan is continuously run in the background, considerable system resources will be consumed.

Press **ENT** on your option to confirm.

Press **ENT** for two seconds to display the confirm message, then press **ENT** to start scanning the array.

```
Execute Media
Scanning       ?
```

Write Policy

From the Main Menu, press the up or down arrow keys to select "View and Edit Logical Drives".

```
View and Edit
Logical Drives
```

The first logical drive displays. If you have more than one logical drive, use the τ or σ keys to select the logical drive you want to change the write policy of; and then press **ENT.**

```
LG0 RAID5 DRV=3
  4095MB GD SB=0
```

Use arrow keys to select "Write Policy" and then press **ENT.**

```
Write Policy
                ..
```

```
Write Policy
Write-Back     ?
```

The Write-Back cache setting is configurable on a per array basis. Setting to the default value means the array setting is coordinated with the controller's general setting. The controller's general setting option can be found in "View and Edit Config Parms" -> "Caching Parameters" -> "Write-Back Cache". Note that cached writes are lost if a power failure occurs unless cached data has been duplicated to a partner controller and a battery is supporting cache memory.

5.10 Viewing and Editing Host LUNs

Viewing and Deleting LUN Mappings

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**.

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

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

```
CH0 ID002
Sec. Ctlr     ..
```

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

```
CH0 ID0 LUN0
Mapto LG0 PRT0
```

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

```
Delete CH0 ID0
LUN=00 Mapping ?
```

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

```
CH0 ID0 LUN0
Not Mapped
```

Pass-through SCSI Commands

Pass-through SCSI commands facilitate functions like downloading firmware for drives or devices (not controller firmware), setting SCSI drive mode parameters, or monitoring a SAF-TE/S.E.S. device directly from the host. To perform such a function, the channel device must be mapped to a host ID.

From the Main Menu, press the up or down arrow keys to select "View and Edit Host LUNs."

```
View and Edit
Host Luns
```

If you have primary and secondary controllers, use the up or down keys to select the controller for the device that you would like to map.

```
Map CH=0
ID=0 Pri Ctlr ?
```

Press the up or down arrow keys to choose to map an ID to "Logical Drive" or other logical unit, and then press **ENT**.

```
Map to
Logical Drive ?
```

WARNING!

- *When a drive/device is mapped to a host ID so that Pass-through SCSI Commands can be used, the data on that drive/device will not be protected by the controller. Users who employ Pass-through SCSI Commands to perform any write commands to drive media do so at their own risk.*

5.11 Viewing and Editing Drives

Scanning a New SCSI Drive (only available with SCSI drives)

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

```
View and Edit
Drives .
```

SCSI drive information will be displayed on the LCD. Press **ENT** on a drive. Use the up or down arrow keys to select "Scan New SCSI Drive," then press **ENT** again.

```
Scan new
Drive          ..
```

Press the up or down arrow keys to select a SCSI channel, then press **ENT** for two seconds.

```
Scan Channel=1 ?
```

Press the up or down arrow keys to select a SCSI ID, then press **ENT** for two seconds.

```
Scan Channel=1
ID= 01         ?
```

The information of the scanned SCSI drive will be displayed on the LCD. If the drive was not detected on the selected SCSI channel and ID, the LCD will display "Scan Fail!"

```
C=1 I=0   1010MB
NEW DRV  SEAGATE
```

```
Scan Channel=1
ID=1 Scan Fail!
```

An empty drive entry is added for this channel/SCSI ID for enclosure

```
C=1 I=1 ABSENT
```

management. The drive status is "ABSENT."

To clear the empty drive entry, press **ENT** and use arrow keys to select "Clear Drive Status," then press **ENT** to proceed.

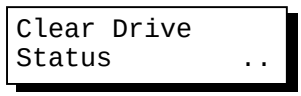

Press **ENT** for two seconds to confirm the drive entry's deletion. Information of other drives will be displayed instead.

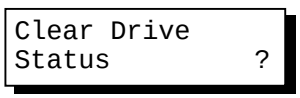

Identifying a Drive

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

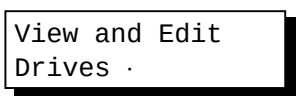

SCSI drive information will be displayed. Press the up or down arrow keys to select a SCSI drive, then press **ENT.**

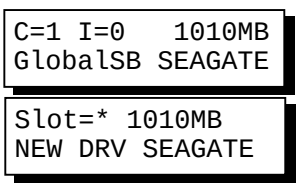

Press the up or down arrow keys to select "Identify Drive," then press **ENT** to continue.

Press the up or down arrow keys to select "Flash All Drives", "Flash Selected Drive", or "Flash All But Selected Drive". Press **ENT** for two seconds to flash the read/write LEDs of all the connected drives.

Or, press the up or down arrow keys to select "Flash Selected SCSI Drives," then press **ENT** for two seconds to flash the read/write LED of the selected drive. The read/write LED will light for a configurable time period from 1 to 999 seconds.

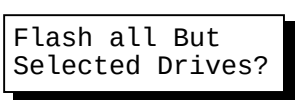

Deleting Spare Drive (Global / Local Spare Drive)

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

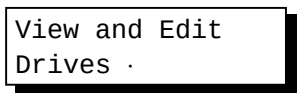

SCSI drive information will be displayed on the LCD. Press the up or down arrow keys to select the spare drive you wish to delete, then press **ENT.**

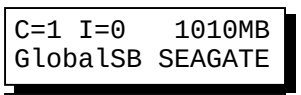

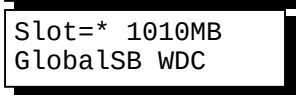

Press the up or down arrow keys to select "Delete Spare Drive," then press **ENT** to continue.

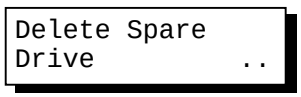

Press **ENT** for two seconds to delete the spare drive.

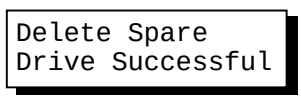

5.12 Viewing and Editing SCSI Channels

Redefining Channel Mode

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

Channel information will be displayed. Press the up or down arrow keys to browse through the information of all channels. Press **ENT** on the channel you wish the channel mode changed.

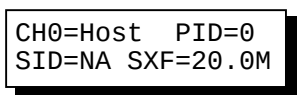

Press the up or down arrow keys to select "Redefine Channel Mode," then press **ENT**.

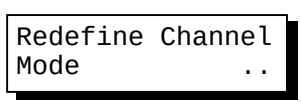

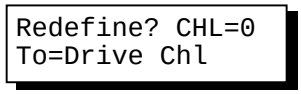

Press **ENT** for two seconds to change the channel mode.

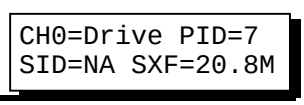

The new setting will be displayed.

IMPORTANT!

- *Every time you change a channel mode, you must reset the controller for the changes to take effect.*

Setting a SCSI Channel's ID - Host Channel

Viewing IDs

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

```
View and Edit
Channels  .
```

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

```
C0=Host  PID=0
SID=NA SXF=20.0M
```

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

```
Set SCSI Channel
ID             ..
```

Press the up or down arrow keys to browse through the existing ID settings. Press **ENT** on any ID combination to continue.

```
CHL=0 ID=0
Primary Ctrl  ..
```

Adding a Channel ID

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

Press the up or down arrow keys to choose "Add Channel SCSI ID", then press **ENT**.

```
Add Channel
SCSI ID       ..
```

Press the up or down arrow keys to choose "Primary Controller" or "Secondary Controller", then press **ENT** for two seconds to confirm.

```
Primary
Controller     ?
```

Press the up or down arrow keys to choose the SCSI ID you wish to add, then press **ENT** for two seconds to complete the process.

```
Add CHL=0 ID=2
Primary Ctlr   ?
```

Deleting a Channel ID

Press **ENT** on the existing host channel ID you want to delete. Press the up or down arrow keys to choose "Delete Channel SCSI ID," then press **ENT**.

```
Delete Channel
SCSI ID       ..
```

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

```
Delete ID=2
Primary Ctlr   ?
```

IMPORTANT!

- *Every time you make changes to channel IDs, you must reset the controller for the configuration to take effect.*
- *The reserved IDs for SCSI-based controllers are shown below:*
  Single controller configuration (SCSI-based controllers):
  Drive channels – "7"
  Redundant controller configuration:
  Drive channels – "8" and "9"

  For IDs reserved in different controller/subsystem configurations, please refer to the ***Installation and Hardware Reference Manua****l that came with your system. For controllers connected through back-end PCBs, firmware can detect their board types and automatically apply the preset IDs. There is no need to set IDs for these models.*
- *In single controller mode, you should set the Secondary Controller's ID to "NA." If a secondary controller exists, you need to set an ID for it on each of your drive channels.*
- *Multiple target IDs can be applied to host channels while each drive channel has only one or two IDs (in redundant mode).*
- *At least one controller's ID has to be present on each channel bus.*

Setting a Channel's Primary ID - Drive Channel

Press **ENT** for two seconds to enter the Main Menu. Press the up or down

```
View and Edit
Channels  .
```

arrow keys to select "View and Edit SCSI Channels," then press **ENT**.

Channel information will be displayed. Press **ENT** on the drive channel for which you wish the ID changed.

```
C1=Drive  PID=7
SID=NA SXF=80.0M
```

Press the up or down arrow keys to select "Set SCSI Channel Pri. Ctlr ID..", then press **ENT**.

```
Set SCSI Channel
Pri. Ctlr ID  ..
```

Press the up or down arrow keys to select a new ID, then press **ENT** for two seconds to confirm.

```
Set Pri. Ctlr
ID: 7 to ID: 8 ?
```

Setting a SCSI Channel's Secondary ID - Drive Channel

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

```
View and Edit
Channels  .
```

Channel information will be displayed. Press **ENT** on the drive channel for which you wish the ID changed.

```
C1=Drive  PID=7
SID=NA SXF=20.0M
```

Press the up or down arrow keys to select "Set SCSI Channel Sec. Ctlr ID..", then press **ENT**.

```
Set SCSI Channel
Sec. Ctlr ID  ..
```

Press the up or down arrow keys to select a new ID, then press **ENT** for two seconds to confirm.

```
Set Sec. Ctlr
ID:NA to ID: 9 ?
```

Setting the Channel Bus Terminator

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

```
View and Edit
Channels  .
```

Channel information will be displayed. Press the up or down arrow keys to select a channel. Press **ENT** on the

```
C0=Host  PID=0
SID=NA SXF=20.0M
```

```
Set SCSI Channel
Terminator    ..
```

channel for which you wish the terminator mode changed.

Press the up or down arrow keys to select "Set SCSI Channel Terminator," then press **ENT**.

Its current status will be displayed on the LCD. Press **ENT** to continue.

```
SCSI Terminator
Enabled       ..
```

Press **ENT** again for two seconds to change the terminator mode to the alternate setting.

```
CHL=0 Disable
Terminator    ?
```

IMPORTANT!

- *You can use the terminator jumpers on the controller board to control SCSI bus termination of the SentinelRAID series controllers. When using jumpers to control, the firmware termination setting must be disabled. To disable SCSI termination of a SCSI bus, the associated terminator jumpers must be left open, and the firmware setting must be disabled.*

Setting Transfer Speed

IMPORTANT!

Transfer rate	Transfer speed in Mhz
SCSI-320	160Mhz
SCSI-160	80Mhz

- According to SCSI Channel Specifications, transfer speed shown in Mhz is the SCSI bus sync frequency. With 16-bit Wide transfer, the actual data rate is the double of Mhz transfer speed.
- Every time you change the Transfer Speed, you must reset the controller for the changes to take effect.

Transfer speed refers to the SCSI bus speed in synchronous mode. Asynchronous mode is also available in this option setting. In Ultra/Ultra Wide SCSI, the maximum synchronous speed is 160MHz.

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

```
View and Edit
Channels  .
```

Channel information will be displayed. Press the up or down arrow keys to

```
CH0=Host  PID=0
SID=NA SXF=80.0M
```

select a channel. Press **ENT** on the channel for which you wish the transfer speed changed.

Press the up or down arrow keys to select "Set Transfer Speed," then press **ENT**.

```
Set Transfer
Speed          ..
```

The current speed of this SCSI channel will be displayed. Press the up or down arrow keys to select the desired speed, then press **ENT** for two seconds to confirm.

```
CHL=0 Clk=80.0M
Change to=40.0M?
```

Setting Transfer Width

The controller supports 8-bit SCSI and 16-bit SCSI. Enable "Wide Transfer" to use the 16-bit SCSI function. Disabling "Wide Transfer" will limit the channel transfer speed to 8-bit SCSI.

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

```
View and Edit
Channels  .
```

Channel information will be displayed. Press the up or down arrow keys to browse through the channels. Press **ENT** on the channel for which you wish the transfer width changed.

```
CH0=Host  PID=0
SID=NA SXF=20.0M
```

Press the up or down arrow keys to select "Set Transfer Width," then press **ENT**.

```
Set Transfer
Width          ..
```

The current mode will be displayed. Press **ENT** to continue.

```
Wide Transfer
Enabled        ..
```

Press **ENT** again for two seconds.

```
Disable
Wide Transfer  ?
```

IMPORTANT!

- *Every time you change the SCSI Transfer Width, you must reset the controller/subsystem for the changes to take effect.*

Maximum Synchronous Transfer Clock

Press the up or down arrow keys to select a SCSI target, then press **ENT**.

To set the maximum synchronous clock of this SCSI target, choose "Max. Synchronous Xfer Clock," then press **ENT**. The current clock setting will be displayed on the LCD.

Press the up or down arrow keys to change the clock, then press **ENT** for two seconds.

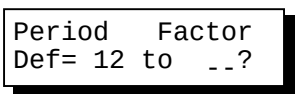

Maximum Transfer Width

Press the up or down arrow keys to select a SCSI target, then press **ENT**.

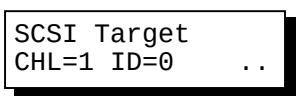

To set the maximum transfer width of this SCSI target, choose "Max. Xfer Narrow Only" or "Max. Xfer Wide Supported," then press **ENT**. The current clock setting will be displayed on the LCD.

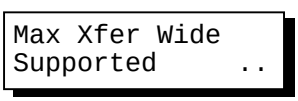

Press **ENT** for two seconds to change the setting.

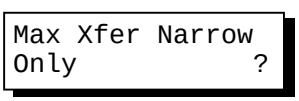

Parity Check

Press the up or down arrow keys to select a SCSI target, then press **ENT**.

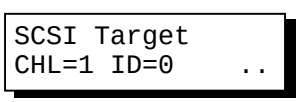

Choose "Parity Check," then press **ENT**. The current clock setting will be displayed on the LCD.

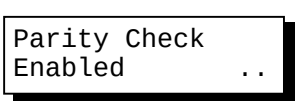

Press **ENT** for two seconds to change the setting.

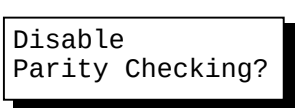

Maximum Tag Count

Press the up or down arrow keys to select a SCSI target, then press **ENT**.

```
SCSI Target
CHL=1 ID=0    ..
```

Choose "Max Tag Count," then press **ENT**. The current clock setting will be displayed on the LCD.

```
Max Tag Count:
Default( 32)  ..
```

Press the up or down arrow keys to change the setting, then press **ENT** for two seconds to change the setting.

```
Tag Cur=32
Set to:Default ?
```

IMPORTANT!

- *Disabling the Maximum Tag Count will disable the internal cache of this SCSI drive..*

Restore to Default Setting

Press the up or down arrow keys to select a SCSI target, then press **ENT**.

```
SCSI Target
CHL=1 ID=0    ..
```

Choose "Restore to Default Setting," then press **ENT**.

```
Restore to
Default Setting.
```

Press **ENT** again for two seconds to restore the SCSI target's default settings.

```
Restore to
Default Setting?
```

Data Rate

This option is available in the configuration menu of Fibre host channel and the drive channel configuration menus of Fibre-, ATA-, or SATA-based subsystems. Default is "AUTO" and should work fine with most drives. Changing this setting is not recommended unless some particular bus signal issues occur.

The host channel data rate setting allows a 2Gbit Fibre Channel to negotiate with devices communicating over 1GHz link if the link speed configuration is determined externally.

Most SATA/ATA-based systems connect only one drive per SATA/ATA channel (4 for multi-lane with SATA-II). This helps to avoid a single drive failure from affecting other drives. The maximum mechanical performance of today's drives can reach around 30MB/second (sustained read). This is still far below the bandwidth of a drive channel bus. Setting the SATA/ATA bus speed to a lower value can get around some problems, but will not become a bottleneck to system performance.

Note that the SATA/ATA speed is the maximum transfer rate of the SATA/ATA bus in that mode. It does not mean the drive can actually carry out that amount of sustained read/write performance. For the performance of each drive model, please refer to the documentation provided by drive manufacturer.

Host Channel:

Select from Main Menu, "View and Edit Channels," and then a host channel you wish to change its data rate. Press **ENT** on the channel and use the arrow keys to find the "Data Rate" option. Press **ENT** on the Data Rate option to display "Set Chl=X Data Rate To AUTO?", where "X" stands for the channel number.

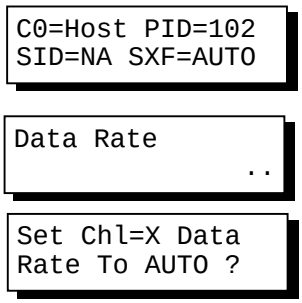

Use your arrow keys to display a different data rate (1GHz or 2GHz). Press **ENT** to confirm a selection.

Drive Channel:

Select from Main Menu, "View and Edit Channels," and then a drive channel you wish to change its data rate. Press **ENT** on the channel and use the arrow keys to find the "Data Rate" option. Press **ENT** on the Data Rate option to display "Set Chl=X Data Rate To AUTO?", where "X" stands for the channel number.

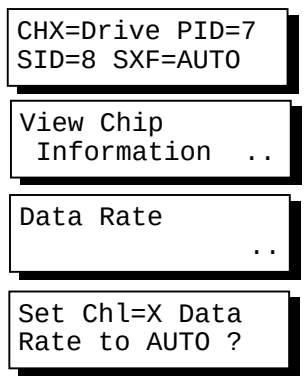

Use your arrow keys to display a data rate value which ranges from 33 to 150MB/s (SATA drive channels). Press **ENT** to confirm a selection.

Issue LIP

This option allows you to manually issue a LIP (Loop Initialization Sequence) on either a host or drive channel. This function is particularly useful when cabling failure occurs between two cascaded drive enclosures.

The condition for applying the "Issue LIP" function is described as follows (see diagram Figure 5-1):

1. Multiple enclosures have been cascaded using FC-AL Fibre links.

2. One cable link fails between drive enclosures or is disconnected.

3. Cabling failure is restored.

4. Since the FC port to which the failed link was attached to has already been bypassed at the time when the fault occurred, a LIP command must be manually exerted by the user for the FC port to participate in the loop again. The dual-loop FC-AL can thus be restored.

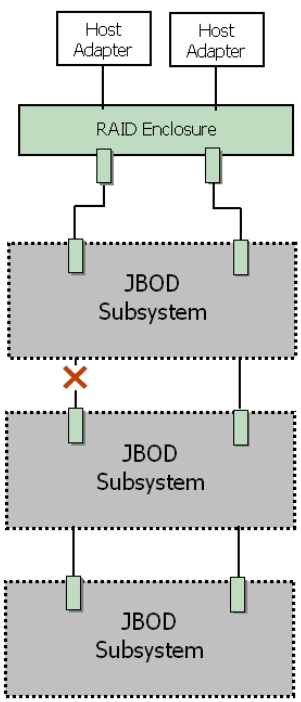

Figure 5 - 2 Cabling Failure between Drive Enclosures

Select from under, "View and Edit Channels," and then a host or drive channel over which you wish to issue a

```
Issue LIP     ?
```

LIP command. Press **ENT** on the channel and use the arrow keys to display the "Issue LIP" command. Press ENT to exert the command.

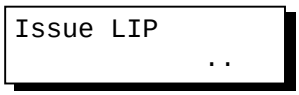

5.13 System Functions

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

Mute Beeper

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.

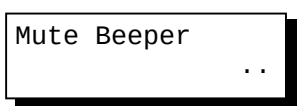

Change Password

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 16-character space. The maximum number of characters for a controller password is 15. If 15 characters are used for a controller name, there will be only one character left for the controller password and vice versa. Since 3.42J09, later firmware revisions will support a 32-character space.*

Changing Password

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

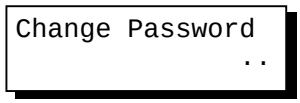

If the password has previously been set, the controller will ask for the old password first. If password has not yet been set, the controller will directly ask for the new password. The password cannot be replaced unless the correct old password is provided.

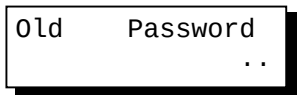

Press the up or down arrow keys to select a character, then press **ENT** to move to the next space. After entering all the characters (alphabetic or numeric), press **ENT** for two seconds to confirm. If the password is correct, or there is no preset password, it will ask for the new password. Enter the password again to confirm.

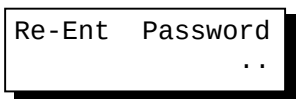

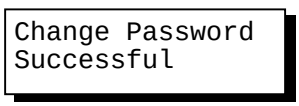

Disabling Password

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.

Reset Controller

To reset the controller without powering off the system, Press the up or down arrow keys to "Reset Controller," then press **ENT**. Press ENT again for two seconds to confirm. The controller will now reset.

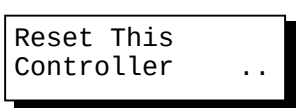

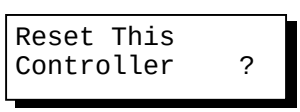

Shutdown Controller

Before powering off the controller, unwritten data may still reside in cache memory. Use the "Shutdown Controller" function to flush the cache content. Press the up or down arrow keys to "Shutdown Controller," then press **ENT**. Press **ENT** again for two seconds to confirm.

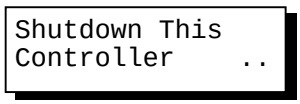

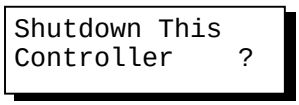

The controller will now flush the cache memory. Press **ENT** for two seconds to confirm and reset the controller or power off the controller.

Controller Maintenance

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

Saving NVRAM to Disks

You can choose to backup your controller-dependent configuration information to disk. 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 data as a file to a computer system drive.

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

From the Main Menu, choose "System Functions." Use arrow keys to scroll down and select "Controller Maintenance," "Save NVRAM to Disks," then press **ENT.** Press **ENT** for two seconds on the message prompt, "Save NVRAM to Disks?".

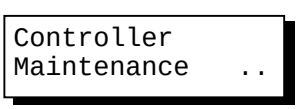

A prompt will inform you that NVRAM information has been successfully saved.

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.

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

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

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.

From the main menu, choose "View and Edit Peripheral Dev." 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. Press **ESC** to return to the previous menu.

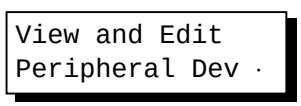

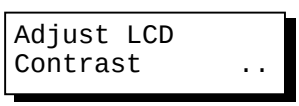

5.14 Controller Parameters

Controller Name

Select "View and Edit Config Parms" from the Main Menu. Choose "View and Edit Configuration Parameters," "Controller Parameters," then press **ENT**. 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.

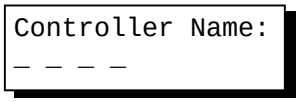

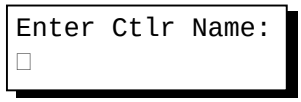

LCD Title Display Controller Name

Choose "View and Edit Configuration Parameters," "Controller Parameters," then press **ENT**. Use the up or down arrow keys to choose to display the embedded controller logo or any given name on the LCD initial screen.

Password Validation Timeout

Choose "View and Edit Configuration Parameters," "Controller Parameters," then press **ENT**. 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 to Always Check. The Always Check timeout will invalidate any configuration change without entering the correct password.

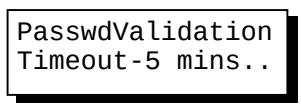

Controller Unique Identifier

Choose "View and Edit Configuration Parameters," "Controller Parameters," then press **ENT**. Press the up or down arrow keys to select "Ctlr Unique ID-,"

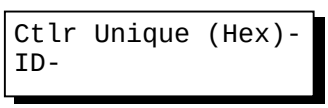

then press **ENT**. Enter any hex number between "0" and "FFFFF" and press **ENT** to proceed.

Enter a unique ID for any RAID controller in a single or dual-controller configuration. The unique ID is recognized by the controller as the following:

1. A controller-specific identifier that helps controllers to identify its counterpart in a dual-active configuration.

2. The unique ID is combined to generate a unique WWN node name for controllers or RAID systems using Fibre Channel host ports. The unique node name helps prevent host computers from mis-addressing the storage system during controller failback/failover processes.

3. MAC addresses for the controller's Ethernet port that should be taken over by a surviving controller in the event of controller failure.

Controller Date and Time

This submenu is only available for controllers or subsystems that come with a real-time clock on board.

Time Zone

Choose "View and Edit Configuration Parameters," "Controller Parameters," then press **ENT**. Press the up or down arrow keys to scroll down and select "Set Controller Date and Time", then press **ENT**.

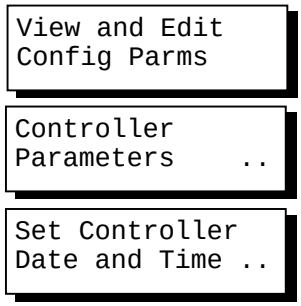

The controller uses GMT (Greenwich Mean Time), a 24-hours clock. To change the clock to your local time zone, enter the hours later than the Greenwich mean time after a plus (+) sign. For example, enter "+9" for Japanese time zone.

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

Use the down key to enter the plus sign and the up key to enter numeric representatives.

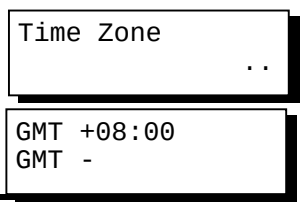

Date and Time

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

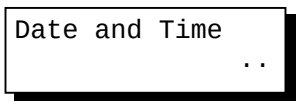

Use the arrow keys to select and enter the numeric representatives in the following order: month, day, hour, minute, and the year.

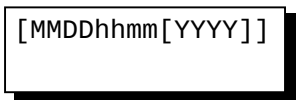

5.15 SCSI Drive Utilities

From the "View and Edit SCSI Drives" menu, select the drive that the utility is to be performed on; then press **ENT**. Select "SCSI Drive Utilities; then press **ENT**. Choose either "SCSI Drive Low-level Format" or "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 drive.

```
View and Edit
Drives
```

```
C=1 I=1  8683MB
NEW DRV  SEAGATE
```

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

```
Drives
  Utilities   ..
```

```
Drive Read/Write
 Test         ..
```

SCSI Drive Low-level Format

Choose "SCSI Drive Low-level Format" and confirm by selecting **Yes**.

```
Drive Low-Level
 Format       ..
```

IMPORTANT!

- *Do not switch the controller's and/or disk drive's power off during the Drive Low-level Format. If any power failure occurs during a drive low-level format, the formatting must be started over again when power resumes.*
- *All of the data stored in the disk drive will be destroyed during a low-level format.*
- *The disk drive on which a low-level disk format will be performed cannot be a spare drive (local or global) nor a member drive of a logical drive. The "SCSI Drive Low-level Format" option will not appear if the drive's status is not stated as a "New Drive" or a "Used Drive".*

SCSI Drive Read/Write Test

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

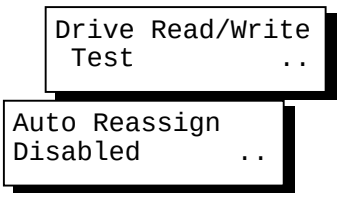

Press the up or down arrow keys to select and choose to enable/disable the following options:

1. Auto Reassign Bad Block
2. Abort When Error Occurs
3. Drive Test for - Read Only/Read and Write.

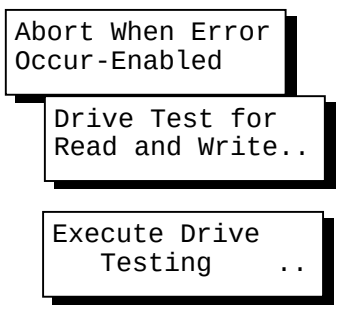

When finished with configuration, select "Execute Drive Testing" and press **ENT** to proceed.

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

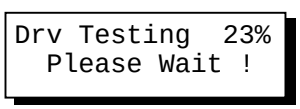

You may press **ESC** and select "Read/Write Test" later and press τ or σ to select to "View Read/Write Testing Progress" or to "List Current Bad Block Table." If you want to stop testing the drive, select "Abort Drive Testing" and press **ENT** to proceed.

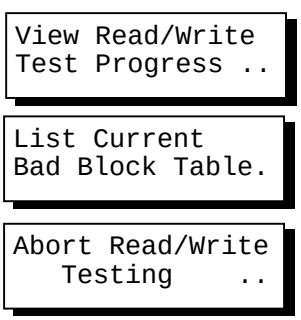

Chapter

6 Terminal Screen Messages

6.1 The Initial Screen

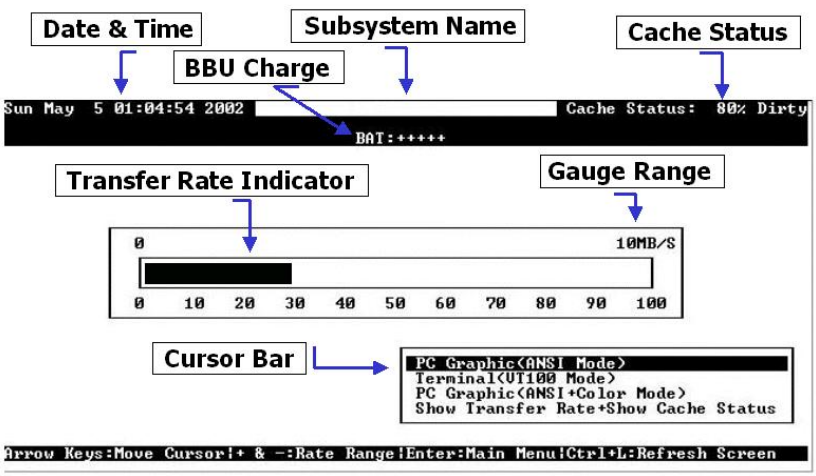

Cursor Bar:	Move the cursor bar to a desired item, then press **[ENTER]** to select
Controller Name:	Identifies the type of controller/subsystem or a preset name
Transfer Rate Indicator:	Indicates the current data transfer rate
Gauge Range:	Use + or - keys to change the gauge range in order to view the transfer rate indicator
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

6.2 Main Menu

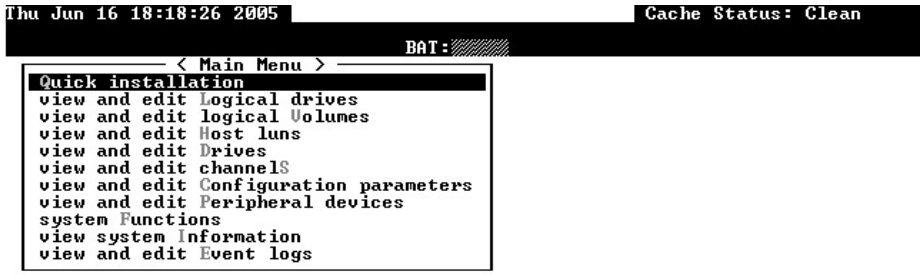

Use the arrow keys to move the cursor bar through the menu items, then press **[ENTER]** to choose a menu, or **[ESC]** to return to the previous menu/screen.

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, or several "+ " (plus) signs will be used to indicate battery charge. A battery fully-charged will be indicated by five plus signs.

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).

6.3 Quick Installation

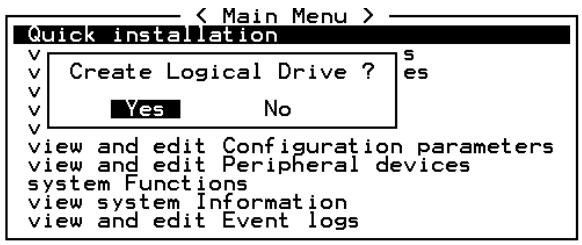

Type **Q** or use the ↑ ↓ keys to select "Quick installation", then press **[ENTER]**. Choose **Yes** to create a logical drive.

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.

The controller will start initialization and automatically map the logical drive to LUN 0 of the first host channel.

6.4 Logical Drive Status

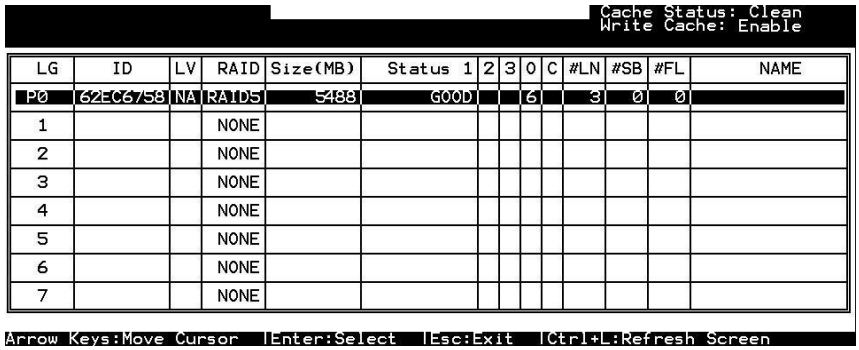

NOTE:

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

LG Logical Drive number

P0: Logical Drive 0 managed by the Primary Controller

S0: Logical Drive 0 managed by the Secondary Controller

LV The Logical volume to which this logical drive belongs

ID Controller-generated unique 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
CREATING	Logical drive is being initiated
DRV ABSENT	An empty drive tray
INCOMPLETE	Two or more drives 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 FAIL	Two 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

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		
	2	4KB	**6**	64KB
	3	8KB	**7**	128KB
	4	16KB	**8**	256KB
	5	32KB		
Column C		Logical Drive Status – Write Policy setting		
	B	Write-back		
	T	Write-through		
#LN		Total 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)		

6.5 Logical Volume Status

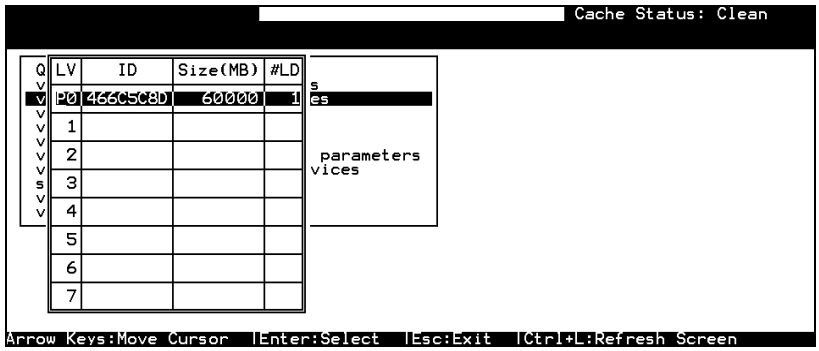

NOTE:

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

LV	Logical Volume number. P0: Logical Volume 0 managed by the Primary Controller S0: Logical Volume 0 managed by the Secondary Controller
ID	Logical Volume ID number (controller randomly generated)
Size(MB)	Capacity of the Logical Volume
#LD	The number of Logical Drive(s) included in this Logical Volume

6.6 Drive Status

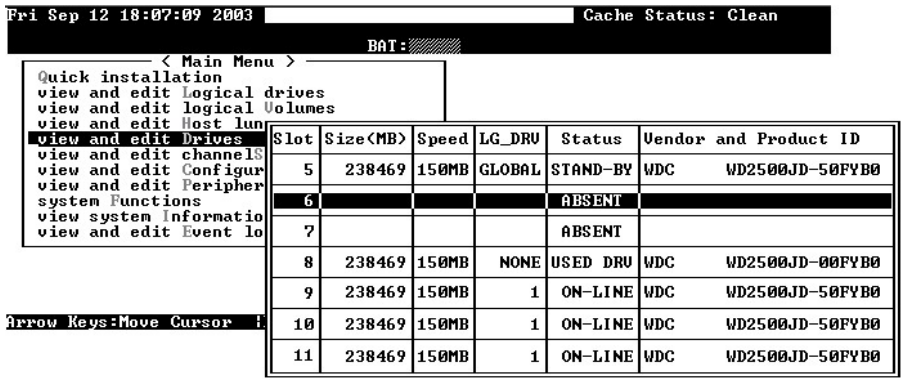

Slot Slot number of the drive;

"S" indicates this is the drive used for passing through SES signals (Fibre drives in a JBOD equipped with SES modules only)

Chl The drive channel where the drive is connected (only available on SCSI- and Fibre-based controllers/subsystems)

"X<Y>" indicates two channels are configured in a dual-loop

ID The channel ID assigned to this drive (only available on SCSI- and Fibre-based controllers/subsystems)

Size (MB) Drive capacity

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

Async The drive is using asynchronous mode.

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

If the Status column shows "STAND-BY", the drive is a Local Spare of logical drive *x*.

Status **Global** The SCSI drive is a Global Spare Drive

INITING Processing array initialization

ON-LINE The drive is in good condition

REBUILD Processing Rebuild

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 configured to any logical drive or as a spare drive

USED DRV An used drive that is not a member of any logical drive or configured as spare

FRMT DRV Formatted drive (drive formatted with a reserved section)

BAD Failed drive

ABSENT	Drive does not exist
MISSING	Drive once existed, but is missing now
SB-MISS	Spare drive missing
Vendor and Product ID	The vendor and product model information of the drive

6.7 Channel's Status

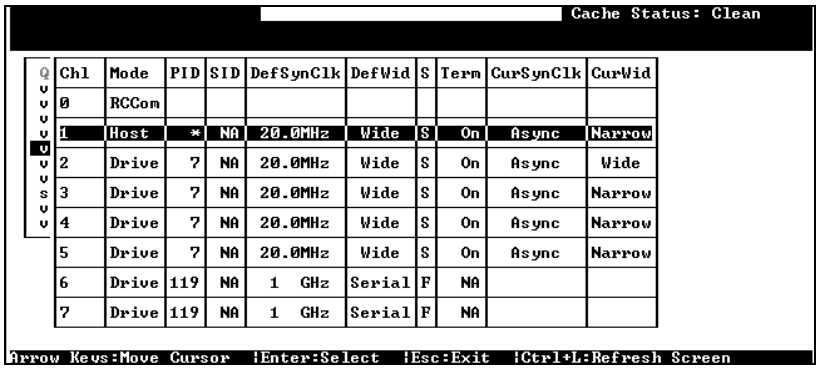

Chl SCSI channel number

Mode Channel mode

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

PID IDs managed by the Primary Controller

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

SID IDs managed by the Secondary Controller

*	Multiple IDs were applied (Host Channel mode only)
(ID number)	Host Channel: Specific IDs managed by the Secondary Controller for host LUN mapping Drive Channel: Specific ID reserved for the channel processor on the Secondary controller; used in redundant controller mode
NA	No SCSI ID applied

DefSynClk Default SCSI 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 SCSI Bus Width:

Wide	16-bit SCSI
Narrow	8-bit SCSI

S Signal:

S	Single-ended
L	LVD
F	Fibre

Term	Terminator Status:	
	On	Terminator is enabled.
	Off	Terminator is disabled.
	Diff	The channel is a Differential channel. The terminator can only be installed/removed physically.
CurSynClk	Current SCSI bus synchronous clock:	
	*??.?*M	The default setting of the SCSI channel is ??.? MHz in Synchronous mode.
	Async.	The default setting of the SCSI channel is Asynchronous mode.
	(empty)	The default SCSI bus synchronous clock has changed. Reset the controller for the changes to take effect.
CurWid	Current SCSI Bus Width:	
	Wide	16-bit SCSI
	Narrow	8-bit SCSI
	(empty)	The default SCSI bus width has changed. Reset the controller for the changes to take effect.

6.8 Controller Voltage and Temperature

Controller Voltage and Temperature Monitoring

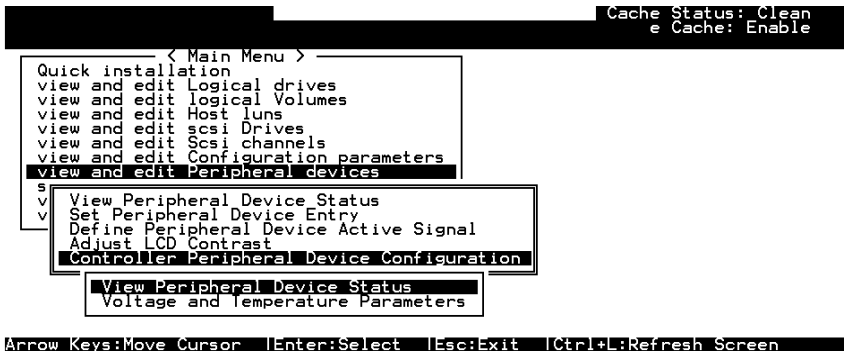

Choose from Main Menu "View and Edit Peripheral Devices," and press [ENTER]. From the submenu, choose "Controller Peripheral Device Configuration," "View Peripheral Device Status", then press [ENTER].

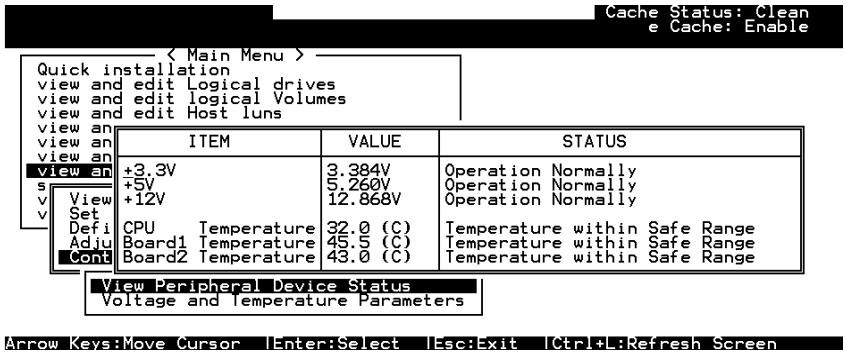

The current status of voltage and temperature detected by the controller will be displayed on-screen and will be stated as normal or out of order.

6.9 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 cause an automatic deletion of all the recorded event logs. To view the events log on-screen, from the Main Menu "View and Edit Event Logs" by pressing **[ENTER]**.

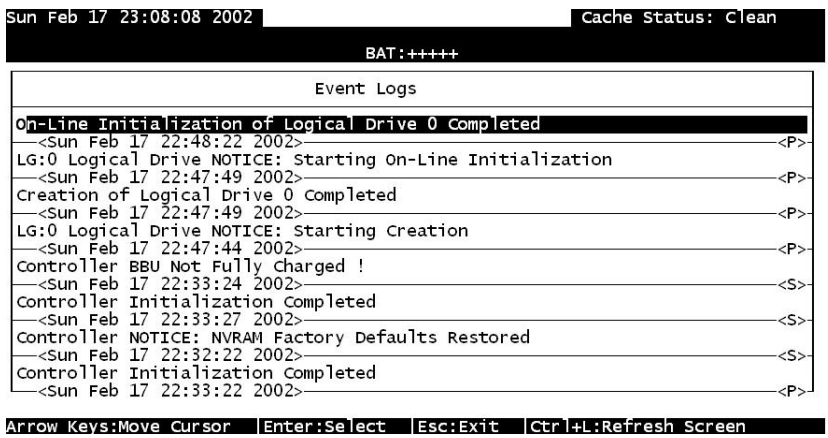

The "P" or "S" icon on the right indicates which one of the controllers (Primary or Secondary) issued an event in a dual-controller configuration.

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.

Chapter

7

Terminal Operation

7.1 Power on RAID Enclosure

Hardware installation should be completed before powering on your RAID enclosure. Drives must be configured and the controller properly initialized before the host computer can access the storage capacity. The 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 return to the previous menu/screen.

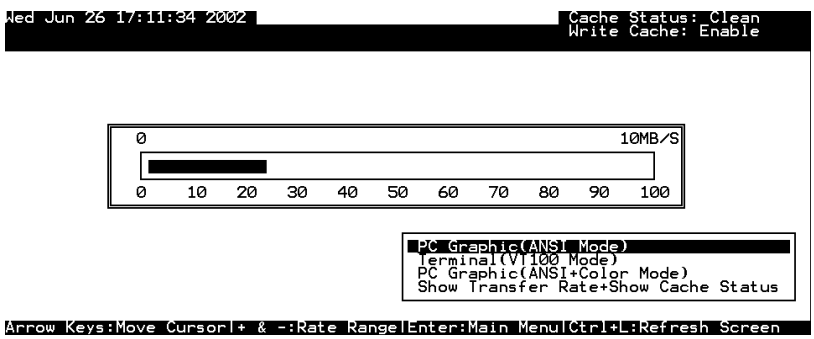

7.2 Caching Parameters

Optimization Modes

Mass storage applications can be roughly divided into two categories database and video/imaging - according to their read/write characteristics. To optimize system operation for these

two categories, the controller has two embedded optimization modes with controller behaviors pre-adjusted to different read/write parameters: "Optimization for Random I/O" and the "Optimization for Sequential I/O."

Limitations:

There are limitations on the use of optimization modes.

1. You can select the stripe size of each array (logical drive) during the initial configuration. However, changing stripe size is only recommended for experienced engineers who have tested the effects of tuning stripe sizes for different applications.

2. The array stripe size can only be changed when creating a logical drive.

3. Think twice before choosing the optimization mode. Once the controller optimization mode is applied, access to different arrays in a RAID system will follow the same optimized pattern. **You can only change the optimization mode after recreating the arrays.**

Database and Transaction-based Applications:

These kinds of applications usually includes SQL server, Oracle server, Informix, or other database services. These applications keep the size of each transaction down to a minimum, so that I/Os can be rapidly processed. Due to its transaction-based nature, these applications do not read or write a bunch of data in a sequential order - access to data occurs randomly. The transaction size usually ranges from 2K to 4K. Transaction performance is measured in "I/Os per second" or "IOPS."

Video Recording/Playback and Imaging Applications:

These kinds of applications usually include video playback, video post-production editing, or applications of a similar nature. These applications have the tendency to read or write large files from and into storage in a sequential order. The size of each I/O can be 128K, 256K, 512K, or up to 1MB. The efficiency of these applications is measured in "MB/sec."

When an array works with applications such as video or image-oriented applications, the application reads/writes from the drive as large-block, sequential threads instead of small and randomly accessed files.

The controller optimization modes have read-ahead buffer and other Read/Write characteristics tuned to obtain the best performance for these two major application categories.

Optimization Mode and Stripe Size

Each controller optimization mode has preset values for the stripe size of arrays created in different RAID levels. If you want a different optimization mode for a configured array, you need a duplicate of the data in your array elsewhere, select a different optimization mode, reset the controller/subsystem, and re-create the array. Once the array is created, stripe size cannot be changed.

Using the default value should be sufficient for most applications. The interrelation among stripe sizes, optimization modes, and RAID levels is shown in the table below:

Table 7 - 1 RAID Levels, Optimization Modes, and Stripe Sizes

	Stripe Size: Opt. for Sequential I/O	Stripe Size: Opt. for Random I/O
RAID0	128	32
RAID1	128	32
RAID3	16	4
RAID5	128	32

Optimization for Random or Sequential I/O

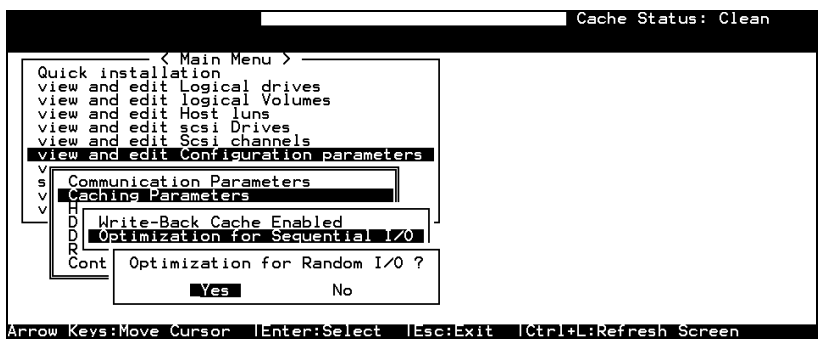

Choose "Optimization for Random I/O" or "Optimization for Sequential I/O," then press **[ENTER].** The "Random" or "Sequential" dialog box will appear, depending on the option you have selected. Choose **Yes** in the dialog box that follows to confirm the setting.

Write-Back/Write-Through Cache Enable/Disable

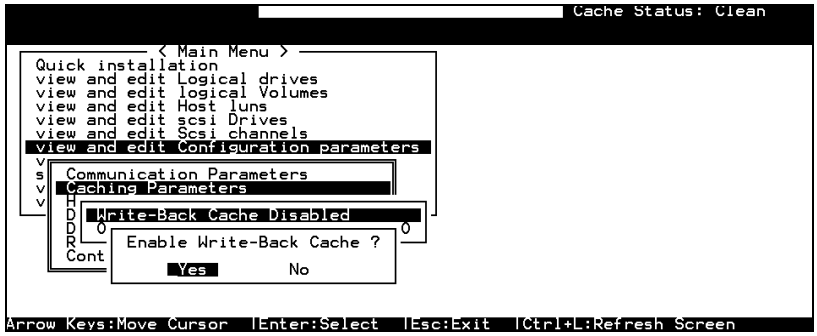

Choose "Caching Parameters", then press **[ENTER].** Select "Write-Back Cache," then press **[ENTER].** "Enabled" or "Disabled" will display the current setting with Write-Back caching. Choose **Yes** in the dialog box that follows to confirm the setting.

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 module can hold cached data for days. If a controller fails in a dual-controller configuration, an exact replica of the data cached in the failed controller is preserved by its counterpart controller and therefore the data can be retrieved once the normal working condition is restored.

Periodic Cache Flush

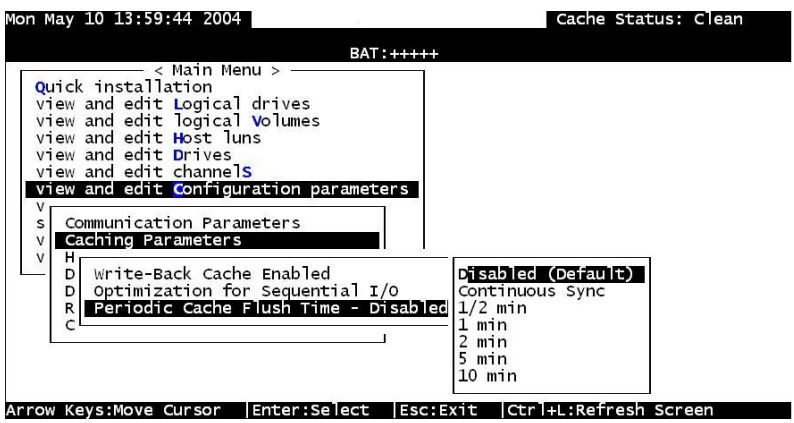

If for a reason Write-Back caching is preferred for better performance and yet data integrity is of the concern, e.g., lack of the battery protection, the system can be configured to flush the cached writes at every preset interval.

Note that the "Continuous Sync" option holds data in cache for as long as necessary to complete a write operation and immediately commits it to hard drives if it does not come in a series of sequential write requests.

Adaptive Write Policy

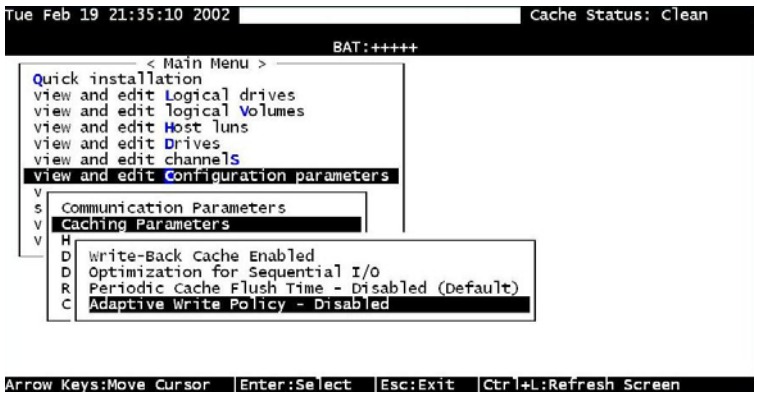

Firmware is embedded with intelligent algorithms to detect and to adapt the array's caching mode to the characteristics of I/O requests. The capability is described as follows:

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 conducting writes 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.

IMPORTANT!

- *If the size of an array is larger than 16TB, only the optimization for sequential I/O can be applied. A logical drive of this size is not practical; therefore, there is actually no limitation on the combination of optimization mode and array capacity.*
- *Every time you change the Caching Parameters, you must reset the controller for the changes to take effect.*
- *In the redundant controller configuration, write-back will only be applicable when there is a synchronized cache channel strung between partner controllers.*
- *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.*

7.3 Viewing the Connected Drives

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

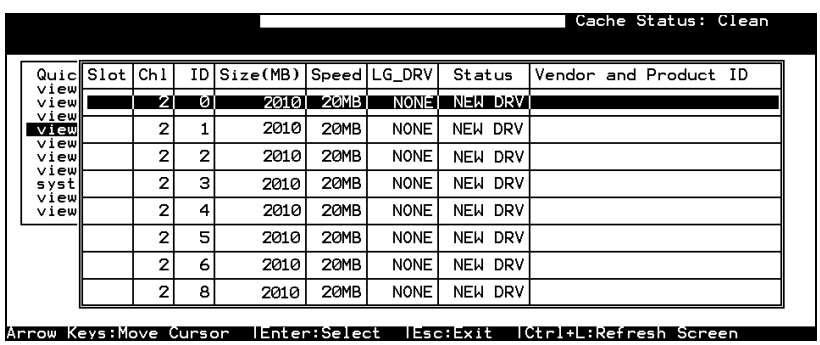

Use the arrow keys to scroll down to "View and Edit Drives." This will display information on all the physical drives installed.

Drives will be listed in the table of "View and Edit Drives." Use the arrow keys to scroll the table. You may first examine whether there is any drive installed but not listed here. If a drive is installed but not listed, the drive may be defective or not installed correctly. Please contact your RAID supplier.

IMPORTANT!

- *Drives of the same brand/model/capacity might 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.*

7.4 Creating a Logical Drive

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

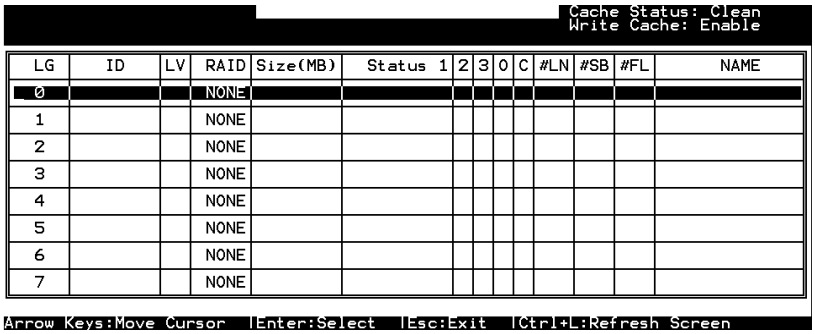

For the first logical drive on the RAID, simply choose LG 0 and press **[ENTER]** to proceed. You may create as many as 128 logical drives from drives on any drive channel or within a cascaded enclosure.

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

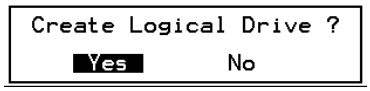

Choosing a RAID Level:

A pull-down list of supported RAID levels will appear. In this chapter, RAID 5 will be used to demonstrate the configuration process. Choose a RAID level for this logical drive.

Choosing Member Drives:

Choose your member drive(s) from the list of available physical drives. The drives can be tagged for inclusion by positioning the cursor bar on the 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 drives.

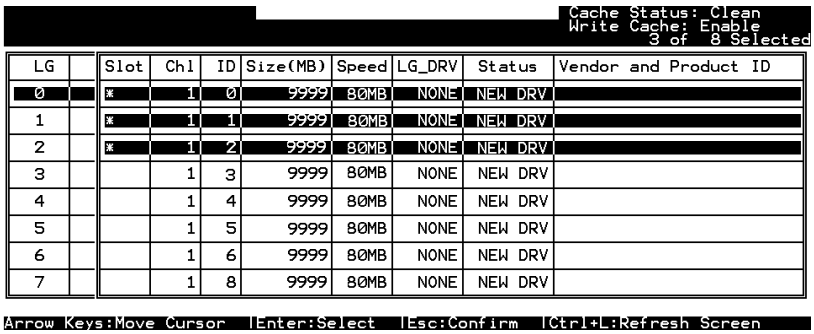

Logical Drive Preferences

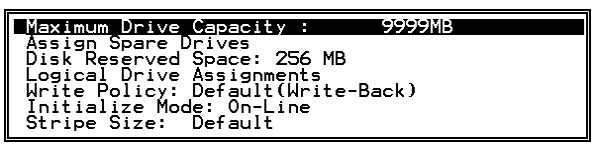

After all member drives have been selected, press **[ESC]** to continue with the next option. A list of array options is displayed.

Maximum Drive Capacity:

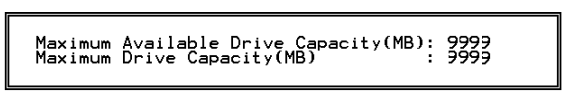

As a rule, a logical drive should be composed of drives of the same capacity. A logical drive can only use the capacity of each drive up to the maximum capacity of the smallest drive selected for the array.

Assign Spare Drives

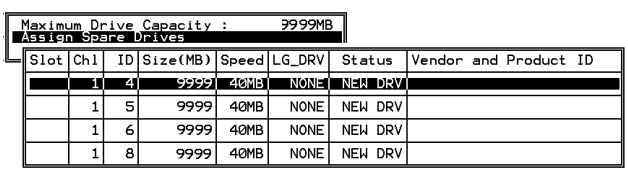

You can add a spare drive from the list of the unused drives. The spare chosen here is a Local spare and will automatically replace any failed drive in the event of drive failure. The controller will then rebuild data onto the replacement drive.

A logical drive composed in a non-redundancy RAID level (NRAID or RAID 0) does not support spare drive rebuild.

Disk Reserved Space

The reserved space is a small section of disk space formatted for storing array configuration and RAIDWatch program data. This item is display only. You can not change the size of reserved space.

Logical Drive Assignments (Dual-Active Controllers)

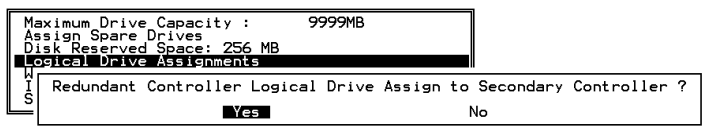

If you use two controllers for a dual-active configuration, a logical drive can be assigned to either of the controllers to balance workload. The default is the Primary Controller; press **[ESC]** if change is not preferred. Logical drive assignment can be changed at any time.

Write Policy

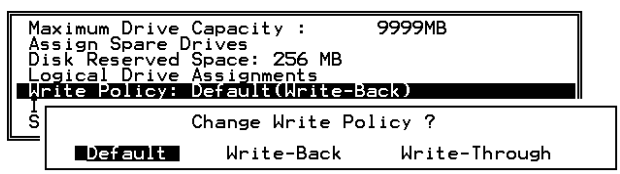

This sub-menu allows you to set the caching mode for this specific logical drive. "Default" is a neutral value that is coordinated with the controller's current caching mode setting, that you can see bracketed in the Write Policy status.

Initialization Mode

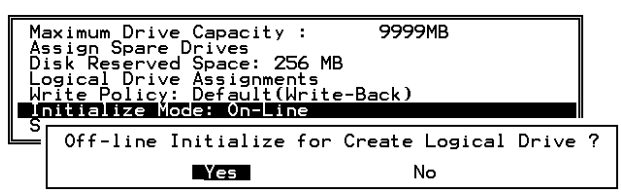

This sub-menu allows you to see if the logical drive is immediately available. If the online (default) mode is used, data can be written onto it and you may continue with array configuration, e.g., including the array into a logical volume, before the array's initialization is completed.

Stripe Size

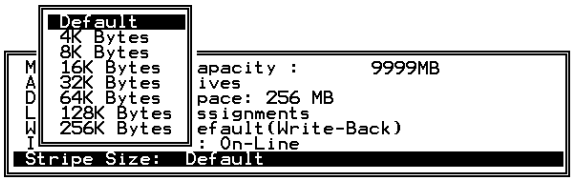

This option should only be changed by experienced engineers. Setting to an incongruous value can severely drag performance. This option should only be changed when you can be sure of the performance gains it might bring you.

The default value is determined by the combination of the controller **Optimization Mode** setting and the **RAID level** used for the array.

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

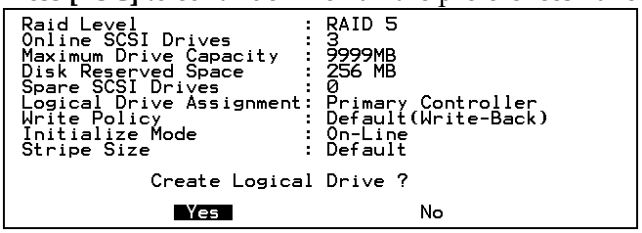

A confirm box will appear on the screen. Verify all information in the box before choosing "**Yes**" to confirm and proceed.

If the online initialization mode is applied, the logical drive will first be created and the controller will find an appropriate time to initialize the array.

The completion of array creation will be indicated by the message prompt above.

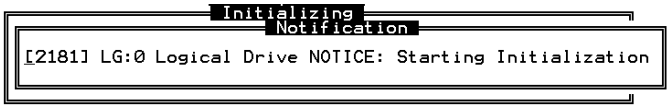

A controller event will then prompt to indicate that the logical drive initialization has begun. Tap **[ESC]** to cancel the "Notification" prompt, and a progress indicator will display on the screen as a percentage bar.

The array initialization runs in the background while you can start using the array or continue configuring your RAID system.

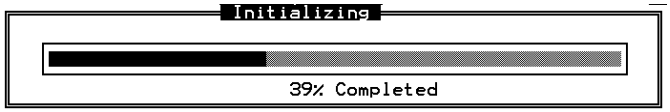

When a fault-tolerant RAID level (RAID 1, 3, or 5) is selected, the controller will start initializing parity.

Use the **[ESC]** key to view the status of the created logical drive.

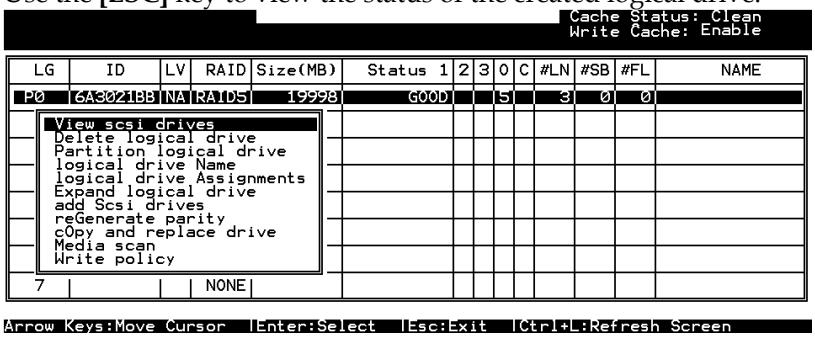

IMPORTANT!

- *Note that only logical drives with RAID levels 1, 3, or 5 will take the time to initialize the logical drive. Logical drives with RAID level 0 and NRAID do not perform logical drive initialization; the drive initialization will be finished almost immediately.*

7.5 Creating a Logical Volume

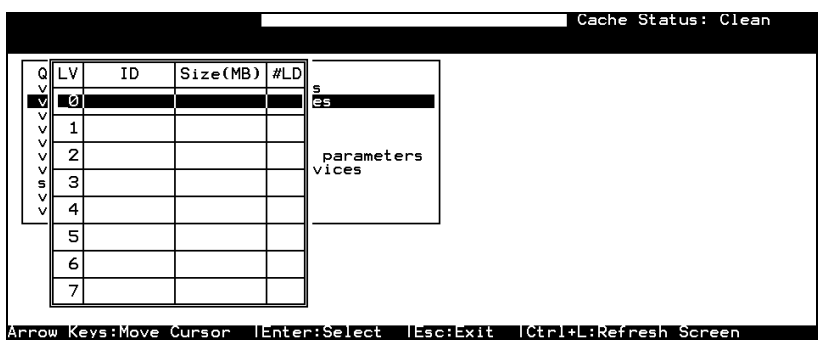

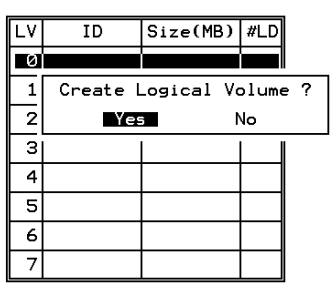

A logical volume consists of one or several logical drives. Choose "View and Edit Logical Volumes" in the Main Menu. The current logical volume configuration and status will be displayed on the screen. Choose a logical volume number (0-7) that has not yet been defined, then press **[ENTER]** to proceed. A prompt

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

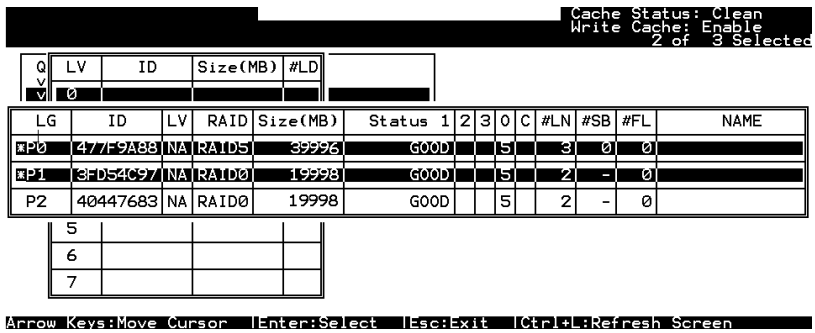

Select one or more logical drive(s) available on the list. The same as creating a logical drive, the logical drive(s) can be tagged for inclusion by positioning the cursor bar on the desired drive and then press **[ENTER]** to select. An asterisk (*) mark will appear on the selected drive. Press **[ENTER]**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

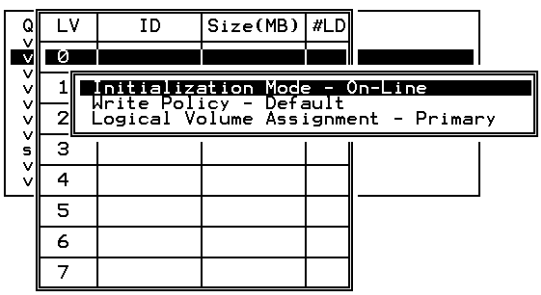

Use arrow keys to select a sub-menu and make change to the initialization mode, write policy, or the managing controller.
Logical volumes can be assigned to different controllers (primary or secondary). Default is primary.

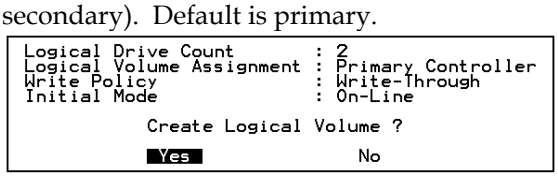

Note that if a logical volume is manually assigned to a specific controller, all its members' assignment will also be shifted to that controller.

As all the member logical drives are selected, press **[ESC]** to continue. The confirm box displays. Choose **Yes** to create the logical volume.

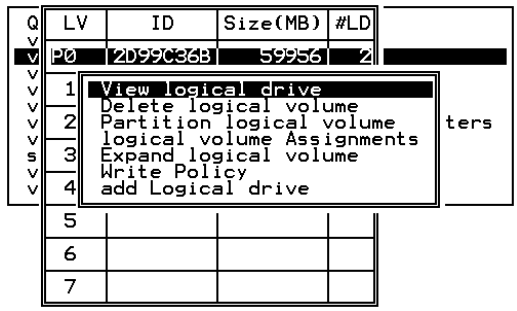

Press **[ENTER]** and the information of the created logical volume displays.

LV:	Logical Volume ID
P0:	Logical Volume 0 managed by the primary controller
S0:	Logical Volume 0 managed by the secondary controller
ID:	Unique ID for the logical volume, randomly generated by the controller
Size:	Capacity of this volume
#LD:	Number of the included members

7.6 Partitioning a Logical Drive/Logical Volume

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.

Please note that partitioning can be very useful when dealing with a very large capacity; however, partitioning a logical drive or logical volume is not a requirement for RAID configuration.

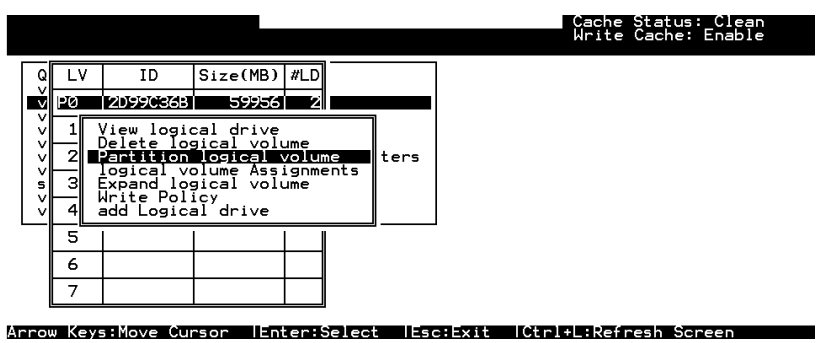

Choose the logical volume you wish to partition, then press **[ENTER]**. Choose "Partition logical volume", then press **[ENTER]**. Select from the list of undefined partitions and press **[ENTER]**.

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

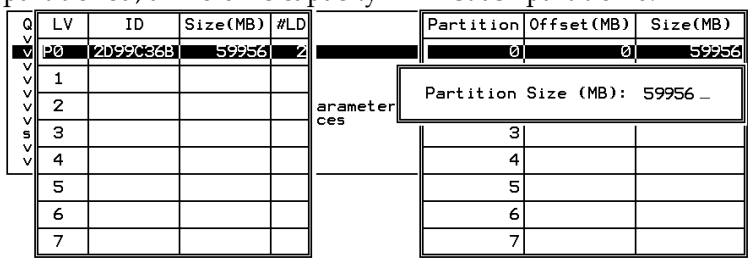

Press **[ENTER]** and type the desired size for the selected partition, and then press **[ENTER]** to proceed. The remaining size will be automatically allotted to the next partition.

Choose **Yes** to confirm when prompted to the "Partition Logical Volume?" message. Press **[ENTER]** to confirm. Follow the same procedure to partition the remaining capacity of your logical volume.

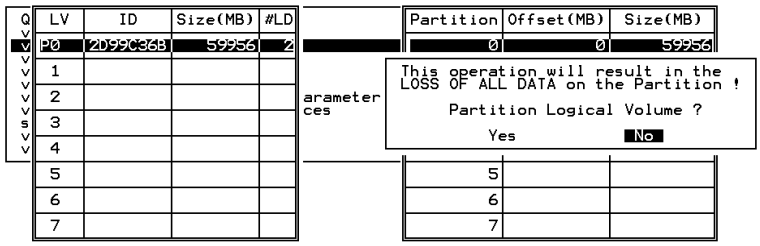

When a partition of a logical drive/logical volume is deleted, the capacity of the deleted partition will be added to the last partition.

WARNING!

- *Whenever a partition is changed, it is necessary to re-configure all host LUN mappings. All data in it will be lost and all the host LUN mappings will be removed with any change to partition capacity.*
- *If operated with a Unix-based system, reset the subsystem for the configuration changes to take effect if any changes have been made to partition sizes and partition rearrangement.*

7.7 Mapping a Logical Volume to Host LUNs

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

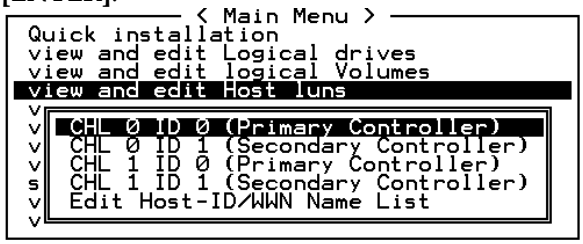

A list of host channel/ID combinations appears on the screen. The diagram above shows two host channels and each is designated with both a primary and a secondary ID.

Multiple IDs on host channels are necessary for redundant controller configuration. Details on creating multiple IDs and changing channel modes will be discussed later. Choose a host ID by pressing **[ENTER]**.

Several details are noticeable here:

1. A logical group of drives (logical drive/logical volume) previously assigned to the primary controller cannot be mapped to a secondary ID. Neither can those assigned to the secondary controller be mapped to a primary ID.
2. For a SCSI-based controller, ID 7 is reserved for the controller itself. If there are two controllers, controllers might occupy ID6 and ID7, or ID8 and ID9. Please check your system Hardware Manual for details on reserved IDs.

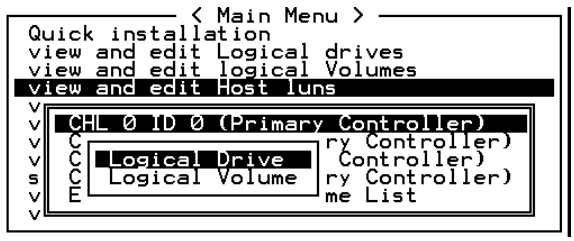

Choose the "channel-ID" combination you wish to map, then press **[ENTER]** to proceed. Choose mapping a "Logical Drive" or a "Logical Volume" on the drop box.

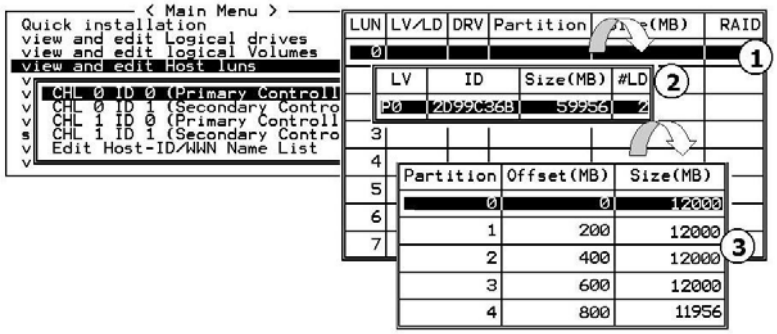

1. A list of LUN entries and their respective mappings will be displayed. To map a host LUN to a logical volume's partition, select an available LUN entry (one not mapped yet) by moving the cursor bar to the LUN, then press **[ENTER]**.

2. A list of available logical volumes displays. Move the cursor bar to the desired logical unit, then press **[ENTER]**.

3. A list of available partitions will prompt. Move the cursor bar to the desired partition, then press **[ENTER]**. If you have not partitioned the logical volume, the whole capacity will be displayed as one logical partition.

4. When prompted to "Map Host LUN," press **[ENTER]** to proceed.

5. When prompted to "Map Logical Volume?," select **Yes** to continue.

A prompt will display the mapping you wish to create. Choose **Yes** to confirm the LUN mapping you selected.

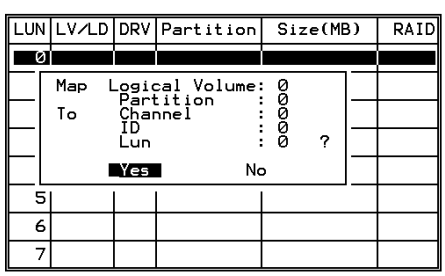

The detail in the confirm box reads: partition 0 of logical volume 0 will map to LUN 0 of SCSI ID 0 on host channel 0.

Continue to map other partitions to host LUNs.

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

If your controller has not been configured with a host channel and assigned with SCSI ID, please move on to Section ***7.12 Viewing and Editing Channels***."

7.8 Assigning Spare Drive, Rebuild Settings

Adding Local Spare Drive

A spare drive is a standby drive automatically initiated by controller firmware to replace a failed drive. A spare drive must have an equal or larger capacity than the array members. A local spare should have a capacity equal to or larger than the members of the logical drive it is assigned to. A global spare should have a capacity equal to or larger than all physical drives in a RAID system.

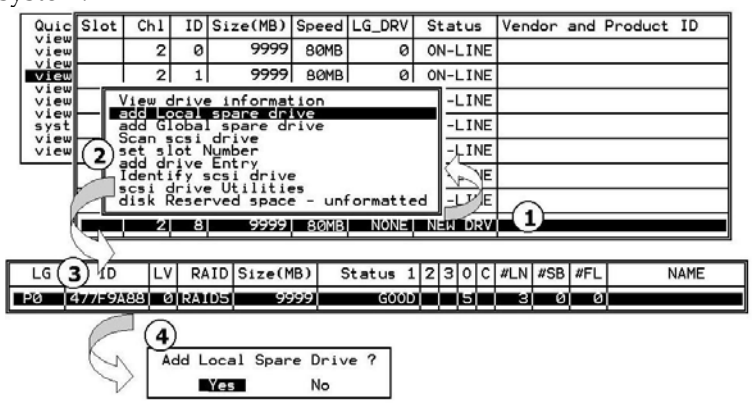

Choose "View and Edit Drives" on the Main Menu, then press **[ENTER]**. Move the cursor bar to a drive that is not assigned to a logical drive or as a spare drive (usually indicated as a "New Drive"), and then press **[ENTER]**.
Choose "Add Local Spare Drive" and press **[ENTER]**. A list of logical drives displays.
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.
When prompted to "Add Local Spare Drive?", choose **Yes** to confirm.

Adding a Global Spare Drive

A global spare replaces the failed drive in any logical drive of a RAID system.

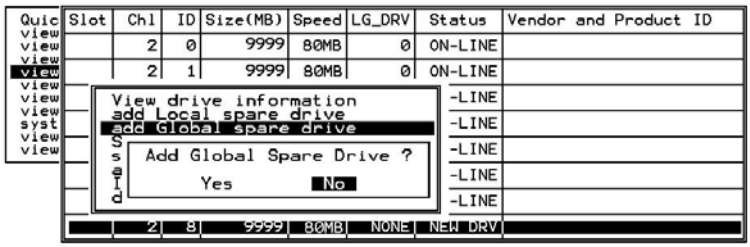

Move the cursor bar to the drive that is not a member drive or a spare (usually indicated as a "New Drive"), and then press **[ENTER]**. Choose "Add Global Spare Drive." When prompted to "Add Global Spare Drive?", choose **Yes**.

7.9 Viewing and Editing Logical Drive and Drive Members

Choose "View and Edit Logical Drives" in the Main Menu. The array status will be displayed. 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]**.

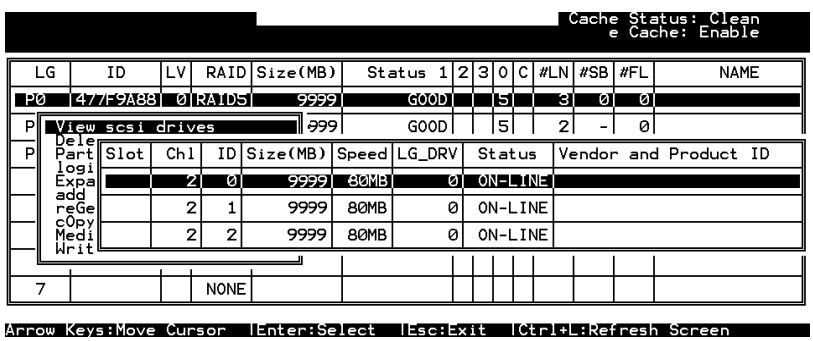

Choose "View Drives." Drive member information will be displayed on the screen.

Deleting a Logical Drive

Choose the logical drive you wish to delete, then press **[ENTER]**. Choose "Delete logical drive." Choose **Yes** when prompted to confirm.

Deleting a Partition of a Logical Drive

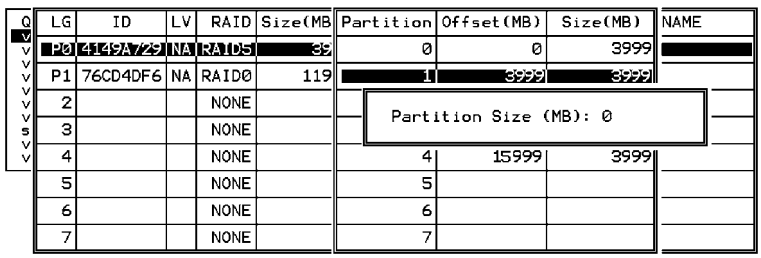

Choose the logical drive which has a partition you wish to delete, then press **[ENTER]**. Choose "Partition logical drive." Partitions of the logical drive will be displayed in tabulated form. Move the cursor bar to the partition you wish to delete, then press **[ENTER]**.

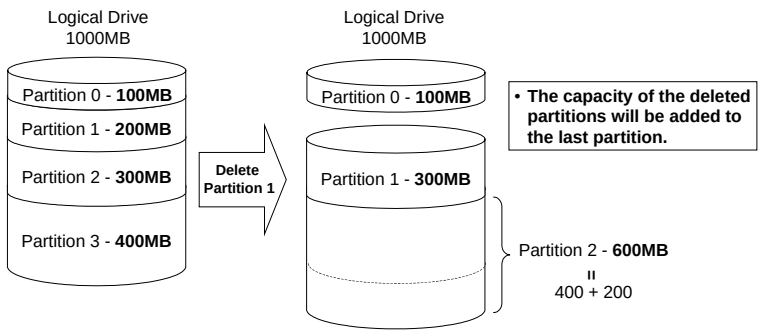

Enter "0" on the partition size to delete the partition.

Figure 7 - 1 Drive Space Allocated to the Last Partition

As illustrated above, the capacity of the deleted partition will be added to the last partition.

WARNING!

- *Whenever a partition is changed, it is necessary to reconfigure all host LUN mappings. All data kept in the partition and the host LUN mappings will be removed with any partition change.*

Assigning a Name to a Logical Drive

Naming can help identify different arrays in a multi-array configuration. This function is also useful in special situations. For example, when one or more logical drives have been deleted, the array indexing is changed after system reboot. The second logical drive might become the first on the list.

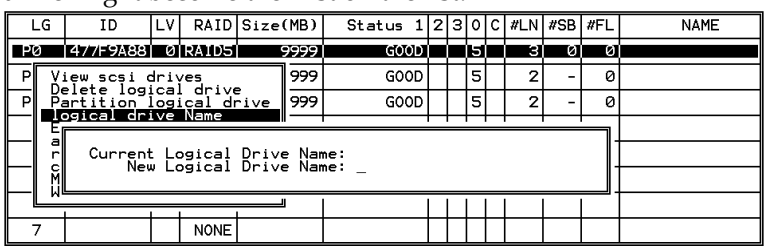

Choose the logical drive for which you wish to assign a name, then press **[ENTER]**. Choose "logical drive name," then press **[ENTER]** again. The current name will be displayed. You may now enter a new name in this field. Enter a name, then press **[ENTER]** to save the configuration.

Rebuilding a Logical Drive

If there is no spare drive in the system, a failed drive should be immediately replaced by a drive known to be good. Once the failed drive is replaced, the rebuild process can be manually initiated.

If you want the controller to auto-detect a replacement drive, make sure you have the following items set to "enabled":

1. Periodic Drive Check Time
2. Periodic Auto-Detect Failure Drive Swap Check Time

These two configuration options can be found under "View and Edit Configuration Parameters" -> " Drive-Side SCSI Parameters".

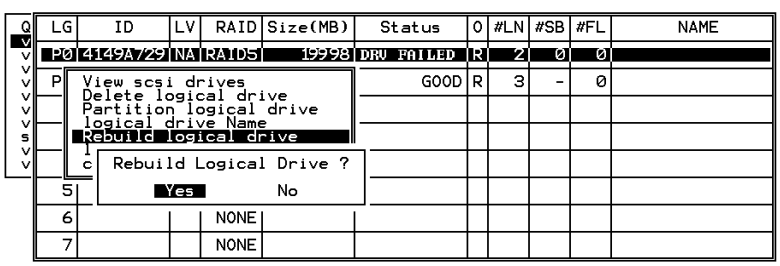

Choose the logical drive that has a failed member drive, then press **[ENTER].** Choose "Rebuild logical drive", then press **[ENTER].** When prompted to "Rebuild Logical Drive?," select **Yes**.

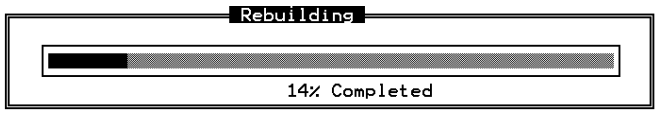

The rebuild progress will be displayed.

When rebuild has already started, choose "Rebuild progress" to see the rebuilding progress.

IMPORTANT!

- *The Rebuild function is only available when a logical drive (with RAID level 1, 3 or 5) has a failed member. NRAID and RAID 0 configurations provide no data redundancy.*

Regenerating Logical Drive Parity

(Applies to RAID Levels 1, 3, and 5)

If no verifying method is applied to data writes, this function can be often performed to verify parity blocks of a selected array. This function compares and recalculates parity data to correct parity errors.

LG	ID	LV	RAID	Size(MB)	Status 1	2	3	0	C	#LN	#SB	#FL	NAME
P0	4295529B	NA	RAID5	476000	GOOD			4	B	3	0	0	
	Execute Regenerate Logical Drive Parity Overwrite Inconsistent Parity - Enabled Generate Check Parity Error Event - Enabled												
3			NONE										
4			NONE										
5			NONE										
6			NONE										
7			NONE										

Choose the logical drive that you want to regenerate the parity for, and then press **[ENTER].** Choose "Regenerate Parity," then press **[ENTER]**. When prompted to "Regenerate Parity?", select **Yes**.

Please refer to ***Chapter 11 Data Integrity*** for more information on Parity Regeneration.

Media Scan

Media Scan is used to examine drives and is able to detect the presence of bad blocks. If any data blocks have not been properly

committed and are found during the scanning process, data from those blocks are automatically recalculated, retrieved and stored onto undamaged sectors. If bad blocks are encountered on yet another drive during the rebuild process, the block LBA (Logical Block Address) of those bad blocks will be shown. If rebuild is carried out under this situation, rebuild will continue with the unaffected sectors, salvaging the majority of the stored data.

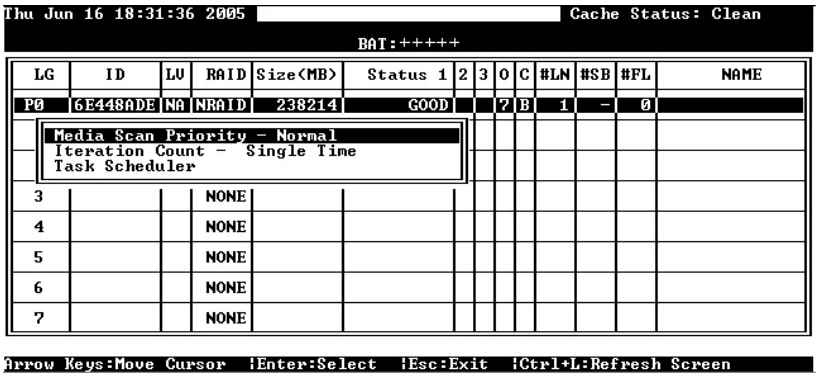

There are two options with performing the Media Scan:

1. **Media Scan Priority:** determines how much system resources will be used for the drive scanning and recalculating process.

2. **Iteration Count**:

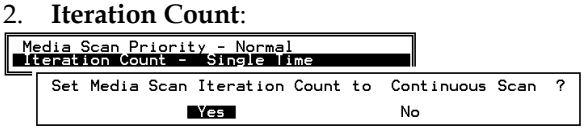

The iteration setting determines how many times the scan is performed. If set to "continuous," the scan will run in the background continuously until it is stopped by user.

System can automatically perform a media scan according to a preset task schedule. For more details, please refer to **Chapter 11 Data Integrity.**

Write Policy

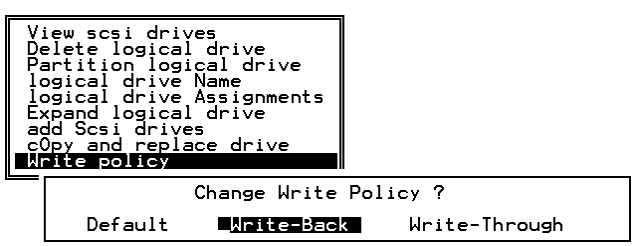

The Write-Back cache setting is configurable on the per array basis. Setting to the default value means the array setting is coordinated with the controller's general setting. The controller's general setting option can be found in "View and Edit Configuration Parameters" -

> "Caching Parameters" -> "Write-Back Cache". Note that cached writes are lost if a power failure occurs.

7.10 Viewing and Editing Host LUNs

Viewing or Deleting LUN Mappings

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

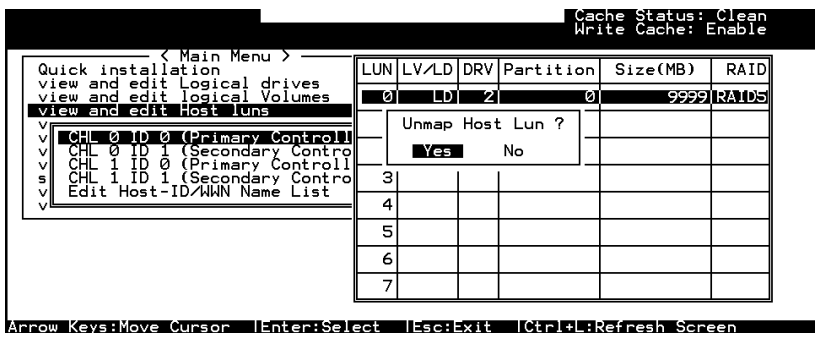

A list of the current LUN mapping will be displayed on the screen. Move the cursor bar to the LUN mapping you wish to delete, then press **[ENTER].** Select **Yes** to delete the LUN mapping, or **No** to cancel.

Edit Host-ID/WWN Name List

This is a specific item used for systems communicating over Fibre host loops. Please refer to ***Chapter 8 Fibre Operation*** for more details.

Pass-through SCSI Commands

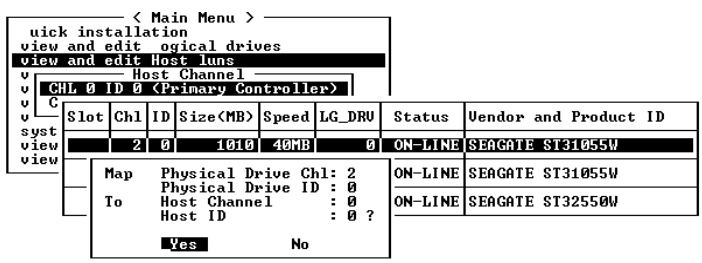

If you have primary and secondary controllers, move the cursor to the controller for the device that you wish to map, then press

[**ENTER**]. You will be prompted to map a SCSI ID to a physical drive.

WARNING!

- *Pass-through SCSI Commands are only intended to perform maintenance functions for a drive or device on the drive side.* ***Do not*** *perform any destructive commands to a disk drive (i.e., any commands that write data to a drive media). This will result in inconsistent parity among drives included in a logical configuration of drives. If a disk drive is a spare drive or a member of a logical drive, such a destructive command may cause a data inconsistency.*
- *When a drive/device is mapped to a host SCSI ID so that Pass-through SCSI Commands can be used, the data on that drive/device will not be protected by the controller. Users who employ Pass-through SCSI Commands to perform any write commands to drive media do so at their own risk.*

7.11 Viewing and Editing Drives

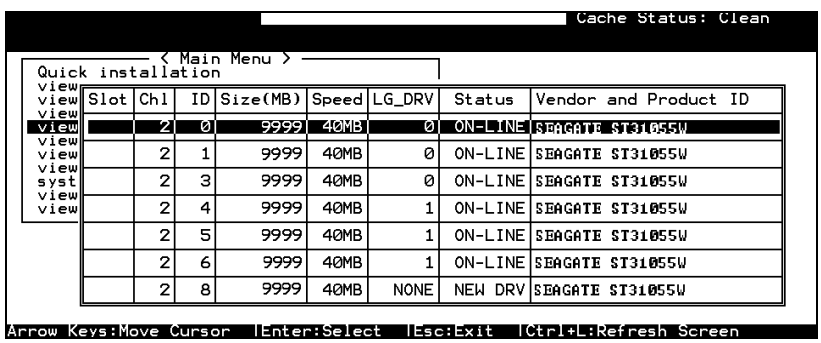

Choose "View and Edit Drives" in the Main Menu. All drives attached to the drive channels will be displayed on the screen.

Scanning New Drive

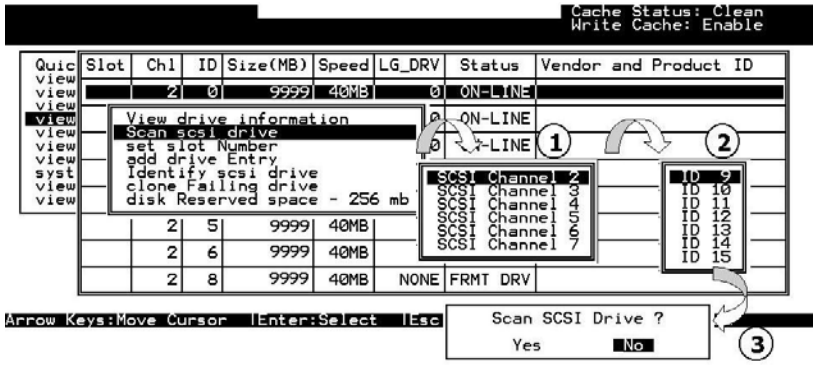

If a drive is connected after the array is started, choose a drive and press **[ENTER].** Choose "Scan drive", then press **[ENTER].** The menu may vary according to the drive status. Choose the drive channel and ID of the drive you wish to scan, then press **[ENTER].**

Identifying Drive (not available with SATA Drives)

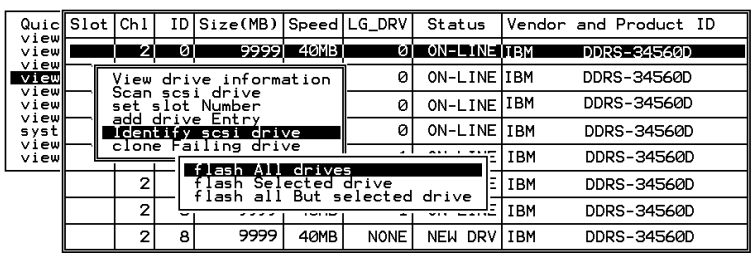

Move the cursor bar to the drive you wish to identify, then press **[ENTER].** Choose "Identify drive," then choose "flash all drives" to flash the read/write LEDs of the all the drives in the drive channel. Choose **Yes**.

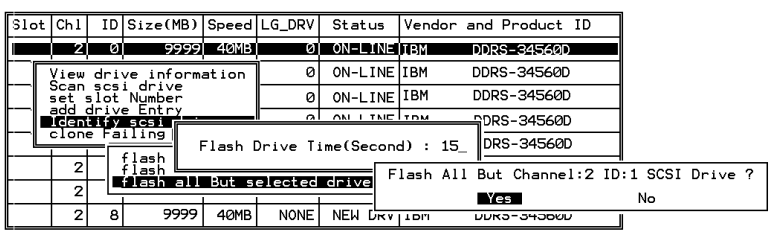

You may also choose "**flash selected drive**" or "**flash all But Selected drives**" to flash the read/write LEDs of the selected drives only, or all the drives except the selected drive. Choose **Yes** and choose an extent of time from 1 to 999 seconds.

Deleting Spare Drive (Global / Local Spare Drive)

Move the cursor to a Local Spare Drive or Global Spare Drive, then press **[ENTER].** Choose "Delete Global/Local Spare Drive," then press **[ENTER]** again. Choose **Yes** to confirm.

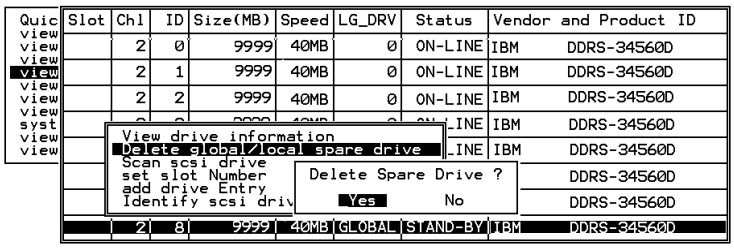

NOTE:

- *The spare drive you deleted or any drive you replaced from a logical unit will be indicated as a "used drive."*

7.12 Viewing and Editing Channels

Except for those shipped in a dual-redundant chassis, SCSI-based controllers use channel 0 as the host channel and also as the communications path between controllers. If a redundant controller configuration is preferred, you may need to assign other channels as host. Flexibility is added so that all channels can be configured as host or drive.

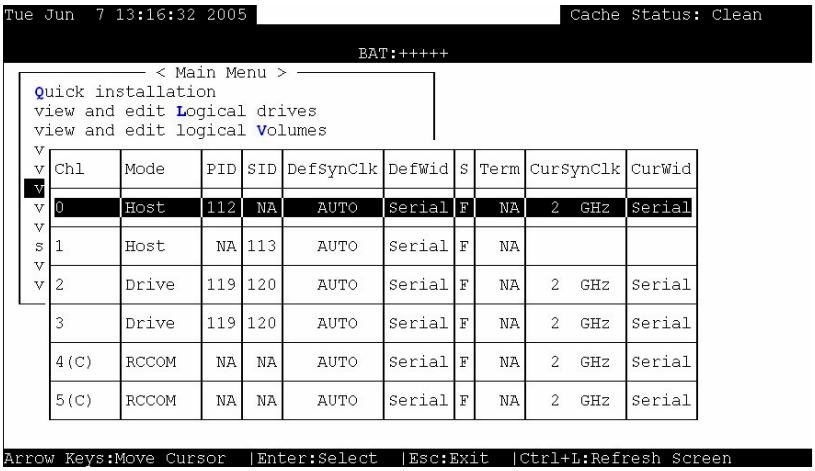

Chl	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid
0	Host	112	NA	AUTO	Serial	F	NA	2 GHz	Serial
1	Host	NA	113	AUTO	Serial	F	NA		
2	Drive	119	120	AUTO	Serial	F	NA	2 GHz	Serial
3	Drive	119	120	AUTO	Serial	F	NA	2 GHz	Serial
4(C)	RCCOM	NA	NA	AUTO	Serial	F	NA	2 GHz	Serial
5(C)	RCCOM	NA	NA	AUTO	Serial	F	NA	2 GHz	Serial

Choose "View and Edit Channels" in the Main Menu. Channel status displays.

Due to the different characteristics of host interfaces, the configuration options for SCSI or Fibre host channels are different.

Fibre host channel

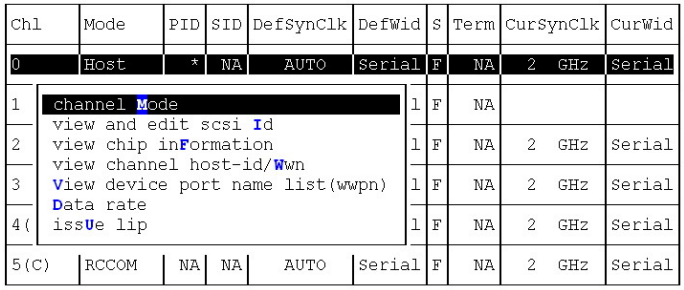

Chl	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid
0	Host	*	NA	AUTO	Serial	F	NA	2 GHz	Serial
1						l	F	NA	
2						l	F	NA	2 GHz	Serial
3						l	F	NA	2 GHz	Serial
4(						l	F	NA	2 GHz	Serial
5(C)	RCCOM	NA	NA	AUTO	Serial	F	NA	2 GHz	Serial

SCSI host channel

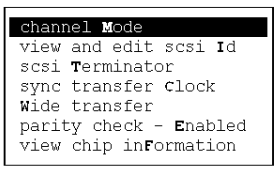

Ch		Clk	DefWid	S	Term	CurSynClk	CurWid
0	Host 0 NA	160.0MHz	Wide	L	On	Async	Narrow
1	Host 0 NA	160.0MHz	Wide	L	On	Async	Narrow
2	Drive 7 NA	160.0MHz	Wide	L	On		
3	Drive 7 NA	160.0MHz	Wide	L	On		

Redefining Channel Mode

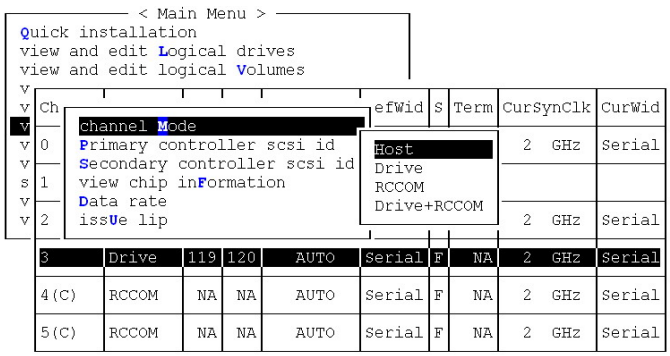

For Fibre- and SCSI-based controllers, channels can be operated in host or drive mode. Choose the channel you wish to change, then press **[ENTER]**. Choose "Channel Mode," then press **[ENTER]**. A dialog box will appear asking you to confirm the change.

IMPORTANT!

- *Every time you change the channel mode, you must reset the controller for the change to take effect.*
- *Only the Fibre channel subsystems feature the RCCOM options. Most EonStor subsystems have preset RCC channels that can not be changed.*

Viewing and Editing IDs - Host Channel

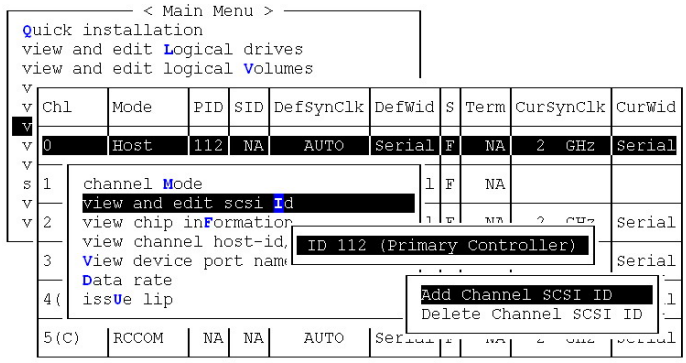

Choose a host channel, then press **[ENTER]**. Choose "View and Edit SCSI ID." A list of existing ID(s) will be displayed on the screen. Select and press **[ENTER]** on one of the IDs. You may then choose to add or delete an existing ID.

Shown below is the display screen of a SCSI host subsystem.

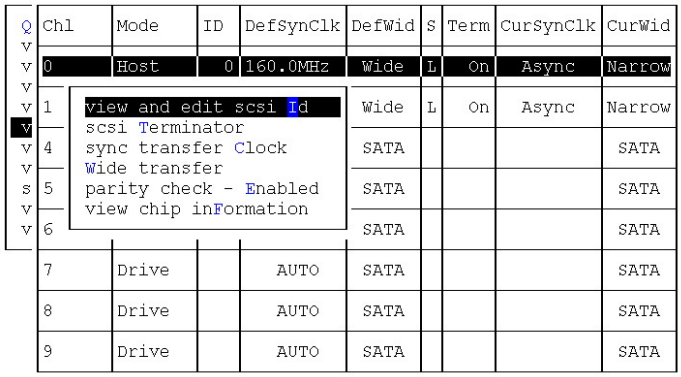

Viewing and Editing IDs

Adding an ID (Primary/Secondary Controller ID)

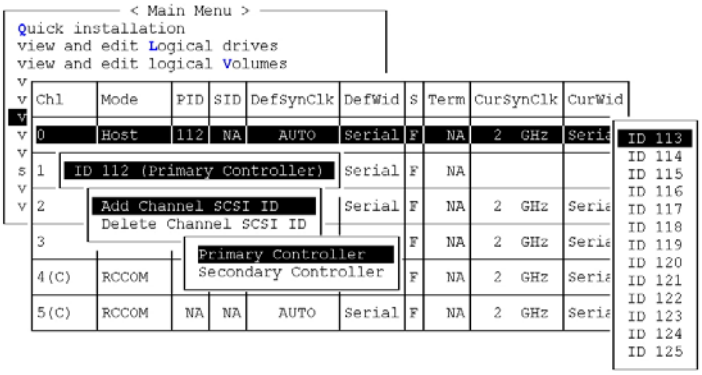

In single controller mode, you should set the secondary controller's ID to "NA" (usually unavailable in a single-controller subsystem). In dual-controller mode, you may need to set an ID for the secondary controller on each of your drive channels.

Press **[ENTER]** on one of the existing IDs. Choose "Add Channel SCSI ID," then choose to assign an ID for either the "Primary Controller" or "Secondary Controller." A list of SCSI IDs will appear. Choose a SCSI ID. **DO NOT** choose a SCSI ID used by another device on the same channel. The defaults are PID=8 and SID=9 (SCSI channel). In redundant mode, logical drives associated with a primary ID will be managed by the primary controller, a logical drive associated with a secondary ID the secondary controller.

Deleting an ID

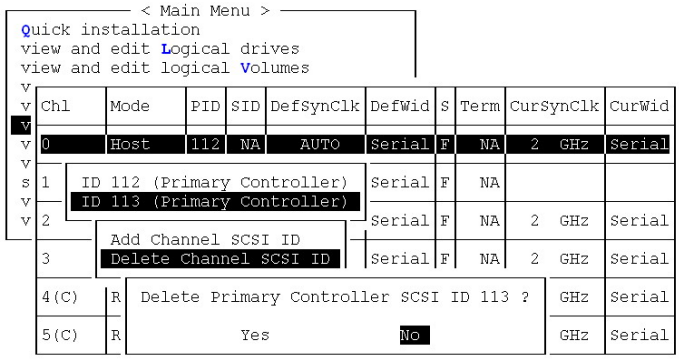

Choose the host bus ID you wish to delete. Choose "Delete Channel SCSI ID." The dialog box "Delete Primary/Secondary Controller SCSI ID#?" will appear. Select **Yes**, then press **[ENTER]** to confirm.

IMPORTANT!

- *Every time you change a channel ID, you must reset the subsystem/controller for the changes to take effect.*
- *The default SCSI ID of the primary controller (single controller configuration) on a host channel is 0, on a Drive channel is 7.*
- *If only one controller exists, you must set the secondary controller's ID to "NA." If a secondary controller exists, you need to set a secondary ID on host and drive channels.*
- *Multiple target IDs can be applied to a host channels while each drive channel has only one or two IDs (in redundant mode).*
- *At least one controller's ID has to be present on each channel bus.*

Setting a Primary Controller's ID - Drive Channel

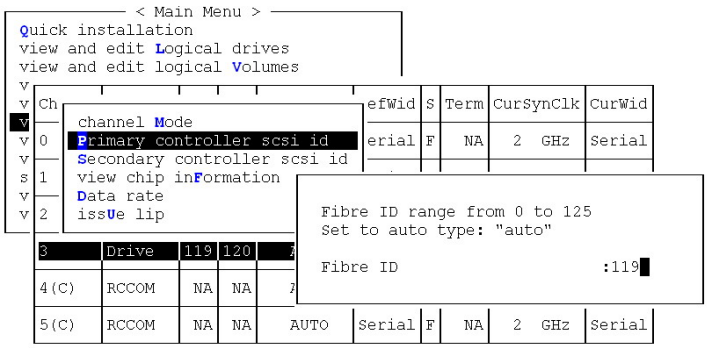

Choose a drive channel, then press **[ENTER]**. Choose "Primary Controller ID." Manually enter an ID or key in "auto". A dialog box displays. Select **Yes** by pressing **[ENTER]** if you want to reset the subsystem for the configuration to take effect immediately. If you select **No,** remember to reset the subsystem later.

For more details on ID settings in redundant mode, please refer to ***Chapter 10***.

NOTE:

- *It is unnecessary to assign IDs on the drive channels of SATA-based subsystems. Either one of the RAID controllers in a dual-controller SATA-based subsystem can access a disk drive through a common backplane and the MUX board behind each drive.*
- *Shown below is the drive channel options of a SATA-based subsystem:*

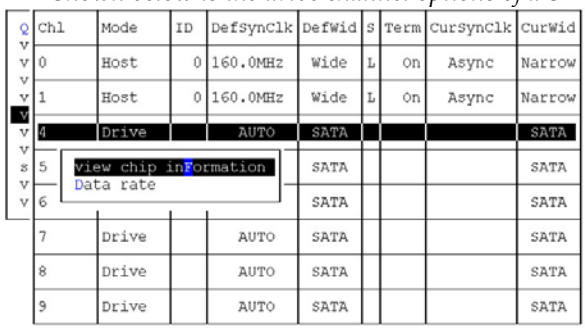

Setting a Secondary Controller's ID - Drive Channel

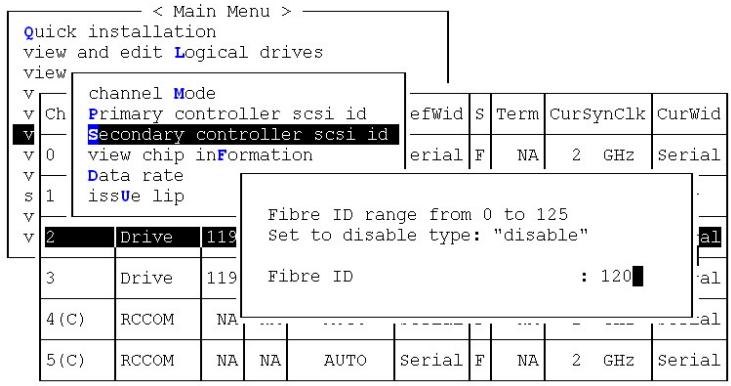

Choose a drive channel, then press **[ENTER].** Choose "Secondary Controller ID." A list of channel IDs displays. Assign an ID to the chip processor managing the secondary controller's drive channel. Key in an ID. A dialog box appears. Select **Yes** by pressing **[ENTER]** if you want to reset the subsystem for the configuration to take effect. If you select **No,** remember to reset the subsystem later.

NOTE:

The EonStor subsystems have preset IDs for the drive channels and there is no need to change those IDs. For all subsystems that come with drives attached through back-end PCBs, system vendors should have configured the drive IDs and there is usually no need to make any changes.

Setting Channel Terminator

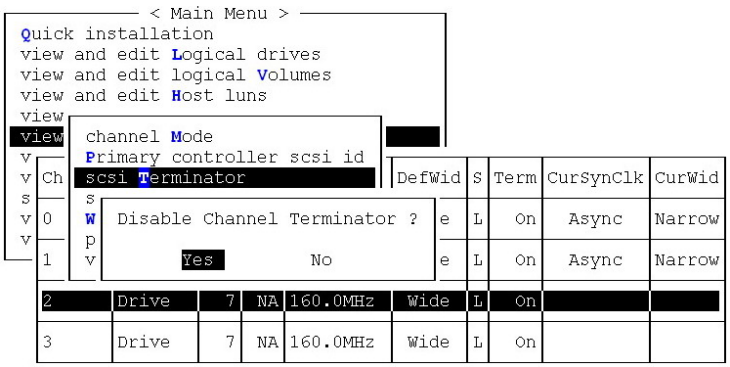

Choose the channel for which you wish the terminator enabled or disabled, then press **[ENTER].** Choose "SCSI Terminator", then press **[ENTER].** A dialog box will appear. Choose **Yes**, then press **[ENTER].** The terminator can also be enabled by switch jumpers. Please refer to the controller hardware manual for more details.

Setting a Transfer Speed (SCSI Channel)

Drive Channel

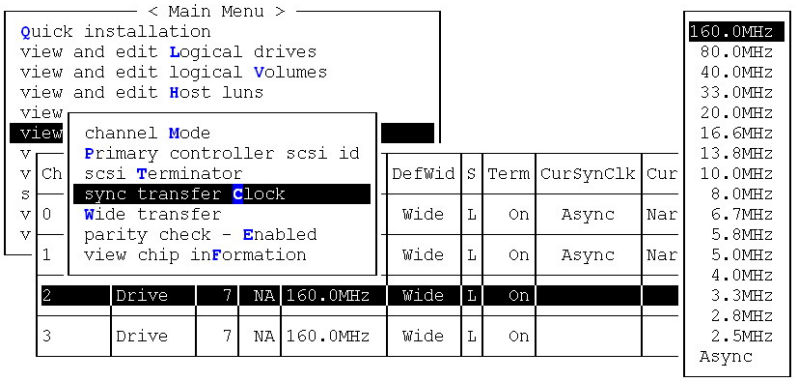

Host Channel

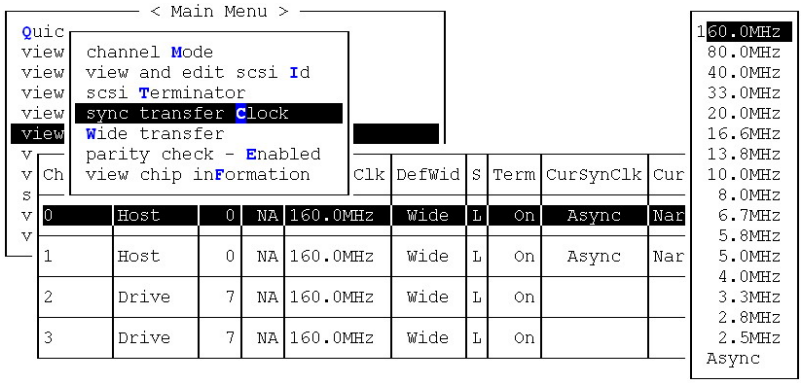

Move the cursor bar to a channel, then press **[ENTER].** Choose "Sync Transfer Clock", then press **[ENTER].** A list of the clock speeds will appear. Move the cursor bar to the desired speed and press **[ENTER]**. A dialog box "Change Sync Transfer Clock?" will appear. Choose **Yes** to confirm.

IMPORTANT!

Transfer rate	Transfer speed in Mhz
SCSI-320	160Mhz
SCSI-160	80Mhz

- *According to SCSI Channel Specifications, transfer speed shown in Mhz is the SCSI bus sync frequency. With 16-bit Wide transfer, the actual data rate is the double of Mhz transfer speed.*
- *Every time you change the Transfer Speed, you must reset the controller for the changes to take effect.*

Setting the Transfer Width

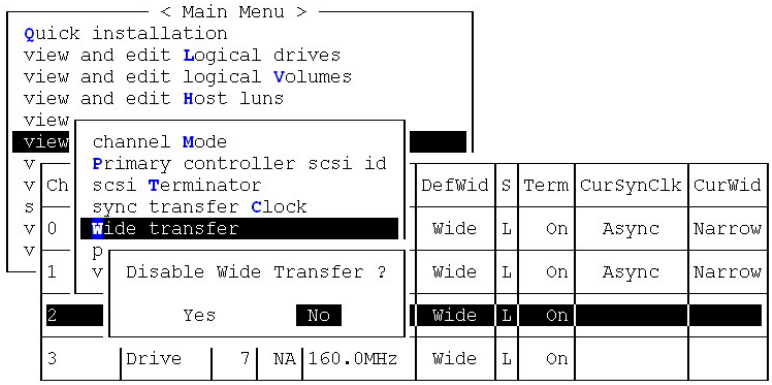

Move the cursor bar to a channel, then press **[ENTER].** Select "Wide Transfer," then press **[ENTER]**. A dialog box "Disable Wide

Transfer?" or "Enable Wide Transfer?" will appear. Choose **Yes** to confirm.

IMPORTANT!

- *Every time you change the SCSI Transfer Width, you must reset the controller for the changes to take effect.*

Viewing and Editing Drive Channel

Data Rate (Channel Bus)

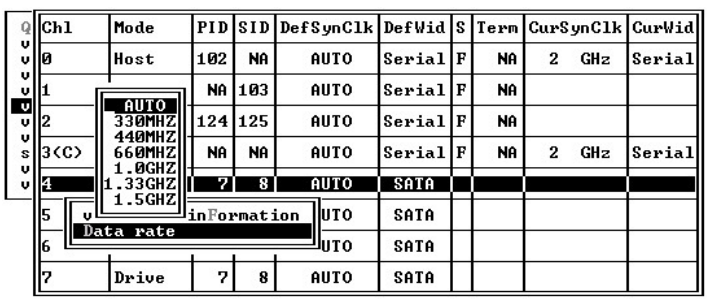

This option is available in the configuration menu of Fibre host channel and the drive channel configuration menus of Fibre-, ATA-, or SATA-based subsystems. Default is "AUTO" and should work fine with most disk drives. Changing this setting is not recommended unless some particular bus signal issues occur.

The host channel data rate setting allows a 2Gbit Fibre Channel to negotiate with devices communicating over 1GHz link if the link speed configuration is determined externally.

Most SATA/ATA-based systems connect only one drive per SATA/ATA channel (4 for multi-lane with SATA-II). This helps to avoid a single drive failure from affecting other drives. The maximum mechanical performance of today's drives can reach around 30MB/second (sustained read). This is still far below the bandwidth of a drive channel bus. Setting the SATA/ATA bus speed to a lower value can get around some problems, but will not become a bottleneck to system performance.

Note that the SATA/ATA speed is the maximum transfer rate of the SATA/ATA bus in that mode. It does not mean the drive can actually carry out that amount of sustained read/write performance. For the performance of each drive model, please refer to the documentation provided by drive manufacturer.

Issue LIP

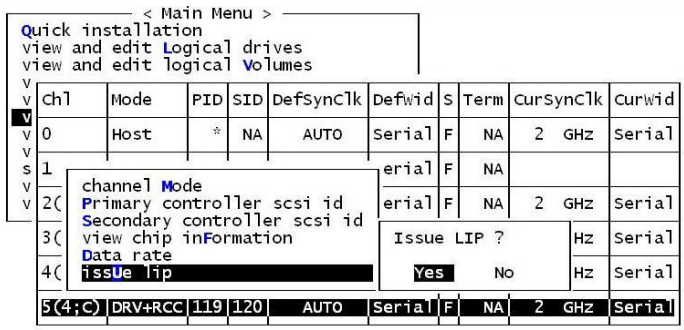

This option allows you to manually issue a LIP (Loop Initialization Sequence) on either a host or drive channel. This function is particularly useful when cabling failure occurs between two cascaded drive enclosures.

The condition for applying the "Issue LIP" function is described as follows (see diagram Figure 7-2):

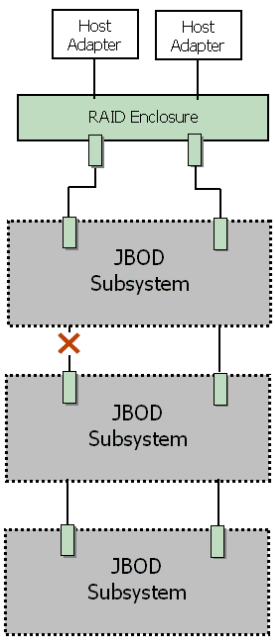

1. Multiple enclosures have been cascaded using FC-AL Fibre links.
2. One cable link fails or is disconnected between enclosures.
3. Cabling failure is restored.
4. Since the FC port to which the failed link was attached to has already been bypassed at the time when the fault occurred, a LIP command must be manually exerted by the user for the FC port to participate in the loop again. The dual-loop FC-AL can thus be restored.

Figure 7 - 2 Cabling Failure between Drive Enclosures

7.13 System Functions

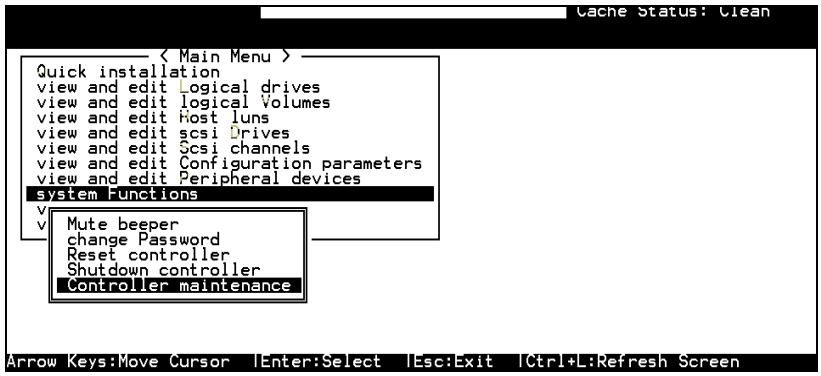

Choose "System Functions" in the Main Menu, then press **[ENTER].** The System Functions menu displays. Move the cursor bar to an item, then press **[ENTER].**

Mute Beeper

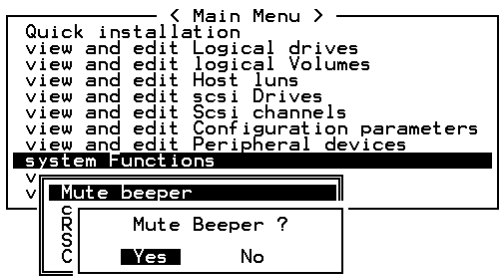

When the controller's beeper has been activated, choose "Mute beeper," then press **[ENTER].** Choose "**Yes**" and press **[ENTER]** in the next dialog box to turn the beeper off temporarily for the current event. The beeper will still be activated on the next event.

Change Password

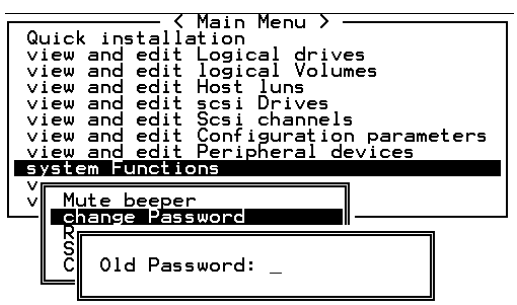

Use the controller's password to protect the array from unauthorized entry. Once the controller's password has been set, regardless of whether the front panel, the RS-232C terminal

interface or RAIDWatch Manager is used, you can only access the RAID controller by providing the correct password.

IMPORTANT!

- *The controller verifies the password when entering the Main Menu from the initial screen or when making a configuration change. If the controller is going to be left unattended, the "Password Validation Timeout" can be set to "Always Check." Setting the validation timeout to "Always Check" will protect the controller configuration from any unauthorized access.*

- *The controller password and controller name share a 16-character space. The maximum characters for the controller password is 15. If the controller name occupies 15 characters, there is only one character left for the controller password, and vice versa. Since 3.42J09, later firmware revisions will support a 32-character space.*

Changing the Password

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 can not be replaced unless a 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 go back to the previous menu.

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

Setting a New Password

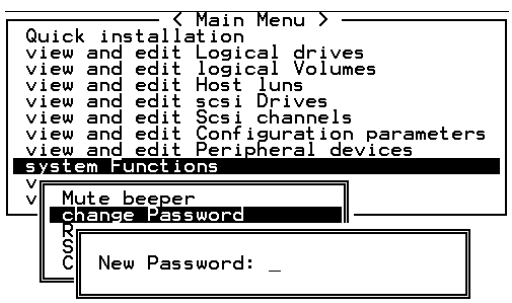

Enter the desired password in the column, then press **[ENTER].** The next dialog box will display "Re-Enter Password". Enter the password again to confirm and press **[ENTER]**.

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.

Disabling the Password

To disable or delete the password, press **[ENTER]** in the empty column that is used for entering a new password. The existing password will be deleted. No password checking will occur when entering the Main Menu or when making a configuration change.

Reset Controller

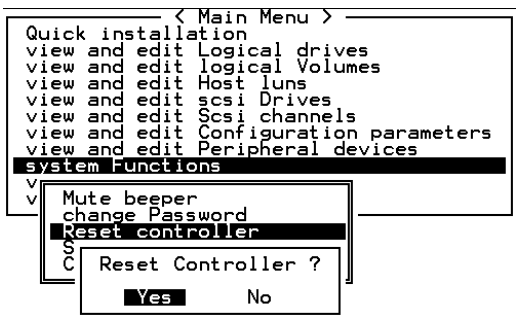

To reset the controller without powering off the system, move the cursor bar to "Reset Controller," then press **[ENTER]**. Choose **Yes** in the dialog box that follows, then press **[ENTER].** The controller will now reset as well as power-off or re-power-on.

Shutdown Controller

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.

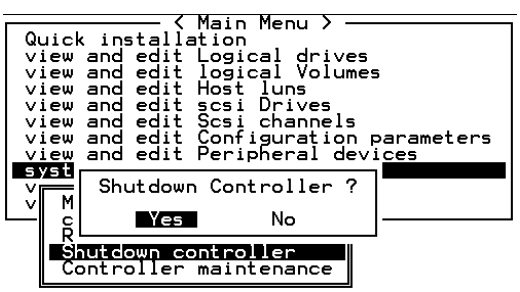

For "Controller Maintenance" functions, such as "Download Firmware," please refer to ***Appendix B***.

Adjust LCD Contrast

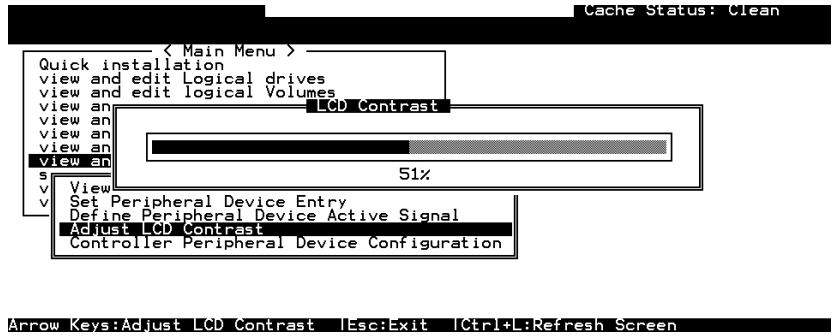

The controller LCD contrast is set at the factory to a level that should be generally acceptable. The controller is equipped with an LCD contrast adjustment circuit, allowing the contrast to be adjusted either via the RS-232 Terminal Emulation Menus or by the LCD User Interface.

7.14 Controller Parameters

Controller Name

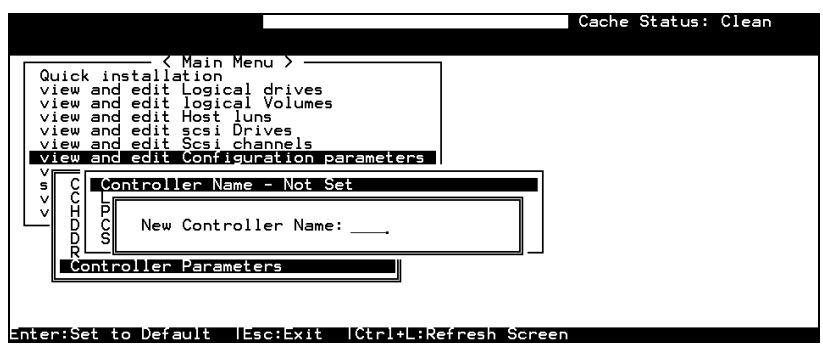

Choose "View and Edit Configuration Parameters," "Controller Parameters," then press **[ENTER]**. The current name displays. Press **[ENTER]**. Enter a name in the dialog box that prompts, then press **[ENTER]**.

LCD Title Display - Controller Name

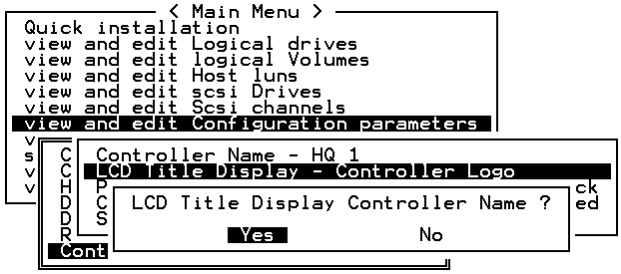

Choose "View and Edit Configuration Parameters," "Controller Parameters," then press **[ENTER]**. Choose to display the embedded controller logo or any given name on the LCD. Giving a specific name to each controller will make them easier to identify if you have multiple RAID systems remotely monitored.

Saving NVRAM to Disks

You can choose to backup your controller-dependent configuration information to disks. We recommend using this function to save configuration information whenever a configuration change is made. The information will be duplicated and distributed to all logical configurations of drives.

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].**

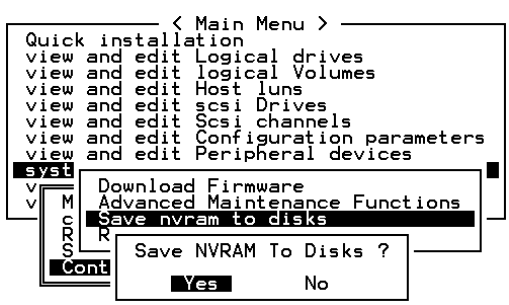

Choose **Yes** to confirm.

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

Restore NVRAM from Disks

When you want to restore your NVRAM information from what you previously saved onto disk, use this function to restore the configuration information.

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].**

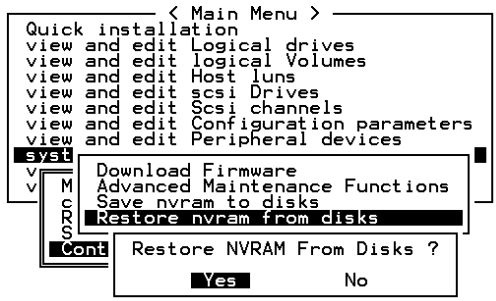

Press **Yes** to confirm.

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

Password Validation Timeout

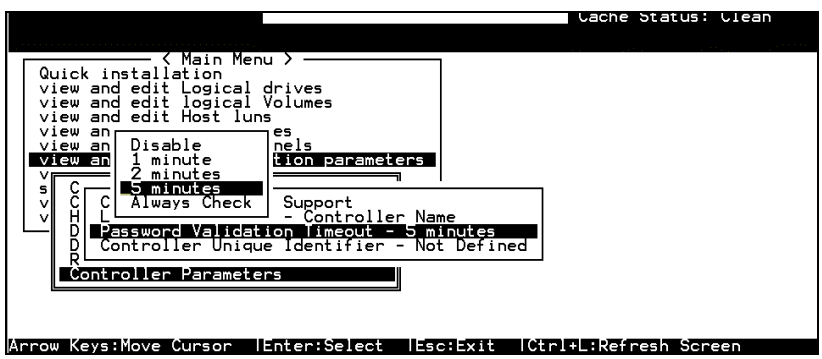

Choose "View and Edit Configuration Parameters," "Controller Parameters," then press **[ENTER]**. Select "Password Validation Timeout," and press **[ENTER]**. Choose to enable a validation timeout from one minute to Always Check. The Always Check timeout will disable any configuration change made without entering the correct password.

Controller Unique Identifier

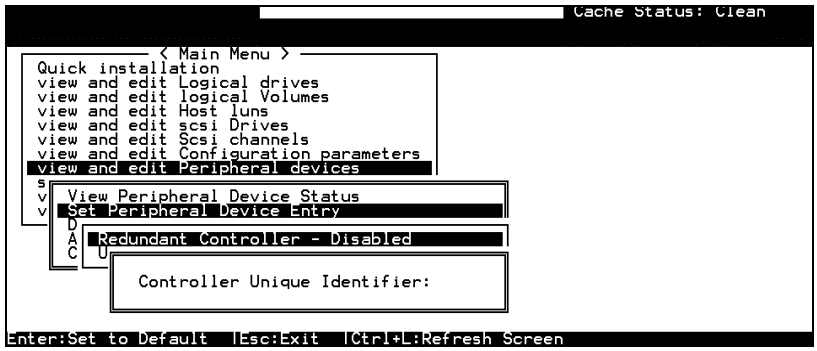

Enter any hex number between "0" and "FFFFF" for the unique identifier. The value you enter ***MUST*** be different for each controller.

The Controller Unique Identifier is ***required*** for configuring every RAID controller. The controller automatically notifies users to enter a unique identifier when the first logical drive is created in a dual-controller system.

Enter a unique ID for any RAID controller whether it is configured in a single or dual-controller configuration. The unique ID is necessary for the following:

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

2. The unique ID is generated into a Fibre Channel WWN node name for controllers or RAID systems using Fibre Channel host ports. The node name prevents host computers from mis-addressing the storage system during the controller failover/failback processes.

3. MAC addresses for the controller's Ethernet port that will be taken over by a surviving controller in the event of single controller failure.

In redundant mode, configuration data is synchronized between controllers. Host ports on the partner controllers appear with the same node name but each has a different port name (WWPN).

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

The unique identifier setting can be accessed from "View and Edit Configuration Parameters" ∏ "Controller Parameters" ∏ "Controller Unique ID."

Set Controller Date and Time

This sub-menu only appears when the controller is equipped with a real-time clock.

Time Zone

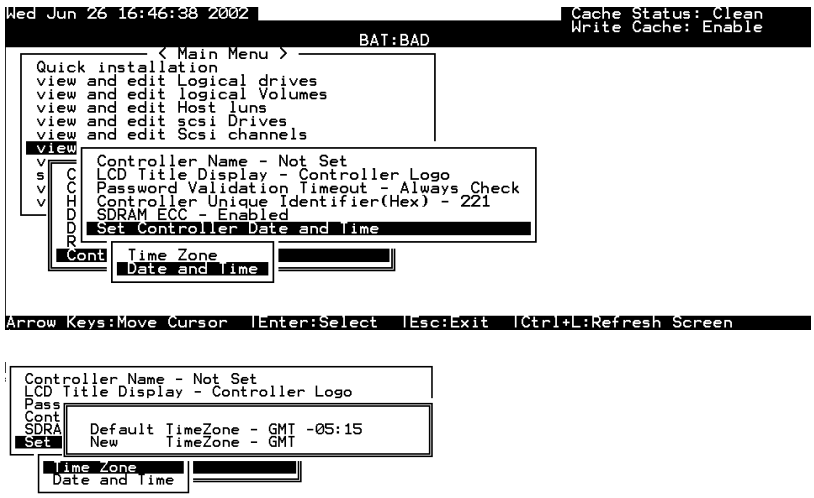

The controller uses GMT (Greenwich Mean Time), a 24-hours clock. To change the clock to your local time zone, enter the hours later than the Greenwich mean time following a plus (+) sign. For example, enter "+9" for Japanese time zone.

Date and Time

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

7.15 Drive Information

View Drive Information

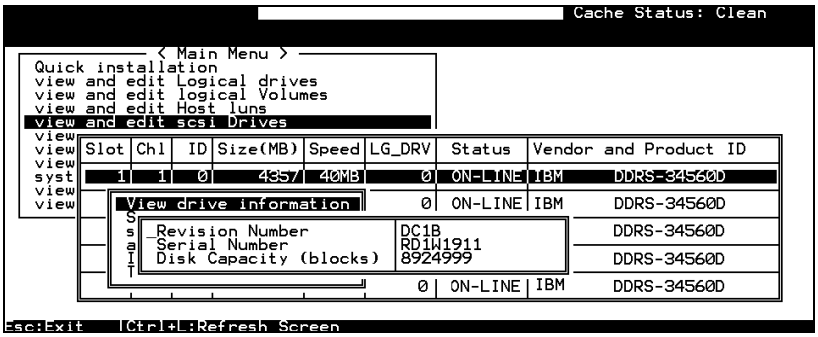

From the "View and Edit Drives" menu, select the drive that the utility is to be performed on, then press **[ENTER]**. Select "View drive information," then press **[ENTER]**.

7.16 Drive Utilities

From the "View and Edit Drives" menu, select the drive that the utility is to be performed on, then press **[ENTER]**. Select "Drive Utilities, then press **[ENTER]**. Choose either "Drive Low-level Format" or "Read/Write Test."

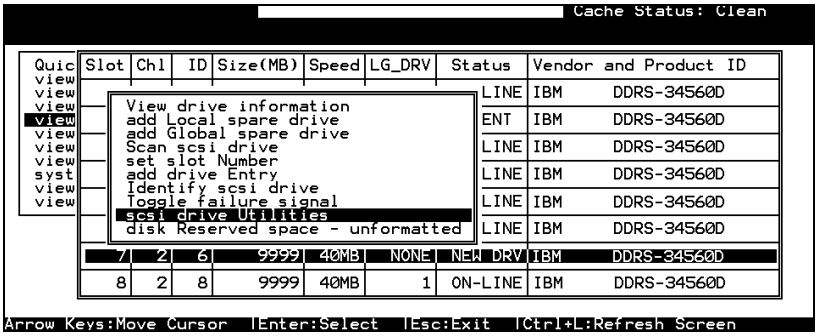

Drive Low-level Format

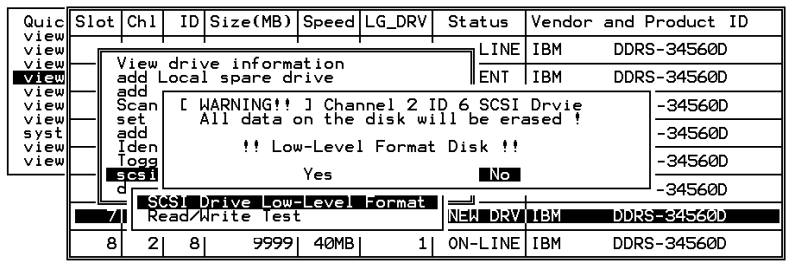

Choose "Drive Low-level Format" and confirm by selecting **Yes**.

IMPORTANT!

- *Do not switch the controller's and/or disk drive's power off during the Drive Low-level Format process. If any power failure occurs during the formatting process, the formatting must be performed again when power resumes.*

- *All data stored in disk drives will be destroyed during the low-level format.*

- *The disk drive on which a low-level disk format will be performed cannot be a spare drive (local or global) nor a member of a logical drive. The "Drive Low-level Format" option will not appear if the drive is not indicated as a "New Drive" or a "Used Drive." Also, a drive formatted with a 256MB reserved space is also excluded from selection.*

Drive Read/Write Test

From the "View and Edit Drives" menu, select a new or used drive that the utility is to be perform on, then press **[ENTER]**. Select "Drive Utilities," then press **[ENTER]**. Choose "Read/Write Test" and press **[ENTER]**. You can choose to enable/disable the following options:

1. Auto Reassign Bad Block
2. Abort When Error Occurs
3. Drive Test for - Read Only/Read and Write

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

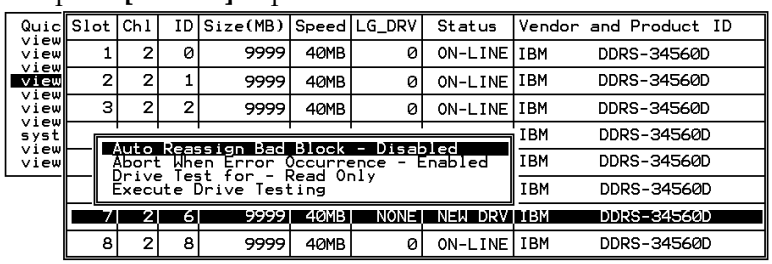

The Read/Write test progress will be indicated by a status bar.

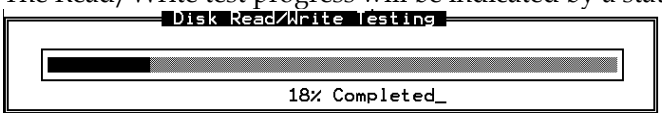

You may press **[ESC]** and select "Read/Write Test" later and choose either to "View Read/Write Testing Progress" or to "List Current Bad Block Table." If you want to stop testing the drive, select "Abort Drive Testing" and press **[ENTER]** to proceed.

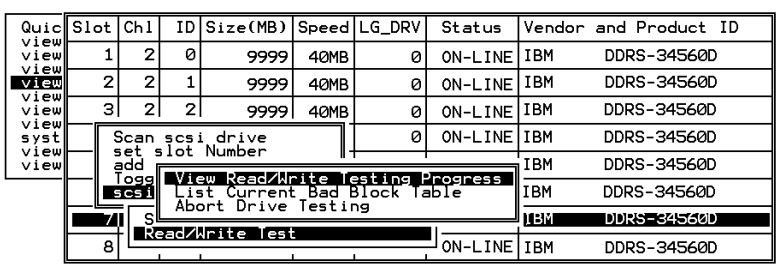

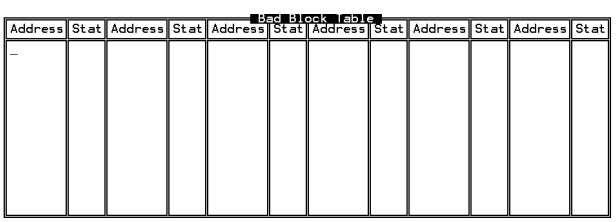

Chapter

8 Fibre Operation

8.1 Overview

This chapter describes the Fibre-specific functions available since firmware release 3.21 and above. Optional functions have been implemented for operations using Fibre Channel and access control under multiple-host environments such as Storage Area Network. Users familiar with Fibre Channel configurations, please move ahead to ***Section 8.4***.

Summary:

8.2 Major Concerns

Things you should know before proceeding with configuration

8.3 Supported Features

List of functionality supported by controller FC chips

8.4 Configuration: Host and Drive Parameters

Configuration options for data bus setting and system drive mapping.

The configuration procedures for changing parameters on FC host and drive channels

8.2 Major Concerns

Most of the configuration options in this chapter are directly related to controller redundancy. Joining two controllers into a dual-active pair can eliminate most possible points of failure. Configuring a controller pair requires careful planning and proper setup, the requirements can be summarized as follows:

▪ **Redundant Cache Coherency Channels (RCC):**	
1. RCC	FC channels can be manually assigned as the dedicated communications loops. Two are recommended for path redundancy and sufficient bandwidth.
2. Drive + RCC	Communications traffic distributed over drive loops
▪ **Connection between Controllers:**	
	Cabling between controllers, hardware link through a common backplane, Fibre hub or switch (for SAN applications and for those models that do not have by-pass chips)
▪ **Channel Mode Assignment**	
	According to the topological plan, your I/O channels can be designated as: • Host • RCC paths • Drive • Drive + RCC
▪ **Host Channel Connection Type:**	
	This depends on the way your RAID system is connected to the host computer(s). The host connection type can be: • FC-AL • Fabric (point-to-point)
▪ **Controller Unique ID:**	
	This ID will be used to generate Fibre ports' node names, and is necessary for addressing the controller during the controller failover/failback operation.
▪ **Dual-Loop:**	
	1. **Drive-side** dual loop provides data path redundancy. Firmware is capable of executing a

	load-sharing algorithm to optimize dual-loop performance. 2. **Host–side** dual loop is passively supported and requires the support of multi-path software on the host computer.

8.3 Supported Features

Fibre Chip

1Gbit Fibre Channel:

Fibre loops (1 Gbit FC-AL) comply with the following standards:

1. (FC-PH) X2.230:1994
2. (SCSI-FCP) X3.269:1996
3. (FC-AL-2) Project 1133-D rev.6.5
4. (SCSI-2) X3.131-1994
5. Support for sustained 1 Gbit/sec (100MB/sec) transfer rates
6. Each Fibre loop can be independently configured for the connection to host or drive

2Gbit Fibre Channel:

1. Fibre Channel Arbitrated Loop (FC-AL-2) working draft, rev 6.4
2. Fibre Channel Fabric Loop Attach (FC-FLA) working draft, rev 2.7
3. Fibre Channel Private Loop SCSI Direct Attach (FC-PLDA) working draft, rev 2.1
4. Fibre Channel Tape (FC-TAPE) profile, T11/98-124vD, rev 1.13
5. Support Fibre Channel protocol-SCSI (FCP-SCSI)
6. Support Fibre Channel Internet protocol (IP)

Multiple Target IDs

Each 2Gbit channel configured as a host loop supports multiple target IDs in the range of 0 to 125.

Drive IDs:

Supported ways to address a Fibre port include Hard assigned and Soft assigned. The controller supports automatic loop ID assignment on drive channels. A hard loop address ID can be assigned to disk drives by enclosure jumper setting. If the AL_PA configuration on a drive enclosure has been set to a neutral status, physical IDs will be automatically assigned to drives.

In-band Fibre and S.E.S. Support

"SCSI Pass-through" commands are supported over host and drive loops just as they are over SCSI channels. The "in-band Fibre" protocol for packaging "External Interface" protocol commands/responses is supported over host Fibre loops (such as the RAIDWatch Manager). Drive-side S.E.S. device identification, monitoring, and control are likewise supported over drive loops.

8.4 Configuration: Host and Drive Parameters

Channel Mode:

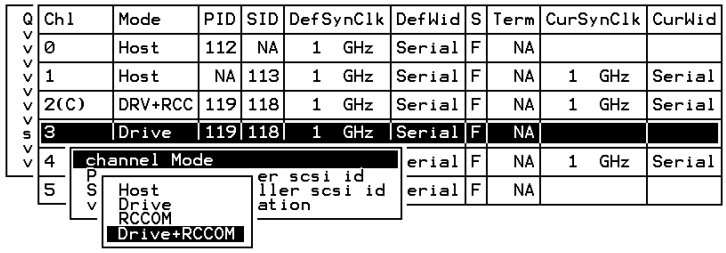

All Fibre Channels can be changed to operate as "Host," "Drive," "RCCOM," or "Drive + RCCOM". Choose the channel for which you wish to change its mode, then press **[ENTER]**. Choose "Channel Mode," then press **[ENTER]**. A dialog box will appear asking you to confirm the change.

Primary and Secondary Controller IDs:

Select a channel by highlighting its status bar and press **[ENTER]**.

Drive Channel In redundant controller mode, each drive channel should be assigned with both a "Primary Controller ID" and a "Secondary Controller ID." The factory defaults for the primary and secondary IDs on drive loops are "119" and "120" respectively.

Host Channel Create host IDs on each specific host channel. Host IDs are designated as the "Primary Controller" or "Secondary Controller" IDs.

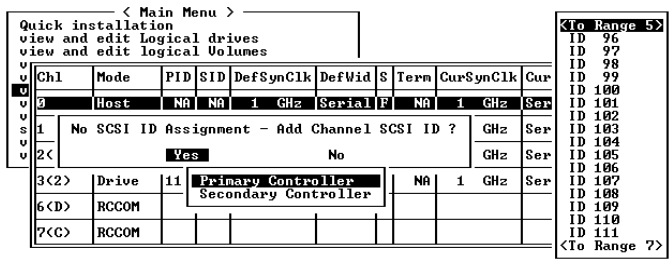

Limitation:

1. If host connection is made in FC-AL mode, there can be a total of "16" Primary and Secondary IDs.

2. Although host connection in point-to-point mode is supported, problems might occur when passing IDs of the failed controller to its counterpart. HA (Host Adapter) failover software is required in this situation to redirect I/O access during the controller failover operation.

Redundant Controller Cache Coherency Channel (RCC Channel)

A host channel can be selected and converted into a communications channel. To convert a drive channel, change it into host mode and then select "RCCOM." Details can be found in the proceeding discussions.

View Channel WWN

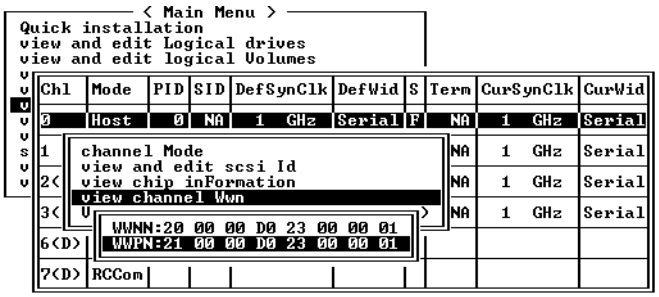

Port name is an unique eight-byte address assigned to an FC device port.

The controller has multiple channels (I/O paths) and each channel is managed by an I/O processor. This function allows users to inspect the node name and port name assigned to these chip processors. Some management software running on host computers needs these names to properly address a storage subsystem.

View Device Port Name List (WWPN)

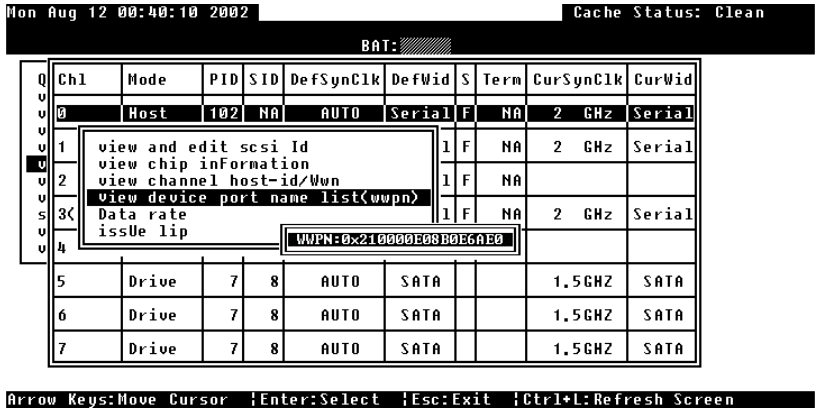

This function displays the device port names (host adapter ID) detected on a host loop or through fabric host connection.

The HBA port names detected can be added to the "**Host-ID WWN name list**" in "View and Edit Host LUN" menu. Adding port names to the list can speed the mapping process that follows.

Each port name should then be assigned a nickname for ease of identification.

Add Host – ID/WWN Label Declaration

A nickname can be added to any host adapter ID for ease of identification in environments, e.g., SAN, where multiple servers are connected.

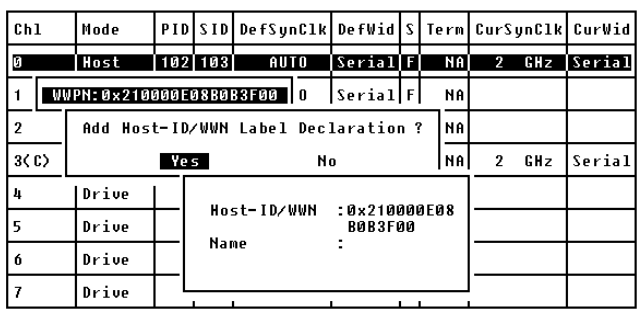

Press Enter on a displayed adapter WWPN and you will be prompted with a confirm box as shown above. Choose Yes and enter a name for the host adapter.

View and Edit Fibre Drive

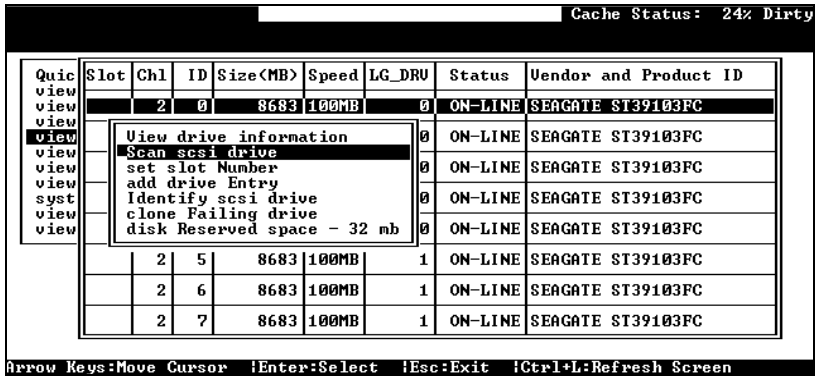

Choose "View and Edit SCSI Drives" on the Main Menu and use the arrow keys to move the cursor bar through the connected drives. Press **[ENTER]** to choose a drive, or **[ESC]** to return to the previous menu/screen.

User-Assigned ID (Scan SCSI Drive)

Select "Scan SCSI drive" to assign an ID to a drive.

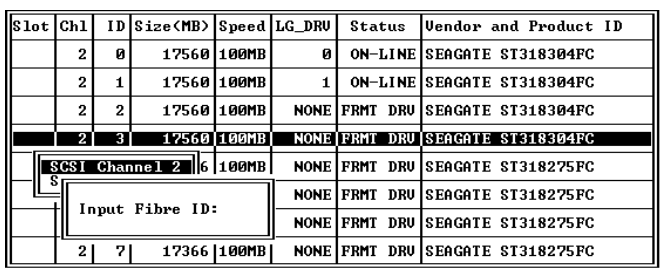

A drive enclosure usually has drive slots pre-assigned with specific IDs. There are occasions when an ID needs to be assigned manually to a device. The "set slot number" and the "add drive entry" functions are reserved for Infortrend's Fault-bus operation.

View Drive Information

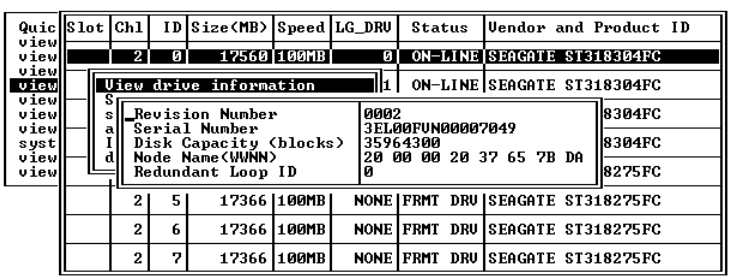

If the selected drive belongs to a drive group that is configured in a dual-loop, the "Redundant Loop ID" will be displayed here.

View and Edit Host-side Parameters

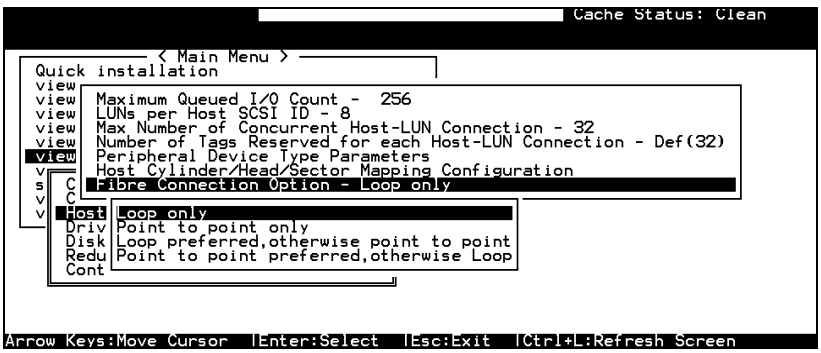

1. Fibre Channel Connection Option:

Use the ↑ ↓ keys to scroll down to "View and Edit Configuration Parameters," "Host-side SCSI Parameters," and then "Fibre Connection Option." A prompt will display all the options. Select one appropriate for your Fibre Channel topology. If the connection to the host is through a Fibre hub, choose "Loop only."

If connection to the host is through a Fibre switch F_Port or directly to a server, choose "Point to point only." Proper selection is necessary and will decrease overhead on data transmission.

A redundant controller configuration should always have its host connection configured in FC-AL mode. For a switched fabric configuration, a redundant controller system can be connected to the FL_ports on an FC switch and then the host computers connect to its F_ports.

IMPORTANT!

- *Note that the host ports between redundant controllers can be connected via switched fabric. However, switch ports connecting the controllers should be operating as FL_Ports. If they are configured as F_Ports, the controller host ports will not support multiple target IDs, especially when IDs on a failed controller's host ports need to be taken over by the surviving controller.*

View and Edit Drive-side Parameters

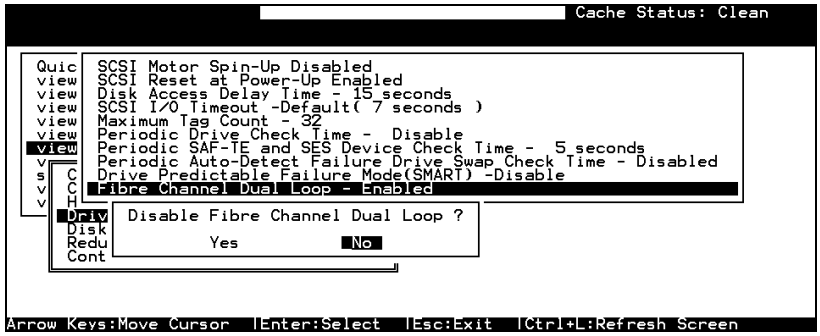

2. Drive-side Dual Loop

Fibre drives are usually configured in a JBOD enclosure. Through the enclosure backplane, these drives form one or two circuit loops. You may choose to assemble a certain number of disk drives into a dual-loop configuration using two of the controller channels.

To configure a dual-loop, connect two of the drive channels each to an FC-AL port on the drive enclosure (JBOD). Please refer to the related documents that came with your drive enclosure for the connection details.

The dual-loop configuration not only doubles traffic bandwidth by separating the transmitting and receiving paths but also provides path redundancy. I/O traffic will be continued if one data path fails.

Controller firmware automatically examines the node names and port names of all the connected drives once initiated. If devices on two different drive channels appear with the same loop ID and port name, the controller will consider these two drive channels as a "dual loop."

< Main Menu >
Quick installation
view and edit Logical drives
view and edit logical Volumes

Chl	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid
0	Host	0	NA	1 GHz	Serial	F	NA	1 GHz	Serial
1	Host	NA	1	1 GHz	Serial	F	NA	1 GHz	Serial
2<3>	Drive	119	120	1 GHz	Serial	F	NA	1 GHz	Serial
3<2>	Drive	119	120	1 GHz	Serial	F	NA	1 GHz	Serial
6<C>	RCCom								
7<C>	RCCom								

The dual loop configuration will be displayed as "channel <pair channel>." For example, channel numbers are displayed as 2<3> and 3<2> if channel 2 and channel 3 are configured as a dual loop.

The data bus will be operating at the bandwidth of up to 400MB/sec (2Gbps Fibre).

Controller Unique Identifier

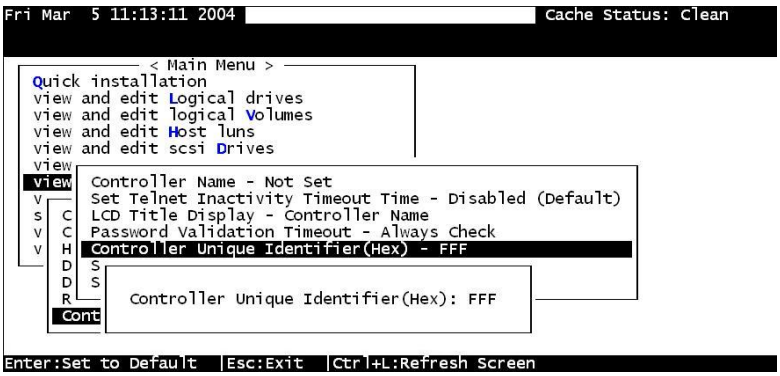

A Controller Unique Identifier is ***required*** for operation with the ***Redundant Controller Configuration***. The controller will automatically notify users to enter a unique identifier when the first logical drive is being created in a dual-controller system.

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, the controller configuration data is continuously synchronized between controllers. Host ports on both controllers appear with the identical node names and each with a different port name (WWPN). 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 controller failback is totally transparent.

Choose "View and Edit Peripheral Devices," "Set Peripheral Device Entry," then enable the "Redundant Controller" configuration. You will be requested to enter a value for the "Controller Unique Identifier." For firmware release 3.25 and above, enter a hex number between **0** and **FFFFF**. The identifier selection box will prompt automatically. The value you enter ***MUST*** be different for each controller.

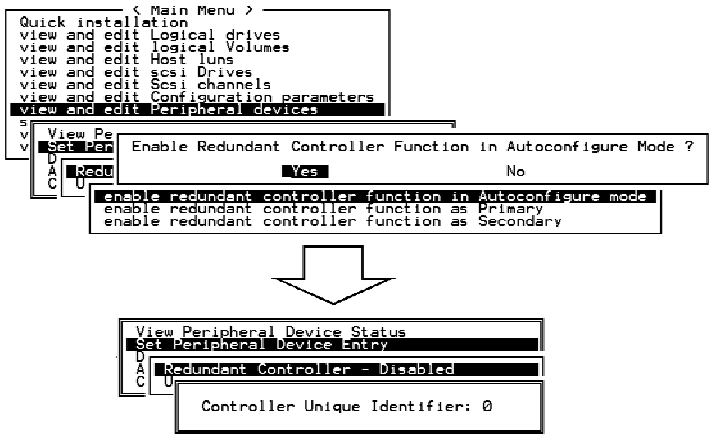

The unique identifier can also be accessed from "View and Edit Configuration Parameters" ∏ "Controller Parameters" ∏ "Controller Unique ID."

Controller Communications over Fibre Loops

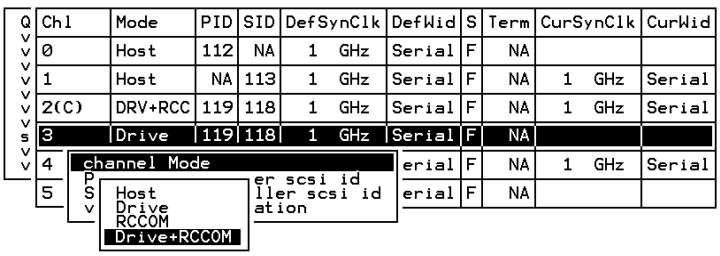

Controllers running firmware version 3.14 and above supports controller communications over Fibre loops.

There are two options for controller communications over Fibre loops. Hardware configuration should be completed before firmware setting.

1. Select from the Main Menu "View and Edit SCSI channels," and configure the selected FC channels into "RCCOM (Redundant Controller Communication)" mode. To ensure the connection with data path redundancy, you may use two channels as the dedicated RCC loops. The dedicated channels should not be attached with any other device.

2. **Communications Traffic Distributed over All Drive Loops**: Select all drive loops and configure them as "Drive + RCCOM (Drive Loops plus Redundant Controller Communications)." The communications traffic between the two controllers will be automatically distributed over all drive loops.

Chl	Mode	PID	SID	DefSynClk	DefWid	S	Term	CurSynClk	CurWid
0	Host	112	NA	1 GHz	Serial	F	NA	1 GHz	Serial
1	Host	NA	113	1 GHz	Serial	F	NA	1 GHz	Serial
2(3;C)	DRV+RCC	119	118	1 GHz	Serial	F	NA	1 GHz	Serial
3(2;C)	DRV+RCC	119	118	1 GHz	Serial	F	NA	1 GHz	Serial
4(5;C)	DRV+RCC	119	118	1 GHz	Serial	F	NA	1 GHz	Serial
5(4;C)	DRV+RCC	119	118	1 GHz	Serial	F	NA	1 GHz	Serial

- As displayed above, channel(s) selected as the communications paths will be displayed as "channel number (C: connected)" or "channel number (D: disconnected)." If channels configured in a dual-loop are selected, channel status will be displayed as "channel number (pair loop; C or D)."
- If any of the communications loops fail, the inter-controller traffic will automatically shift to the remaining Drive/RCC loop(s).

Chapter

9 Host-side and Drive-side Parameters

This chapter discusses the advanced options for configuring and maintaining a RAID system. 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.

9.1 Host-side Parameters

Foreword: SCSI Channel, SCSI ID, and LUN

A SCSI channel (SCSI bus) can connect up to 15 devices (not including the controller itself) when the Wide function is enabled (16-bit SCSI). It can connect up to 7 devices (not including the controller itself) when the Wide function is disabled (8-bit SCSI). Each device has one unique SCSI ID. Two devices contending for the same SCSI ID are not allowed.

Figure 9 - 1 SCSI ID/LUNs

Figure 9-1 illustrates this idea. If you file a document into a cabinet, you must put the document into one of the drawers. From a SCSI's point of view, a SCSI ID is like a cabinet, and the drawers are the LUNs (Logical units). Each SCSI ID enables up to 32 LUNs. Data can be stored into one of the LUNs of the SCSI ID. Most SCSI host adapters treat a LUN like another SCSI device.

The same holds true for a Fibre Channel host interface. Up to 32 LUN's are supported with each host ID and up to 1024 LUNs are configurable per controller.

Maximum Concurrent Host LUN Connection ("Nexus" in SCSI)

The configuration option adjusts the internal resources for use with a number of current host nexus. If there are four host computers (A, B, C, and D) accessing the array through four host IDs/LUNs (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 (for example, host A access 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 SCSI 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."

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.

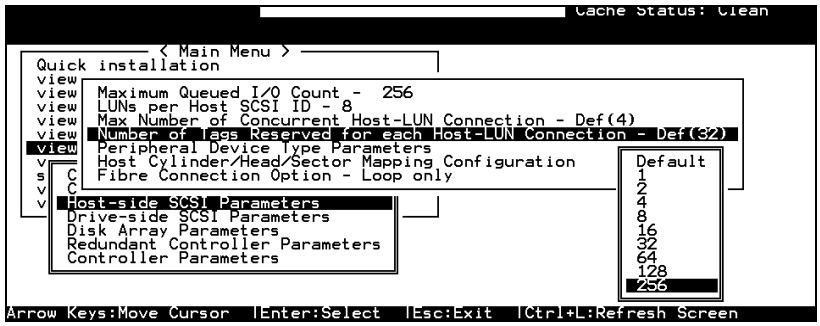

Choose "Host-side SCSI Parameters," then press **[ENTER]**. Choose "Number of Tags Reserved for each Host-LUN Connection," then press **[ENTER]**. A list of available selections will appear. Move the cursor bar to an item, then press **[ENTER]**. Choose **Yes** in the dialog box that follows to confirm the setting.

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.

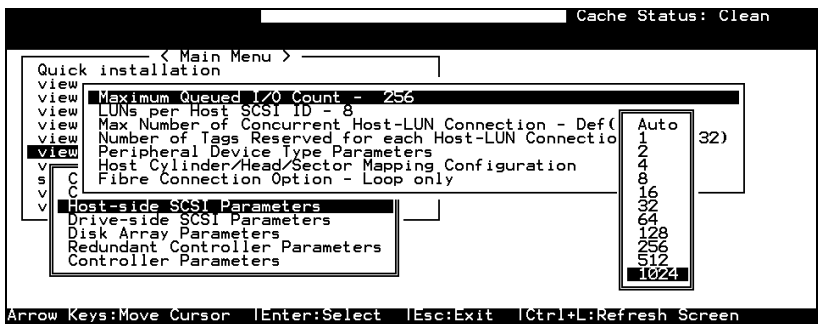

Choose "Host-side SCSI Parameters," then press **[ENTER]**. Choose "Maximum Queued I/O Count," then press **[ENTER]**. A list of available selections will appear. Move the cursor bar to an item, then press **[ENTER]**. Choose **Yes** in the dialog box that follows to confirm the setting.

The controller supports the following Host-side SCSI configurations:
"Maximum Queued I/O Count," "LUNs per Host SCSI ID," "Num of Host-LUN Connect," "Tag per Host-LUN Connect," "Peripheral Dev Type Parameters," and "Cyl/Head/Sector Mapping Config."

LUNs per Host SCSI ID

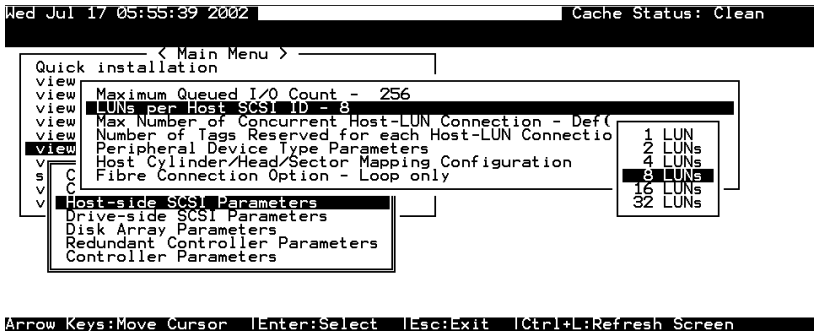

Choose "LUNs per Host SCSI ID," then press **[ENTER]**. A list of selections will appear. Move the cursor bar to an item, then press **[ENTER]**. Choose **Yes** in the dialog box that follows to confirm the setting.

LUN Applicability:

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 SCSI card, 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 RAID using management software from the host. It will be necessary to configure the "Peripheral Device Type" setting for the host to communicate with the controller. 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. If "all undefined LUNs" is selected, each LUN in that host ID will appear as a device with the user-defined peripheral device type.

Different "LUN applicability" selections are available: "Device Type" selection, "Device Qualifier Support," "Support Removable media," "LUN-0's only," and "All undefined LUNs." Please refer to ***Peripheral Device Type Parameters for Various Operating Systems*** for details concerning various operating systems.

Peripheral Device Type:

For connection without a preset logical unit to a host, the in-band SCSI protocol can be used for the host to "see" the RAID controller. Please refer to the reference table below. You will need to make adjustments in the following submenu: Peripheral Device Type, Peripheral Device Qualifier, Device Support for Removable Media, and LUN Application.

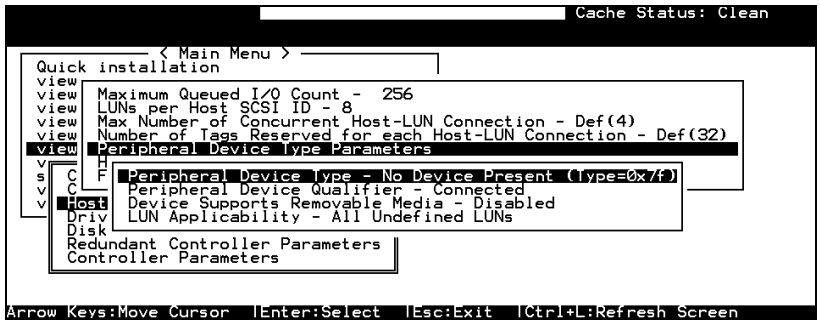

In-band (SCSI or Fibre)

What is In-band?

External devices require communications with the host 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 link, either SCSI or Fibre.

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 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** the following is configured:

1. (a) a logical unit has been created

   (b) at least one logical unit is mapped to a host ID or LUN via the RS-232/LCD keypad interface

2. (a) host bus connection is established.

   (b) the RAID subsystem/controller is configured to appear as a peripheral device on the channel bus connected to a host computer

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 from the host 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 as a peripheral device.

Different host operating systems require different adjustments. See the tables below to find the proper settings for your host operating system. References to "Peripheral Device Qualifier" and "Device Support for Removable Media" are also included.

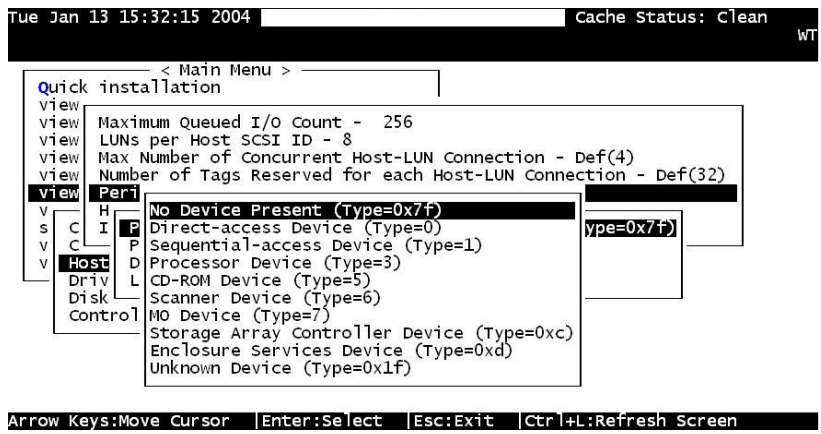

Table 9 - 1 Peripheral Device Type Parameters

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 9 - 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

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 Sun Solaris, 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 9 - 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 9 - 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:

Selecting Sector Ranges

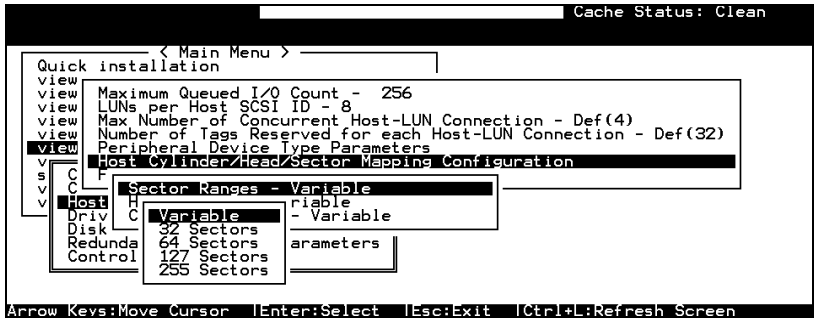

Selecting Head Ranges

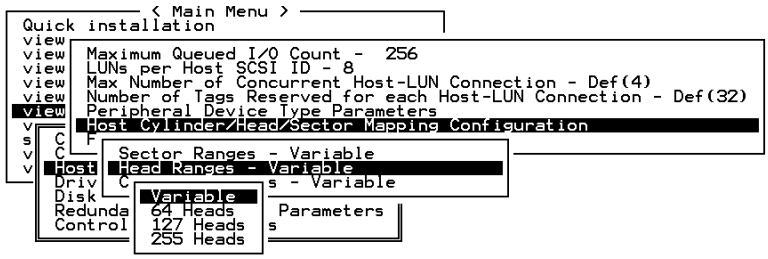

Selecting Cylinder Ranges

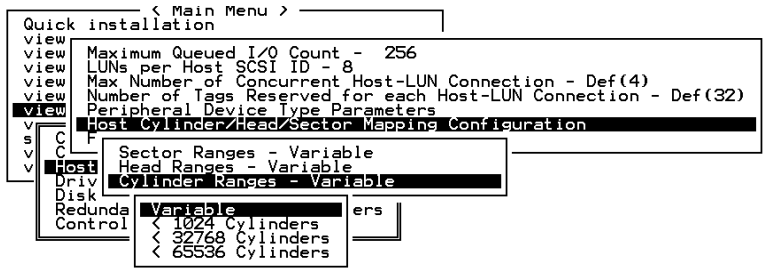

9.2 Drive-side Parameters:

SCSI-based Subsystems

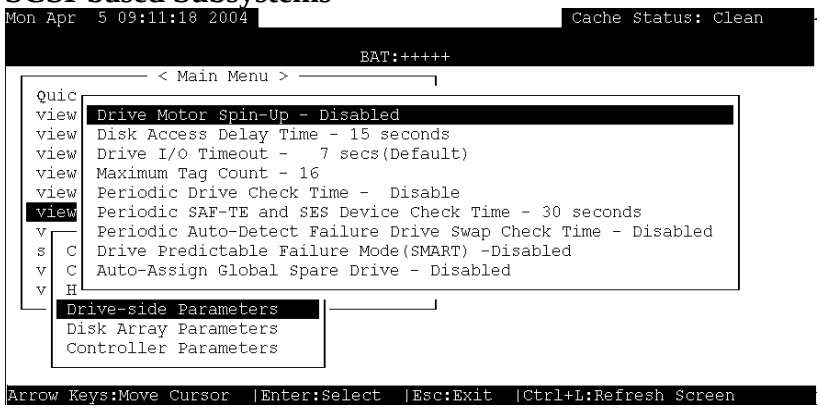

Fibre-based Subsystems

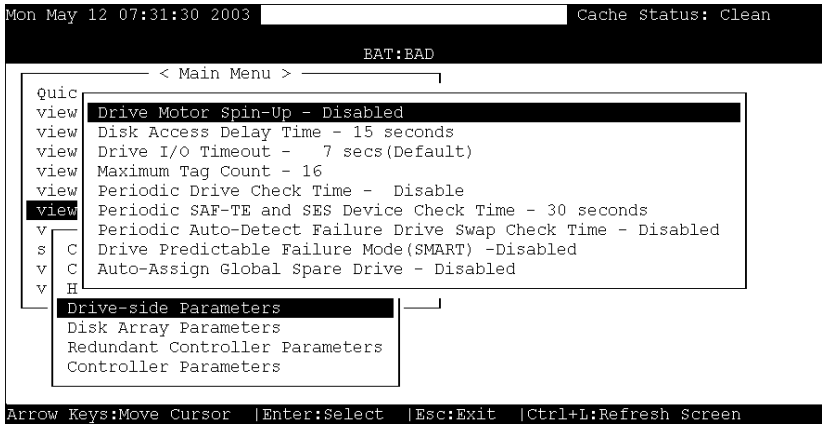

Choose "Drive-side SCSI Parameters," then press **[ENTER]**. The Drive-side SCSI parameters menu displays.

SCSI Motor Spin-Up

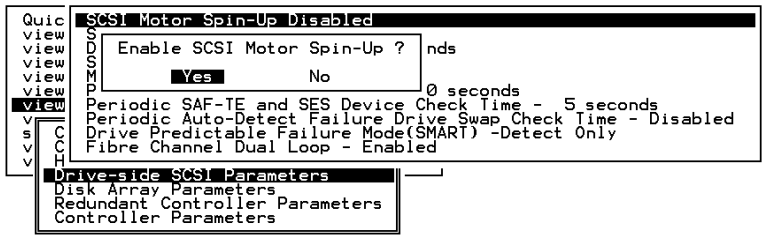

When a power supply is unable to provide sufficient current to start all of the hard drives at once, you may have the drives spin-up in sequence to solve the problem with insufficient power-up current.

By default, all hard drives will start spinning up when powered-on; however, these hard drives can be configured so that drives will not spin-up all at the same time. There are 3 methods for spinning-up the hard drive's motor: "Spin-up at power-on," "Spin-up serially at random sequence" or "Spin-up by SCSI command." Please refer to the hard drive's documentation for instructions on configuring the hard drive using the "Spin-up by SCSI Command." The procedure for each brand/model of hard drive may vary.

Configure all hard drives (usually by switching jumpers on the hard disks). Choose "SCSI Motor Spin-Up," then press **[ENTER]**. Choose **Yes** in the dialog box that follows to confirm the setting.

IMPORTANT!

- *If the drives are configured as "Delay Motor Spin-up" or "Motor Spin-up in Random Sequence," some of these drives may not be ready when the controller accesses them when it is powered up. Increase the disk access delay time so that the controller will wait a longer time for the drives to be ready.*

Disk Access Delay Time

This feature sets the delay time before the controller tries to access the hard drives after power-on. Default is 15 seconds.

SCSI Drives

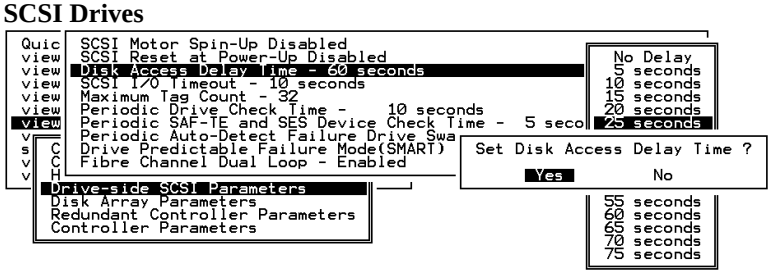

Fibre Drives

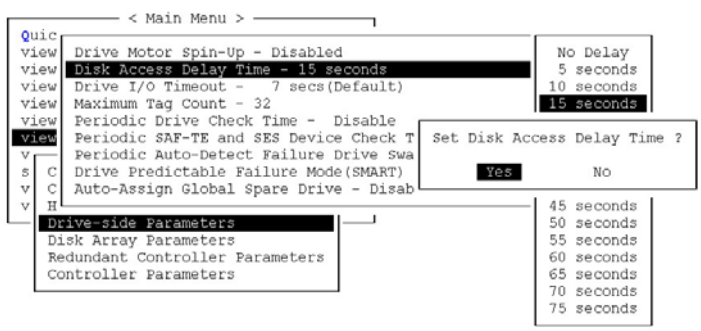

Choose "Disk Access Delay Time," then press **[ENTER]**. A list of selections displays. Move the cursor bar to a selection, then press

[ENTER]. Choose **Yes** in the dialog box that follows to confirm the setting.

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 head. When the drive encounters a bad block on the media, it reassigns the bad block onto a spare block. However, it takes time to perform the above actions. The time to perform these operations can vary between different brands and different models.

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 SCSI 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.

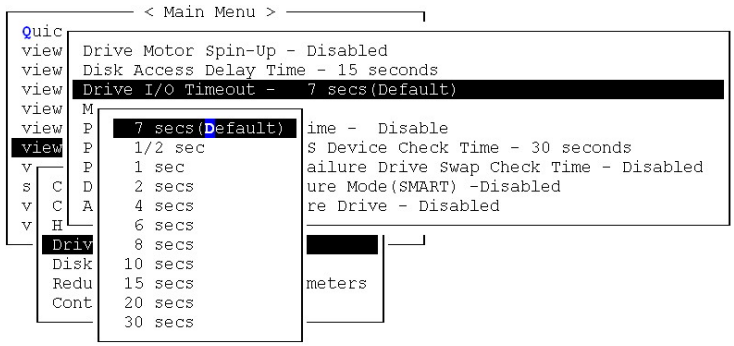

Choose "Drive I/O Timeout –Default (7 seconds)," then press **[ENTER]**. A list of selections will appear. Move the cursor bar to a selection, then press **[ENTER]**. Choose **Yes** in the dialog box that follows to confirm the setting.

Maximum Tag Count:

Tag Command Queuing (TCQ) and Native Command Queuing (NCQ) Support

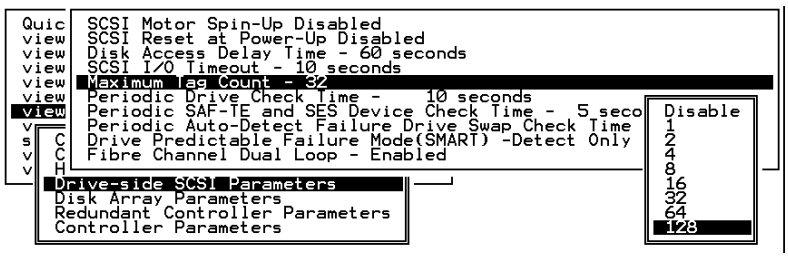

This sub-menu facilitates the support for both Tagged Command Queuing (TCQ) and Native Command Queuing (NCQ). ICQ is a traditional feature on SCSI or Fibre disk drives, while NCQ is recently implemented with SATA disk drives. The queuing feature requires the support of both host adapters and hard disk drives. Command queuing can intelligently reorder host requests to streamline random accesses for IOPS/multi-user environments.

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 or Fibre drives) or 16 (default for 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 for the changes to take effect.*
- *Disabling Tag Command Queuing will disable the hard drives' built-in cache for Write-Back operation.*

Chapter

10 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.

10.1 Enclosure Monitoring

10.1.1 Enclosure Devices

SAF-TE and SES Enclosure Monitoring

(Periodic SAF-TE and SES Device Check Time)

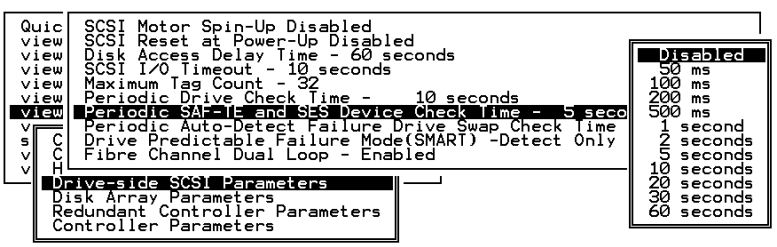

If there are remote components in your RAID enclosure being controlled via SAF-TE/S.E.S. devices, use this function to select how often the controller will check the status of these devices. Choose "Periodic SAF-TE and SES Device Check Time"; then press **[ENTER]**. Move the cursor to the desired interval, then press **[ENTER]**. Choose **Yes** in the dialog box that follows to confirm the setting.

Some Infortrend EonStor drive enclosures (JBODs) come with the SES monitoring service and have a preset value for the device check time. The firmware default is 30 seconds.

Enclosure Devices Status: View Peripheral Device Status

Select "View and edit Peripheral Devices" on the Main Menu and press **[ENTER].** Choose "View Peripheral Device Status," then press **[ENTER]** again. The device list displays.

Below is a list of peripheral devices (enclosure modules) supported by enclosure monitoring. Monitoring of device status depends on enclosure implementation and is accessed through different interfaces, e.g., SAF-TE, S.E.S., or I^2C serial bus.

1. Device type
2. Enclosure descriptor
3. Power supply
4. Cooling element
5. Temperature Sensors
6. Audible alarm
7. Enclosure services controller electronics

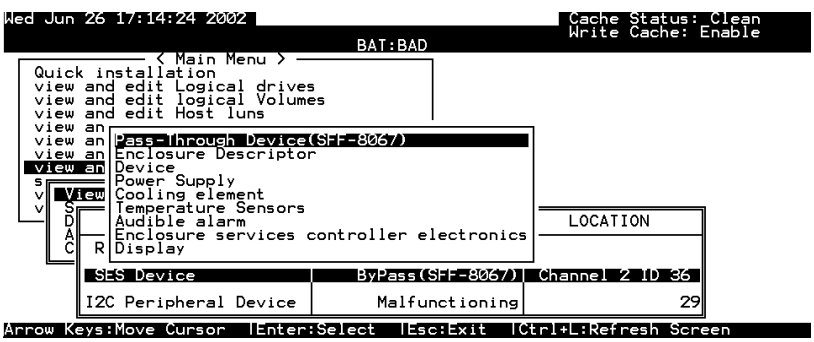

Select the device interface and then press **[ENTER]** on a connected module to examine its operating status as shown below.

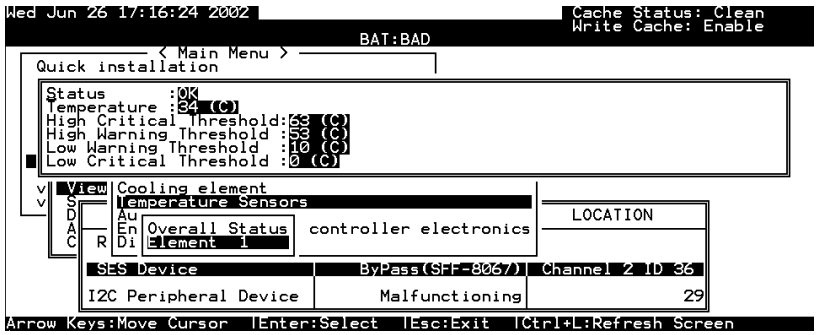

Event Triggered Operations

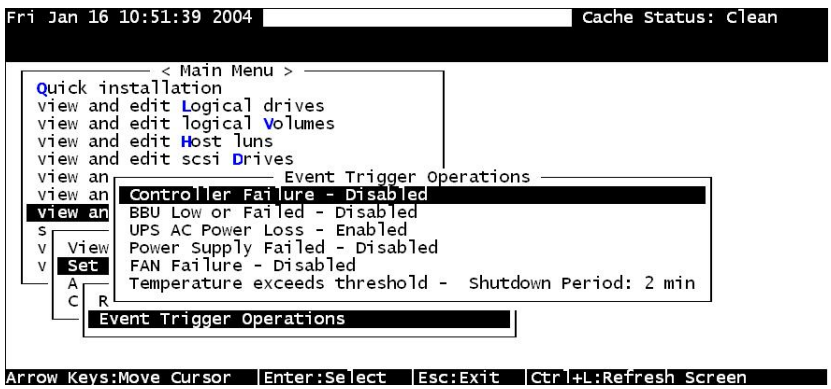

Select "View and Edit Peripheral Devices" on the Main Menu and press **[ENTER].** Choose "Set Peripheral Device Entry", press **[ENTER],** then select "Event Trigger Operations" by pressing **[ENTER]**. The event trigger menu displays.

The Operations

To reduce the chance of data loss, the controller/subsystem automatically commences the following actions to prevent loss of data:

1). Switches its caching mode from write-back to write-through,

2). Flushes all cached data

The Trigger

The mode switching and cache flush operations can be triggered by the detection of the following conditions:

1. Controller failure (in dual-redundant controller configurations)

2. BBU low or failed
3. UPS AC power loss
4. Power supply failure
5. Fan failure
6. Temperature exceeds threshold

When enabled, each of the above conditions forces the controller/subsystem to adopt the write-through caching mode. When the fault condition is corrected, the controller/subsystem automatically restores the previous caching mode.

Note the temperature thresholds refer to those set for "RAID controller board temperature" and the sensors within 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 daughter card.

If a battery is not used, the "BBU Low or Failed " option should be disabled.

You can configure a shutdown period to delay the time when the controller/subsystem commences an automatic shutdown. During this time, system administrators should have been notified of the condition and begun restoring proper cooling of the subsystem.

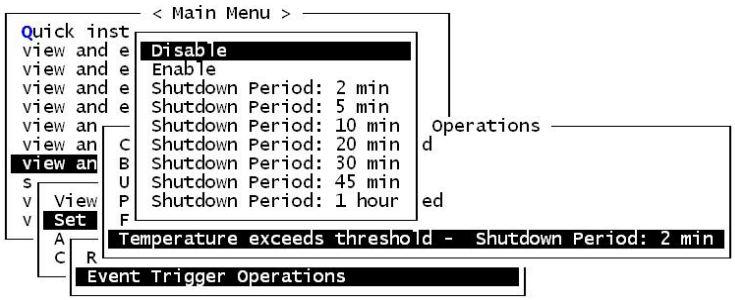

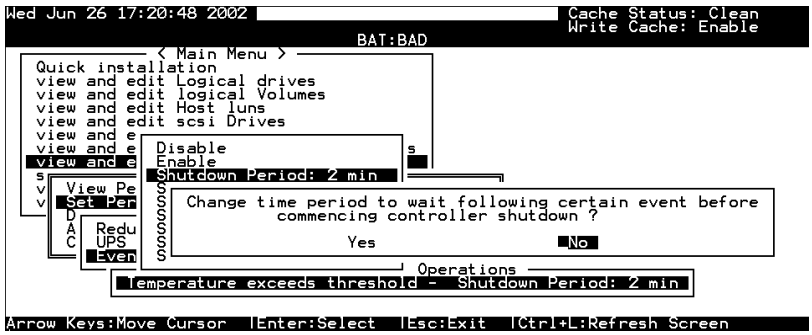

Select "View and Edit Peripheral Devices" on the Main Menu and press **[ENTER].** Choose "Set Peripheral Device Entry" and "Event Trigger Option" by pressing **[ENTER]**. The auto-shutdown option displays.

Select a configurable time span between the detection of exceeded temperature and the controller's commencing an automatic shutdown. Extended operation under critical conditions like elevated temperature greatly reduces system efficiency and may cause component failure.

10.1.2 Controller Self-monitoring

Open your PC Terminal Emulation screen. Enter the Main Menu and select "View and Edit Peripheral Devices." Use the arrow keys to scroll down and select "Controller Peripheral Device Configuration," "View Peripheral Device Status," and then press **[ENTER]**.

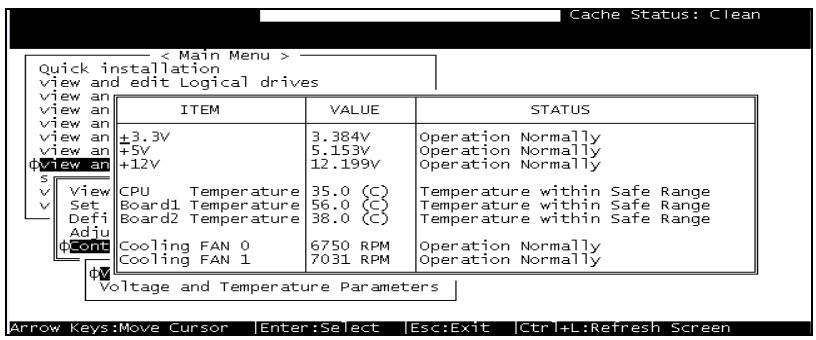

The controller operation status displays. Note the fan status here refers to the 3cm fans in the front panel of RAID controllers that come standard in 5.25" canister. The enclosure fan status is shown in the *Enclosure Devices Status: View Peripheral Device* Status.

Changing Monitoring Thresholds

Open your PC Terminal Emulation utility. Enter the Main Menu and select "View and Edit Peripheral Devices." Use the arrow keys to scroll down and select "Controller Peripheral Device Configuration," "Voltage and Temperature Parameters," and confirm by pressing **[ENTER]**.

Note that 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 high or low, adjusting the default thresholds can coordinate the controller status monitoring with that of the system enclosure.

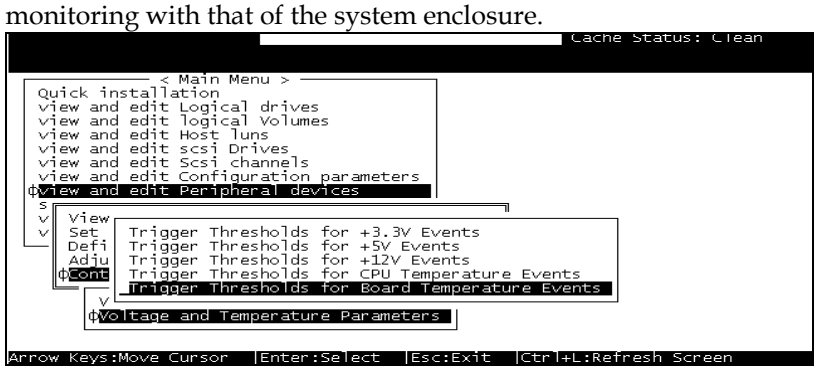

Scroll down and select an item to configure.

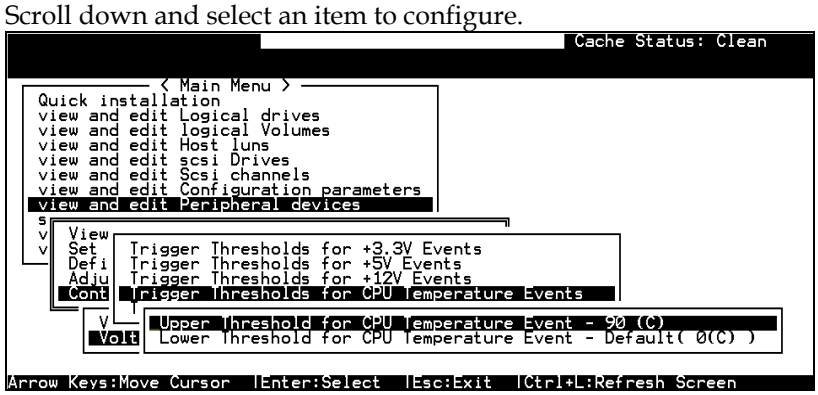

Select an item, such as "Trigger Thresholds for CPU Temperature Events." Press **[ENTER]** and a list of selections will appear. You can change the upper or lower threshold values. Press **[ENTER]** to confirm.

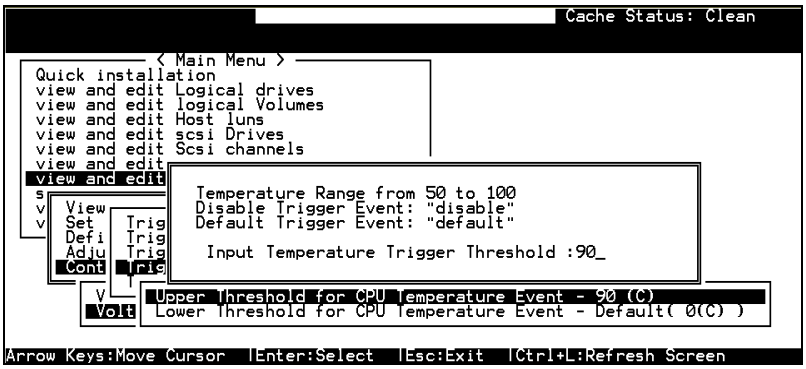

A configuration window will prompt. Enter any value within the safety range. Values exceeding the safety range will be rejected by controller firmware.

Follow the same method to modify other threshold parameters.

10.2 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) by dynamically switching the write policy.

Requirements for UPS Status Monitoring:

1. Connect the 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 3***, RS-232C Serial Port Settings.
3. Set the UPS option in the Event Triggered Operations 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 commences an auto cache-flush and is forced to adopt the Write-Through mode.
2. When the serial port connection is lost or when the UPS is removed, an event is issued to notify the system manager 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.

Chapter

11

Data Integrity

This chapter discusses various firmware mechanisms that help to secure the integrity of array data. Four main topics are included:

- **Maintenance processes** that can be manually or automatically performed onto configured arrays to minimize the chance of data inconsistency due to drive media errors.
- **Data path re-routing capability** to direct data flow through an alternate path in the situation when one of the cabling connections fails.
- **Preventive schemes** to suspend caching activities that are susceptible to hardware faults until the faults are corrected.
- **Drive Fault detection** and **Rebuild settings**.

No system is completely safe from the threat of hardware faults. For example, although the chance of occurrence is considerably low, the occurrences of bad blocks on two 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
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 be implemented through firmware append files and may not be available to all sub-revisions of firmware.

11.1 Protections against Hardware Faults

Event Triggered Operations

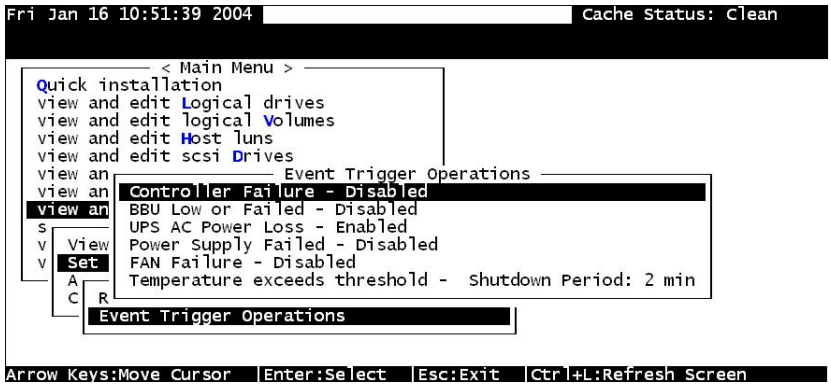

Select "View and Edit Peripheral Devices" on the Main Menu and press **[ENTER].** Choose "Set Peripheral Device Entry", press **[ENTER],** then select "Event Trigger Operations" by pressing **[ENTER]**. The event trigger menu displays.

The Operations

To reduce the chance of data loss, the controller/subsystem automatically commences the following actions to prevent loss of data:

1). Switches its caching mode from write-back to write-through

2). Flushes all cached data

The Trigger

The mode switching and cache flush operations can be triggered by the detection of the following conditions:

1. ***Controller failure***:

   If a controller fails in a dual-redundant controller configuration, the surviving controller no longer has the protection of cached data by storing the replica of unfinished writes to its partner.

2. ***BBU low or failed***:

   If 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 the mains power fails, cached data should be immediately distributed to hard drives.

4. ***Power supply failure***

5. ***Fan failure***

6. ***Temperature exceeds threshold***

   If critical conditions occur, such as the enclosure component failure, chance of system downtime will increase and it is best to temporarily disable write-back caching.

When enabled, each of the above conditions forces the controller/subsystem to adopt the write-through caching mode. When the fault condition is corrected, the controller/subsystem automatically restores the previous caching mode.

Note that 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 trigger automatically applies.

NOTE:

If a battery is not installed in your RAID subsystem, the "BBU Low or Failed " option should be disabled.

Auto Shutdown: 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.

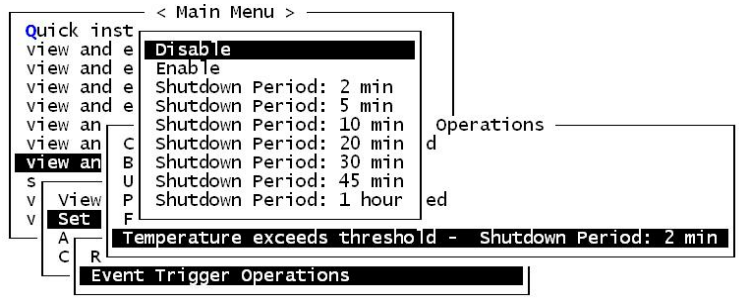

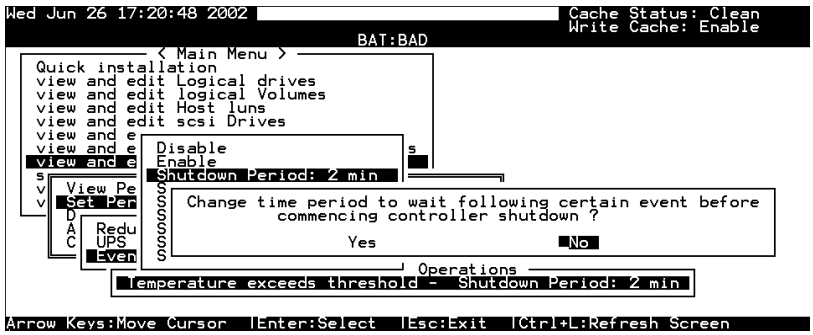

Select "View and Edit Peripheral Devices" on the Main Menu and press **[ENTER].** Choose "Set Peripheral Device Entry" and "Event Trigger Option" by pressing **[ENTER]**. The auto-shutdown options display.

Select a configurable time span between the detection of exceeded temperature and the controller's commencing an automatic shutdown. Extended operation under critical conditions like elevated temperature greatly reduces system efficiency and will eventually cause component failure.

11.2 Maintenance

11.2.1 Failed Drive Detection

Detection of Drive Hot Swap Followed by Auto Rebuild

(Periodic Auto-Detect Failure Drive Swap Check Time)

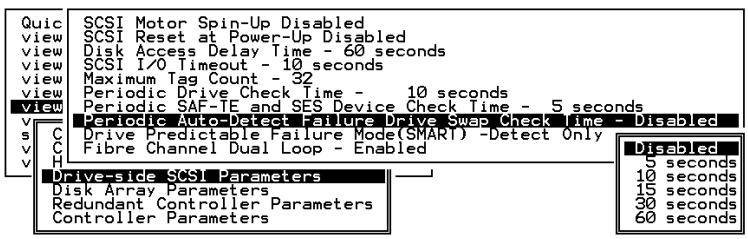

Choose "Periodic Auto-Detect Failure Drive Swap Check Time"; then press **[ENTER]**. Move the cursor to the desired interval; then press **[ENTER]**. Choose **Yes** in the dialog box that follows to confirm the setting.

The controller scans drive buses at this interval to check if a failed drive has been replaced. If a failed drive is replaced, the controller will proceed with the rebuild process.

Periodic Drive Check Time

The "Periodic Drive Check Time" is the time interval for the controller to check all disk drives that were on the SCSI bus at controller startup (a list of all the drives that were detected can be seen under "View and Edit SCSI Drives").

The default value is "Disabled." "Disabled" means that if a drive is removed from the bus, the controller will not be able to know – so long as no host accesses that drive. Changing the check time to any other value allows the controller to check – at the selected time interval – all of the drives that are listed under "View and Edit SCSI Drives." If any drive is then removed, the controller will be able to know – even if no host accesses that drive.

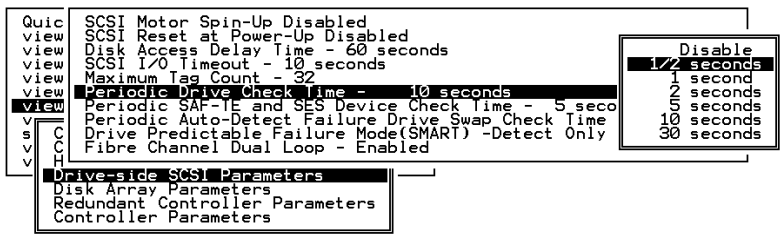

Idle Drive Failure Detection

Periodic Auto-Detect Failure Drive Swap Check Time

The "Drive-Swap Check Time" is the interval at which the controller checks to see if a failed drive has been swapped. When a logical drive's member drive fails, the controller will detect the failed drive (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 "Disabled," meaning that the controller will not Auto-Detect the swap of a failed drive. To enable this feature, select a time interval.

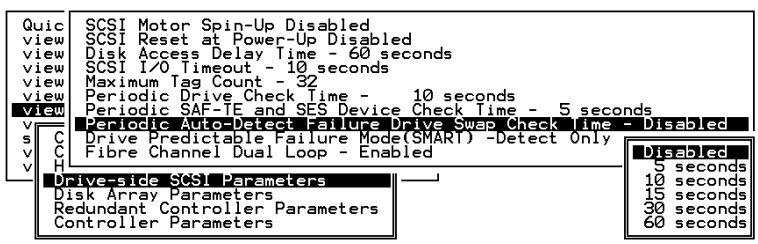

Choose "Periodic Drive Check Time," then press **[ENTER]**. Move the cursor to the desired interval, then press **[ENTER]**. Choose **Yes** in the dialog box that follows to confirm the setting.

IMPORTANT!

- *The "Periodic Drive Check Time" 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 "Periodic Drive Check Time" 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 a host attempts to access data on that drive.*

Auto-Assign Global Spare Drive

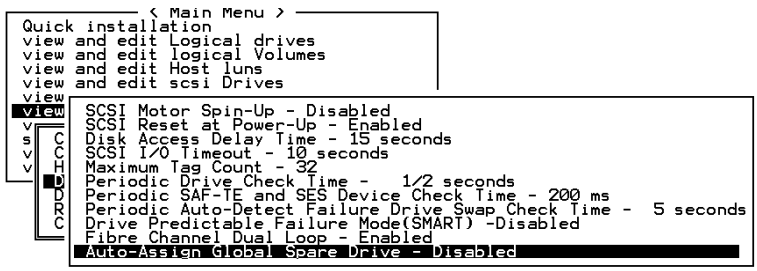

The "Auto-Assign" function automatically assigns any "new" drives that are not included in logical configurations as Global Spares.

The Fault Scenario:

Spare drives accelerate rebuild of a logical drive. In the example as described below, multiple faults can occur at the same time making the array exposed to the risk of data loss:

- There is only one Global Spare in a RAID subsystem.
- That Global Spare has been used to rebuild a logical drive.
- The failed drive is swapped out and replaced by a new one.
- Chances are system administrators forgot to configure the replacement drive as a spare.
- A member of another logical drive fails.

- The subsystem has no spare left. Performance decreases for a considerable portion of system resources has to be conducted to generate data from the remaining members of the logical drive. If another member fails in the logical drive, data is lost.

The chance of failing two drives increases when a failed drive in the array cannot be replaced immediately for the lack of spare drives.

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.

11.2.2 Scheduled Maintenance

Task Scheduler

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 individual drives, all drives of a certain class, all member drives of a specified logical drive, or all member drives of all logical drives. UIs supported are RS232C terminal menus and RAIDWatch GUI manager.

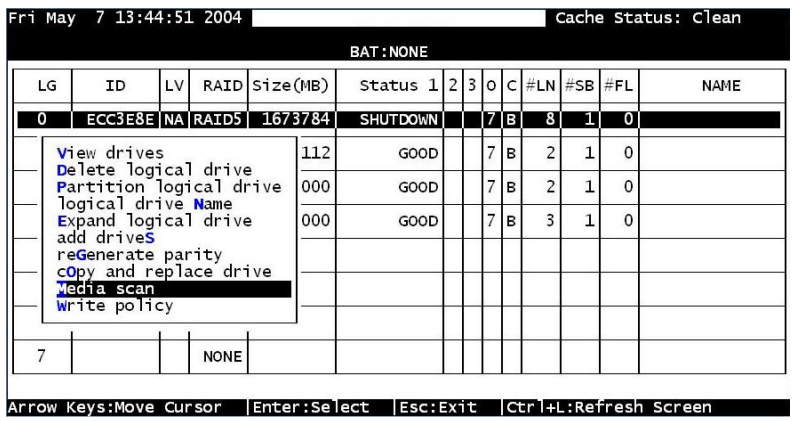

The scheduler is accessed through the "View and Edit Logical Drives" menu by selecting a configured array to display the list of array-related functions. Select "Media Scan" by pressing **[ENTER]**.

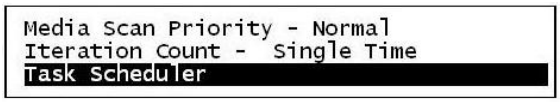

Select "Task Scheduler" by pressing **[ENTER]**.

Creating a New Schedule

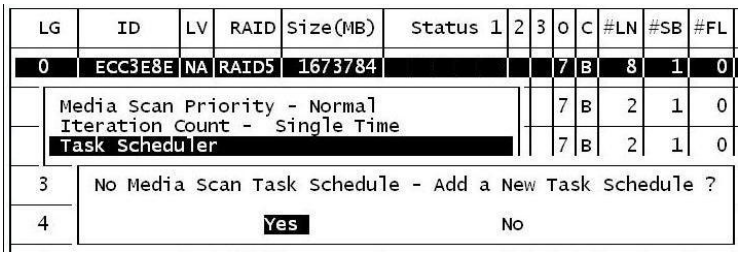

If there is no preset schedule, a confirm box as shown above will prompt.

Adding or Deleting a Schedule

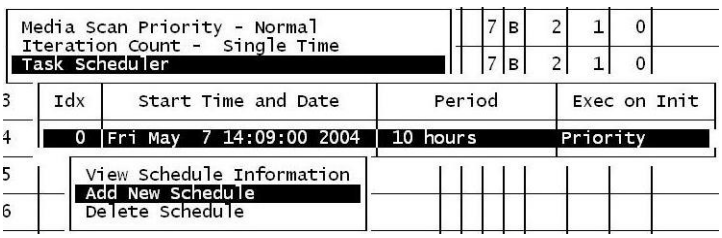

Press **[ENTER]** on an existing schedule to display the configuration options. You may choose to check information of a task schedule, to create a new schedule, or to remove a configured schedule.

Task Scheduler Options

To configure a task schedule, you have the option to do any of the following:

Set the following values/intervals:

Execute on Controller Initialization

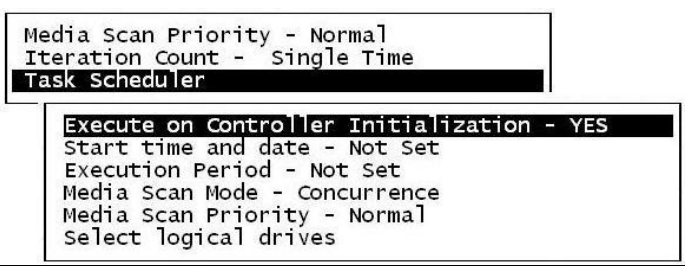

This option determines whether Media Scan is automatically conducted whenever the RAID system is reset or powered on.

Start Time and Date

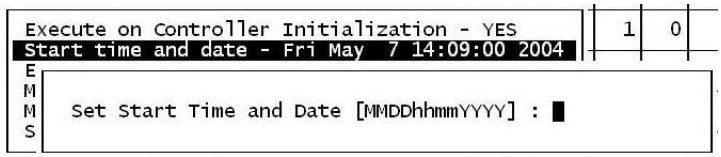

Enter time and date in its numeric representatives in the following order: month, day, hour, minute, and the year.

Execution Period

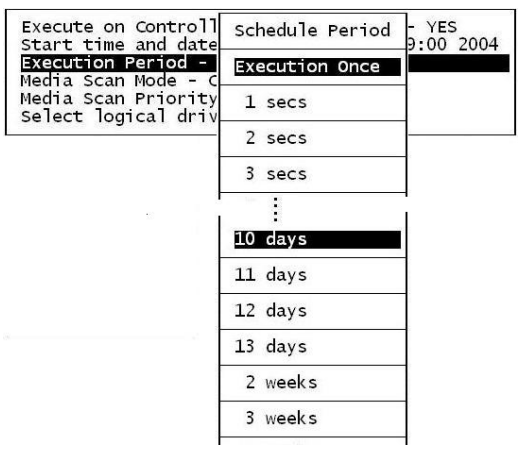

The scheduler memorizes the date and the time the actions are to be executed. Select one of the following:

- 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.

Media Scan Mode

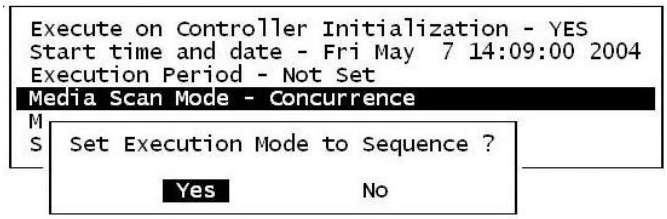

If the maintenance schedule includes more than one logical drive, the scan can be performed simultaneously or separately in a sequential order.

Media Scan Priority

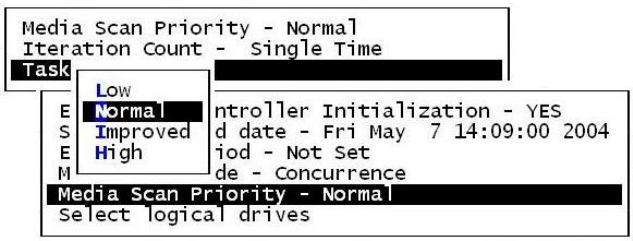

The scan priority determines how much of the system's resources will be consumed to perform the scheduled task. Select "Low" for better array performance and longer time to complete the media scan. Higher priority allows higher scan performance at the cost of reduced array performance.

Select Logical Drives

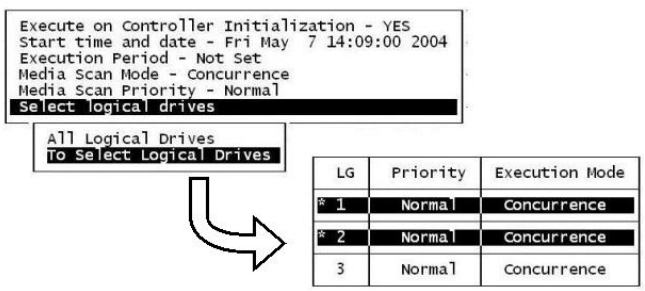

Press **[ENTER]** on "Select Logical Drives" to bring out a sub-menu. From there you may include all configured arrays or press **[ENTER]** on "To Select Logical Drives" to select one or more specific logical drive(s).

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 required.

When selection is done, press **[ESC]** to continue.

Confirming the Creation of a Task Schedule

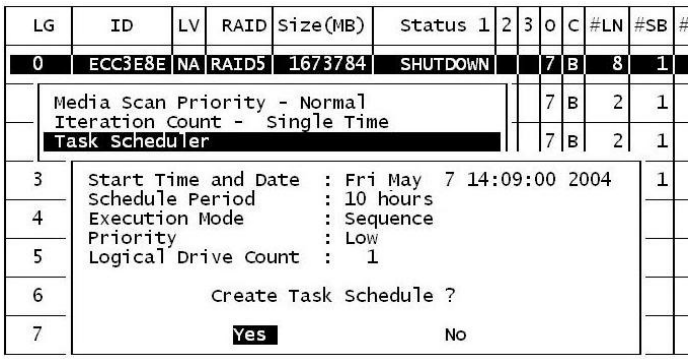

When finished with setting the scheduler options, press **[ESC]** to display a confirm box.

Verify all information in the box before choosing "**Yes"** to confirm and to complete the configuration process.

11.2.3 Regenerating Logical Drive Parity

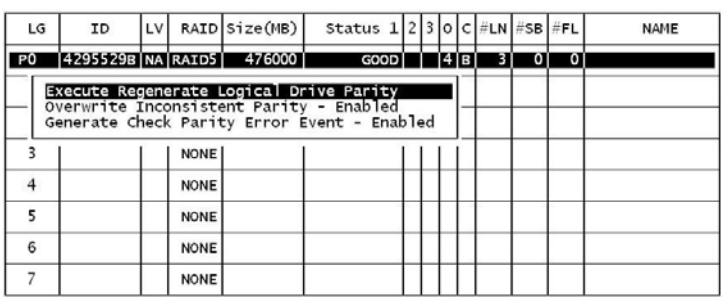

Parity regeneration is a function manually performed onto RAID-1/3/5 arrays to determine whether inconsistency has occurred with data parity.

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.

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.

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.*

11.2.4 Disk Array Parameters

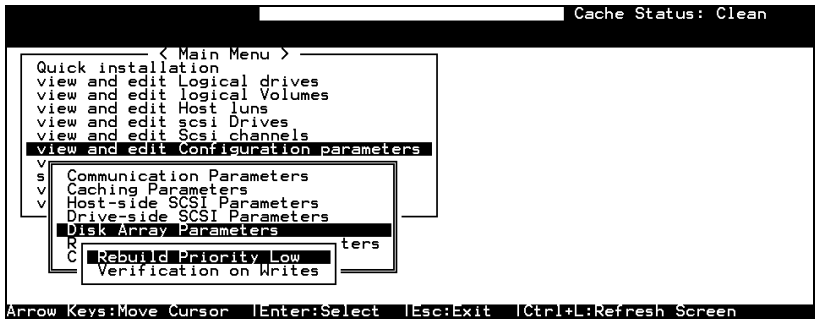

Select "View and Edit Configuration Parameters" on the Main Menu and press **[ENTER].** Choose "Disk Array Parameters," then press **[ENTER]** again. The Disk Array Parameters menu will appear.

Rebuild Priority

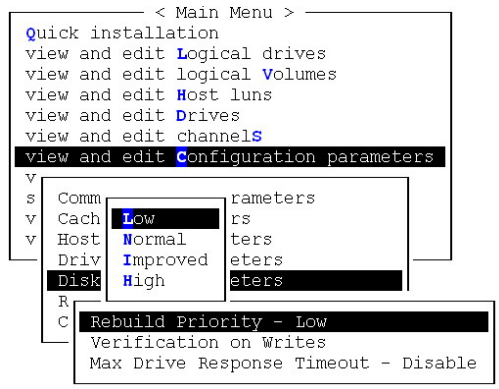

Choose "Rebuild Priority," then press **[ENTER]**. A list of the priority selections (Low, Normal, Improved, or High) displays. Move the cursor bar to a selection, then press **[ENTER]**.

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.

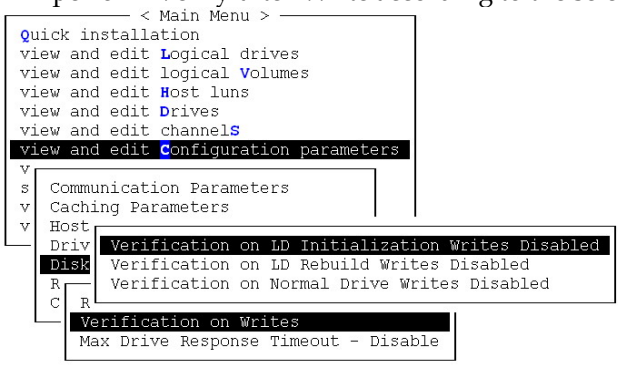

Move the cursor bar to the desired item, then press **[ENTER]**.

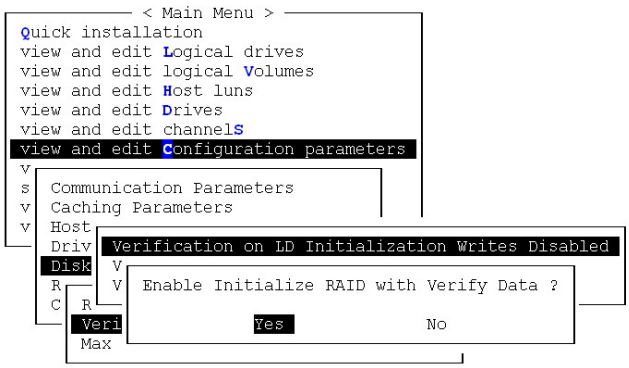

Choose **Yes** in the confirm box to enable or disable the function. Follow the same procedure to enable or disable each method.

IMPORTANT!

- *The "verification on Normal Drive Writes" method will affect the "write" performance of your RAID system.*

Chapter

12 iSCSI Options

This chapter is written for the configuration options exclusively implemented with iSCSI subsystems. iSCSI subsystems provide storage capacity in a network consisting of iSCSI initiators and targets. One or multiple iSCSI storage subsystems can be configured into a storage pool and then divided into iSCSI targets using the ID/LUN mapping techniques similar to traditional SCSI storage.

Infortrend's iSCSI storage supports iSCSI initiators that comply with the IETF iSCSI standard (RFC 3720). The configuration is simple. The configuration efforts required include the following:

1. RAID configuration process is identical to that described in the previous chapters.

2. Configuring host port IPs and connectivity settings. See page ***12-1***.

3. Setting CHAP security. See page ***12-3***.

4. Making RAID storage volumes available through the 1Gbps Ethernet host ports. See page ***12-4***.

12.1 Network Interface Settings

Acquiring An IP Address for the Host Ports

To access the Internet Protocol submenu, press **[ENTER]** to confirm a selection and use the arrow keys to navigate. The iSCSI port configuration options can be found under "View and Edit Configuration Parameters," "Communication Parameters," and then the "Internet Protocol (TCP/IP)".

Press **[ENTER]** on a host port you wish to configure. The identity of a host port is presented as:

"Port number [chip ID] MAC address – IP address"

CAUTION!

Avoid configuring the same number as the unique ID on different subsystems.

The MAC address of a host port is automatically acquired by referring to the subsystem's the pre-assigned serial number or the unique identifier that is manually assigned. Unless you assign a unique ID, the firmware refers to the serial number that is unique for each subsystem.

A manually assigned unique ID can cause problems in the following condition:

- *DHCP (Dynamic Host Configuration Protocol) severs automate network IP configuration and may "remember" these IPs. DHCP servers may remember the MAC addresses of DHCP clients and attempt to assign the same IP address to a client they consider as had been temporarily disconnected.*

- *If multiple subsystems in a local network are configured an identical ID, the DHCP server will assign the same MAC address to these subsystems causing troubles with network connection.*

As for the multiple Ethernet ports on a subsystem, the firmware automatically designates different ports with different MAC bits.

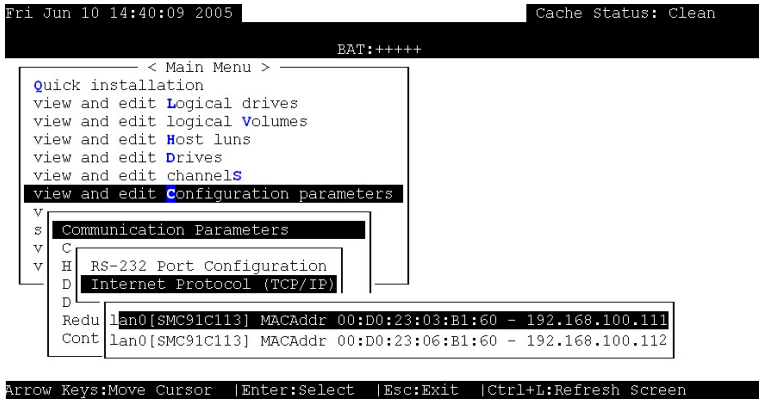

Press **[ENTER]** to select "Set IP Address".

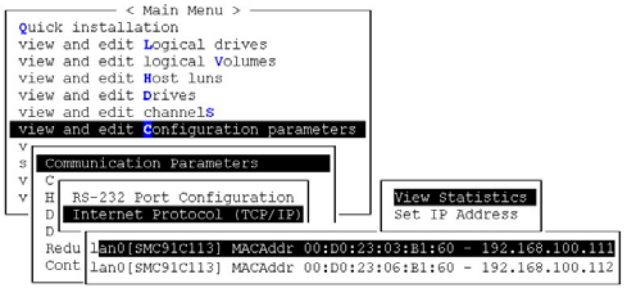

Press **[ENTER]** on the address line and enter the preferred address. Contact your network administrator to obtain a list of valid IP addresses. Provide the adequate NetMask and Gateway values accordingly.

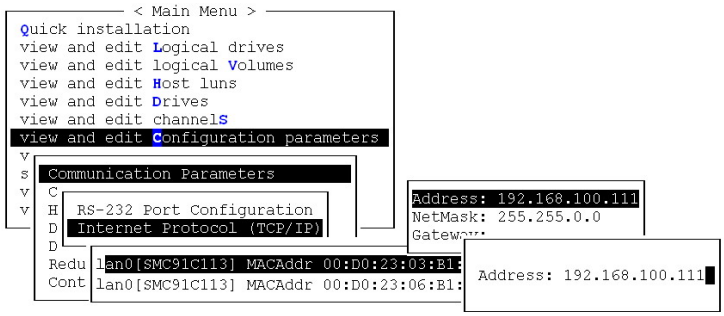

DHCP

Key in "DHCP" in the address field if the DHCP configuration method is preferred and supported by the local network.

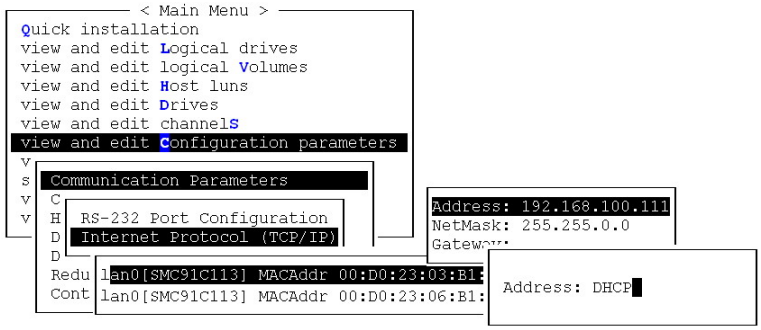

WARNING!

It is not recommended using the DHCP method. In cases of server failure, users of the DHCP method may risk the chance of data loss.

12.2 CHAP Login Authentication

CHAP stands for Challenge Handshake Authentication protocol. With this protocol, networked computers use the encrypted password to authenticate each other remotely.

1. To enable the CHAP authentication, select "View and Edit Configuration Parameters" in the Main Menu, "Host-side Parameters," and then press **[ENTER]** on "Login Authentication with CHAP" to activate the selection prompt.

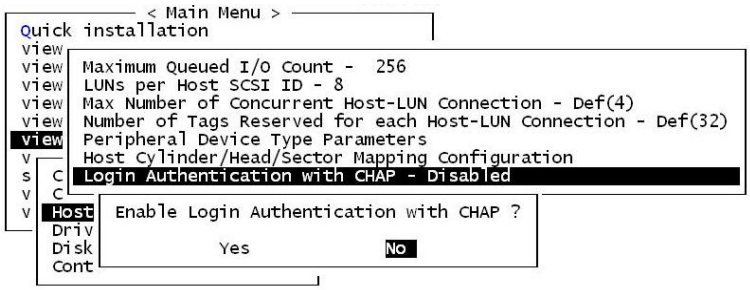

2. Configure the controller password. It is the password that will be used in the process of CHAP authentication. The controller password will be encrypted as the CHAP authenticated password. The password configuration dialog can be found in the "System Functions" submenu.

Note the following when configuring the CHAP authentication:

- The controller password and the controller name share a 15-character space. That means if you use 13 characters for the password, you will be left with only 2 characters for the controller name.

- Some login authentication utilities provided with iSCSI HBAs on Windows systems require a CHAP password of the length of at least 12 characters.

12.3 Mapping Storage Volumes to Host ID/LUNs

Infortrend's implementation of iSCSI protocols is fully compliant with the standardized SCSI architecture. A logical configuration of physical disk drives (a logical drive, a logical volume, or a logical partition) can be associated with one or more host channel IDs and LUN numbers. The methods used to present storage volumes to host ports are identical to the original SCSI-based RAID subsystems.

Please refer to ***Chapter 7*** for complete discussions of RAID configuration process.

IDs on Host Ports

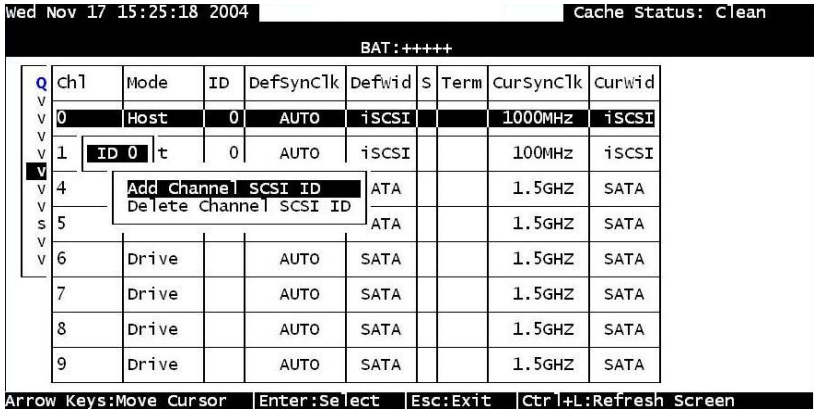

By default, an ID (ID 0) is available on each host port. If you prefer more IDs on the host ports, select a host channel by pressing **[ENTER]**. The ID edit mode dialog will appear. Up to 16 IDs are supported on a single host port.

LUN Mapping

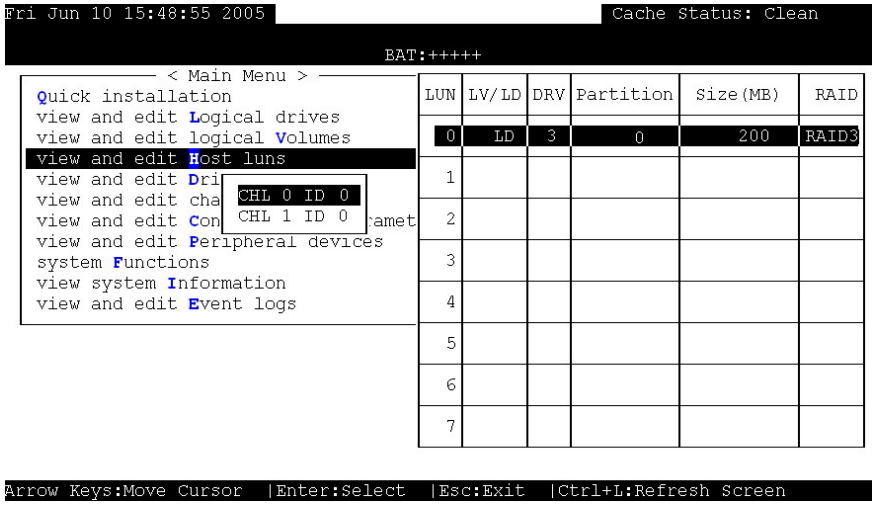

The LUN mapping process associates a storage volume with a host channel ID and LUN.

The configuration process is listed below:

1. Select "View and Edit Host luns" in the Main Menu, then press **[ENTER]**.
2. A list of host channel/ID combinations appears on the screen. Select a channel-ID combination by pressing **[ENTER].**
3. Choose mapping a "Logical Drive" or a "Logical Volume" on the drop box.

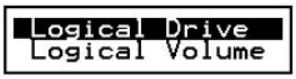

4. A list of LUN entries and their respective mappings will be displayed. To map a logical partition of a logical volume to a host LUN, select an available LUN entry (one not mapped yet) by moving the cursor bar to the LUN, then press **[ENTER]**.

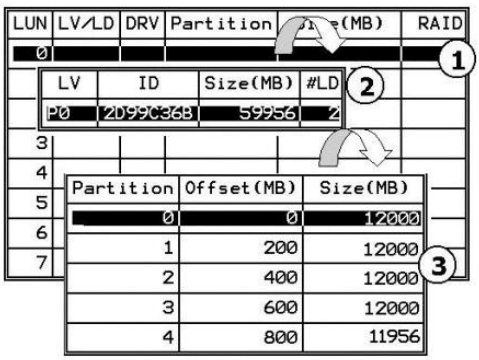

5. A list of logical volume(s) displays. Move the cursor bar to the desired logical unit, then press **[ENTER]**.

6. A list of logical partitions will prompt. Move the cursor bar to the desired partition, then press **[ENTER]**. If you have not partitioned the logical volume, the whole capacity will be displayed as one logical partition.

7. When prompted to "Map Host LUN," press **[ENTER]** to proceed.

8. When prompted to "Map Logical Volume?," select **Yes** to continue.

9. A prompt will display the mapping you wish to create. Choose **Yes** to confirm the LUN mapping you selected.

10. The detail in the confirm box reads: partition 0 of logical volume 0 will map to LUN 0 of ID 0 on host port 0.

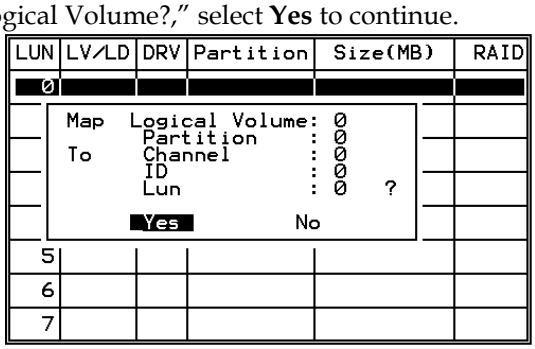

Chapter

13 Redundant Controller

This chapter is written for system integrators using Infortrend's board-only controllers and users who are interested in options directly related to dual-controller configurations. The EonStor dual-controller RAID subsystems come ready with related configuration settings. No user's configuration is required to bring up a controller pair.

13.1 Operation Theory

Sample topologies using redundant controllers can be found in the ***Installation and Hardware Reference Guide*** that came with your RAID subsystems. The proceeding discussions will focus on the theories behind and the configuration procedures for readying a redundant controller system.

Today's high-speed I/O interfaces pose increasing demands on signal quality. We assume that reliable inter-connections have already been made by combining controllers' I/O channels either through a common backplane or cabling connections. Depending on the enclosure design, signal paths for communications should have been strung between controllers over a common backplane. Your RAID subsystems, such as Infortrend's EonStor series, may come with preset IDs and channel mode settings, and require no further configurations.

Users who are familiar with the theories behind dual-controller configurations, please move ahead to ***Section 13.3 Configuration***.

This chapter covers the following topics:

Operation Theory

Preparing Controllers

- 13.2.1 Requirements
- 13.2.2 Limitations
- 13.2.3 Configurable Parameters

Configuration Procedures

- 13.3 Configuration
- 13.3.1 Setup Flowchart
    - Redundant Configuration Using Manual Setting
    - Controller Unique ID
    - Creating Primary and Secondary ID
    - Assigning a Logical Drive/Logical Volume to the Secondary Controller
    - Mapping a Logical Drive/Logical Volume to the Host LUNs
    - Terminal Interface View of Controller Failure
    - Forcing Controller Failover for Testing
- 13.3.4 When and How Is the Failed Controller Replaced?

13.1.1 Considerations Related to Physical Connections

Using SCSI-based Controllers

Figure 13 - 1 Dual-controller Setting Using SCSI-Based Controllers

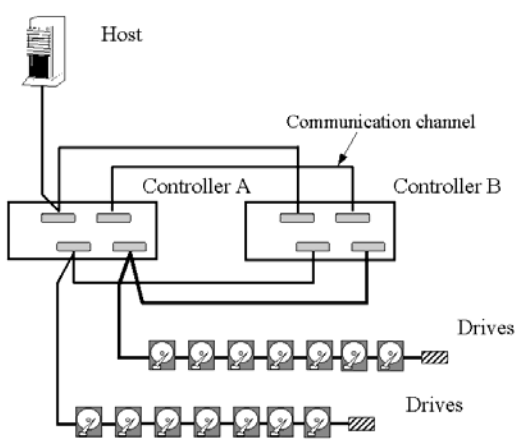

The physical connections between redundant controllers should be similar to the one shown above. The basic configuration rules are:

1. All channels should be connected to both controllers as diagrammed above either using SCSI cabling or over lines strung across a common backplane. Disk drives or cascaded drive enclosures are connected to both controllers.

2. Cached writes are constantly duplicated and distributed to both controllers' memory over a dedicated SCSI channel. The default path for controller communications (SCSI-based controllers) is channel 0.

   Channel 0 is also the default for the host interface; therefore, avail another channel(s) for host connection by changing its channel mode. See ***Chapter 5*** or ***Chapter 7*** for details about channel mode configuration.

3. SCSI channels should be terminated on both ends. It is recommended to use the termination jumpers on the RAID controllers to configure the SCSI bus termination setting. SCSI terminators are provided on board-only controllers. This design allows a failed controller to be removed after the controller failover process. As diagrammed above, terminators on Controller A should be disabled.

Using Fibre-based Controllers

The connection between dual-redundant controllers is more flexible with the Fibre-based controllers.

Figure 13 - 2 Dual-controller Setting Using Fibre-based Controllers

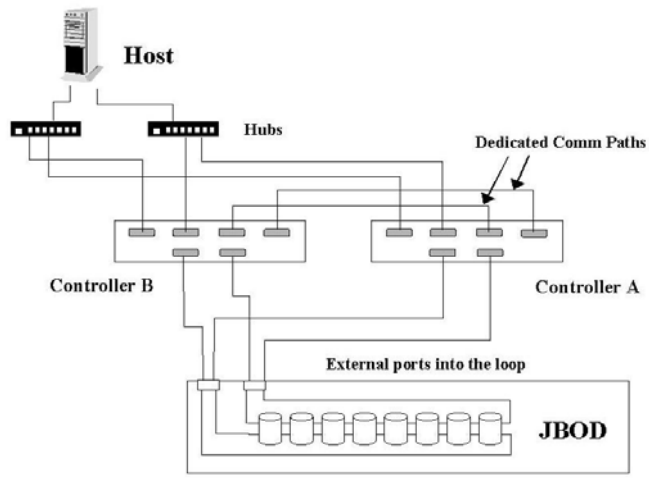

The basic configuration rules are:

1. All channels should be connected to both controllers as diagrammed above.

2. To reduce the chance of downtime, more than one FC switch can be used to connect to the host computer(s) for path redundancy. Drive-side dual loop is supported. Host-side dual loop requires the support of host management (multi-pathing or high availability) software.

3. For the Fibre-to-Fibre RAID controllers, there are two options with the communications loops between controllers:

1). Dedicated Communications Loops – "Dedicated RCC"

The first option is choosing one or two Fibre loops as the dedicated communications paths. Two is recommended for the sufficient bandwidth and path redundancy it provides.

Using two channels for communications offers greater throughput and therefore better performance.

2). Communications over Drive Loops – "Drive + RCC"

Configure all drive loops into the "Drive + RCC" mode to let them share the communications traffic. The controllers can automatically distribute the communications traffic across all drive loops.

Workflow is balanced among loops. Using the drive + RCC mode allows more channels to be used for drive connection. With six-channel controllers, for instance, there can be as many as two channels for host and four channels for drives (drive + RCC). All channels can be used for I/O traffic while the system is still benefited from controller communications.

13.1.2 Grouping Hard Drives and LUN Mapping

Listed below are the array setting options that need to be considered when planning a dual-controller RAID subsystem:

1. **H**ow many logical drives, logical volumes, or logical partitions, and each of what sizes?

2. **S**ystem drive mapping (Primary/Secondary controller IDs): how many storage volumes will appear to which host ports, and managed by which controller?

3. **W**ill those storage volumes be accessed in a multi-host or multi-path configuration? Will they be used as shared as volumes?

4. **F**ault Tolerance: Configure the RAID controllers so that they can failover and failback in the event of single controller failure. See ***13.1.3 Fault Tolerance*** for more details.

Different Controller Configuration Rules

With subsystems built on the new ASIC266 chipset, a slot A-slot B controller position principle is applied. The new position principle helps to ensure the location of a dominating Primary controller. Array management, ID/LUN mapping and array operation remain basically unchanged using these two principles.

The old principle applies to all controller/subsystems built on the previous ASIC models. The old principle places no significant regulations on the locations of a Primary or Secondary controller.

The new principle defines the RAID controller installed in Slot A, usually the upper controller slot, as the Primary controller. The factory defaults ensures the Slot A controller behaves as a Primary controller. In normal operating conditions, the Slot A controller should always be the Primary controller. In the conditions described below, a slot A controller might temporarily serve as a Secondary controller.

1. If the controller in slot A fails, the original Secondary controller takes over and becomes the Primary controller.

2. If the slot A controller is replaced by a new controller, the slot A controller temporarily serves as the Secondary controller.
3. Once the subsystem is reset, the slot A controller resumes the role of the Primary.

Logical Drives, Logical Volumes, and Logical Partitions

Listed below are the basics about configuring a logical drive for a dual-controller subsystem:

- All configuration options are accessed through the Primary RAID controller. With the new ASIC 266 subsystem models, the RAID controller installed in the upper controller slot, Slot A, is the Primary controller.

  With the older models, a controller showing a flashing <P> digit on the upper right corner of its LCD screen is the Primary. In redundant mode, two controllers behave as one, and there is no need to repeat the configuration on another controller.

- Disk drive and array configuration processes are the same for subsystems using single or dual-active controllers.

- Using logical drives as the base units, system workload can be distributed to different RAID controllers. Logical units can be manually assigned to different (Primary or Secondary) controllers to facilitate the active-active configuration.

  Users can designate a host ID either as a Primary or as a Secondary ID. The Primary and Secondary IDs differentiate the workload assigned to different RAID controllers.

- There is no limitation on drive allocations. The members of a logical drive do not have to come from the same drive channel. Grouping drives from different drive channels helps reduce the chance of downtime caused by a channel bus or cable failure.

- Each logical drive can be configured in a different RAID level and several logical drives can be striped together to compose a larger logical volume.

- 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 options. Each of them can be "**associated**" with one or more host ID/LUNs. Each of these

associated host ID/LUNs appears to the host operating system as a virtual hard drive.

Figure 13 - 3 Grouping Hard Drives

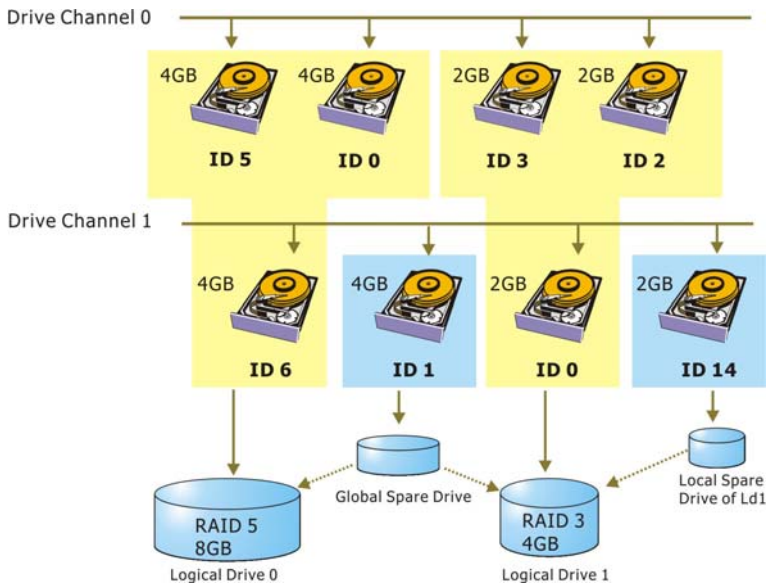

- As diagrammed above, array composition can be very flexible. You may divide a logical drive or logical volume into several partitions as diagrammed below, or use the entire logical drive as a single partition, with or without the support of one or several spare drives.

Figure 13 - 4 Partitioning of Logical Units

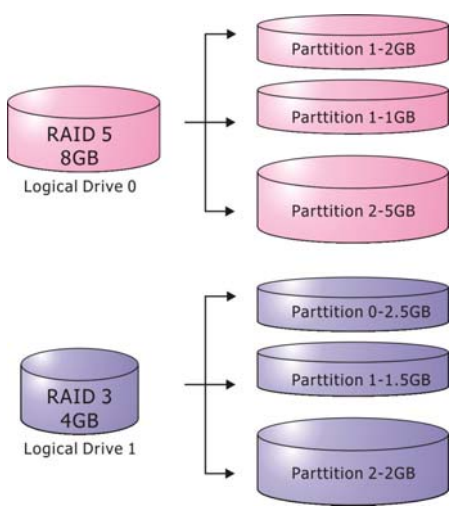

- Each logical unit can be associated (mapped) with one or more host IDs (pre-configured as a Primary or Secondary controller ID) or the LUN numbers under these host IDs.

System Drive Mapping

Primary and Secondary IDs

- **Host Channel**

  When controllers are successfully combined, the array capacity is available through host port IDs, and these IDs are available as "Primary" or "Secondary" IDs. In addition to the default IDs, more IDs can be added through the "Primary ID" or "Secondary ID" pull-down menus from the user interface.

- **Drive Channel**

  Since all channels are strung between two controllers, each channel is connected to two chip processors, and each processor must occupy one channel ID. In redundant mode, both a Primary ID and a Secondary ID must be present on every drive channels.

- **Relationship between Controllers**

  - New Configuration Rule:

  The controllers' identification is determined by the controller slots into which they are installed. A new and normally functioning subsystem is managed by the controller installed in Slot A, the Primary controller. If the Primary controller fails, the Secondary controller takes over. The RAID controller in slot B becomes the Primary. The slot B controller returns the control when the failed controller in Slot A is replaced and the subsystem is reset.

  - Old Configuration Rule:

  There are no specific regulations on the locations of the Primary/Secondary controllers, before or after a controller failure. They are indicated by the flashing digit showing <P> or <S> on the upper right corner of the LCD screen.

- **Create IDs**

  You may have to create Primary and Secondary controller IDs separately on the host and drive channels if these IDs are not available. The configuration procedure will be discussed in ***Section 13.3Configuration.***

Mapping

- A logical unit made available through a Primary controller ID will be managed by the Primary controller, and that through a Secondary controller ID by the Secondary controller.

Figure 13 - 5 Mapping Array Volumes to Different Controllers

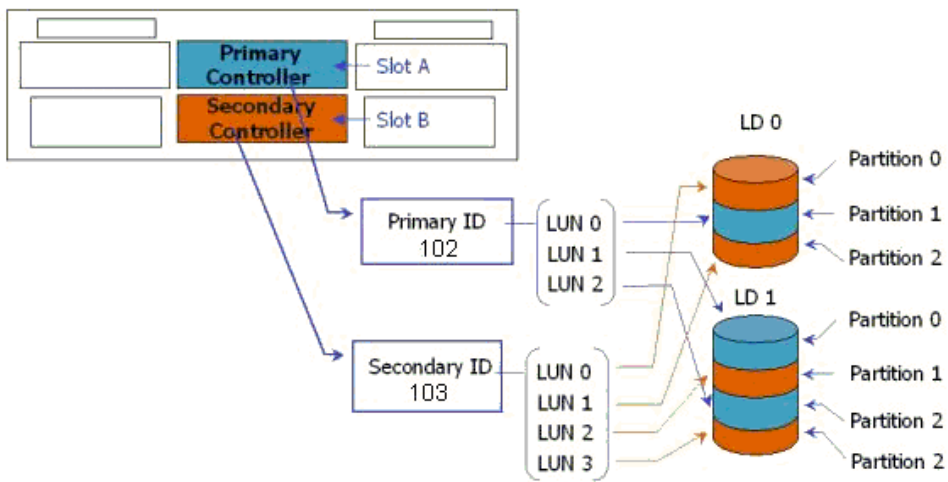

- Each channel ID (or an LUN under target ID) will appear as one virtual storage volume to the host operating system.

Figure 13 - 6 Mapping System Drives (Mapping to LUNs)

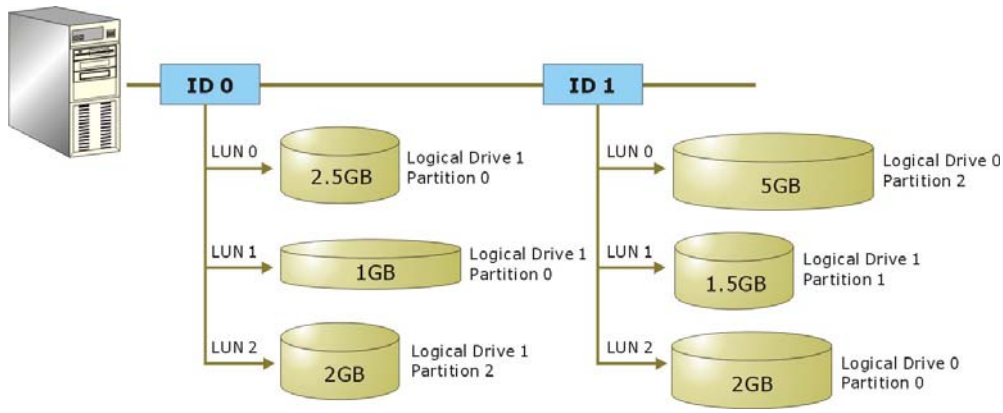

- The diagram above displays a single host computer equipped with two HBA cards allowing the connection of dual I/O paths. A host port ID is presented on each host port as either a Primary ID or a Secondary ID. Users may then map (associate) any logical configuration of drives equally to these IDs or LUN numbers. Workload can thus be distributed across two host ports and managed by both controllers.

Figure 13 - 7 Mapping System Drives (an Example on a SCSI bus)

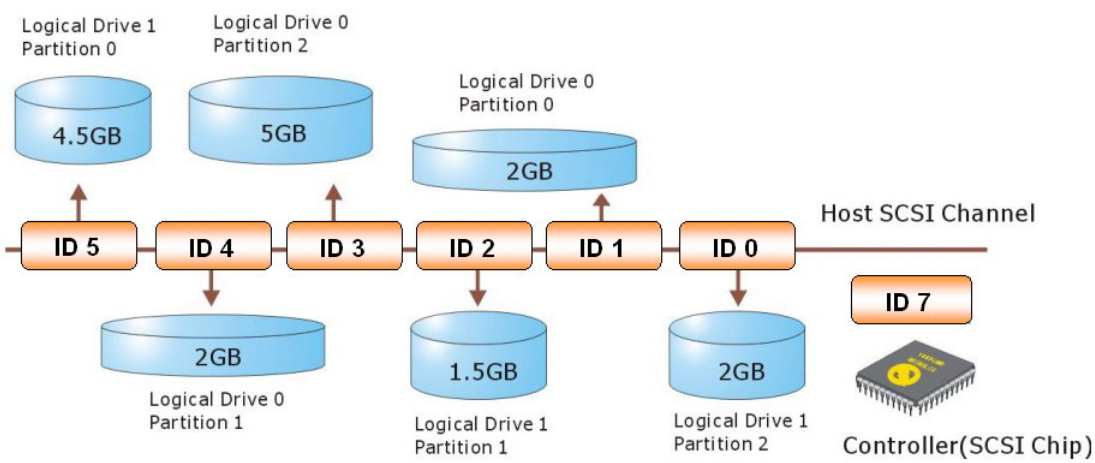

- Some older operating systems do not read multiple LUNs under a target ID. As diagrammed above, you may have the host channel to present several IDs and map logical configurations to these IDs. Each of these IDs can be identified as "Primary ID" or "Secondary ID." As a rule for most operating systems, each configuration unit will be mapped to LUN 0 under each ID.

13.1.3 Fault Tolerance

What Is a Redundant Controller Configuration?

1. Hardware failures can occur. A simple parity error can sometimes cause a RAID system to completely hang up.
2. Having two controllers working together will guarantee that at least one controller will survive catastrophes and keep the system working.
3. 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 and require only minimum efforts to restore the original configuration.

How Does Failover and Failback Work?

A. Channel Bus

Below is a sample illustration of the redundant controller operation:

Figure 13 - 8 Redundant Controller Channel Bus

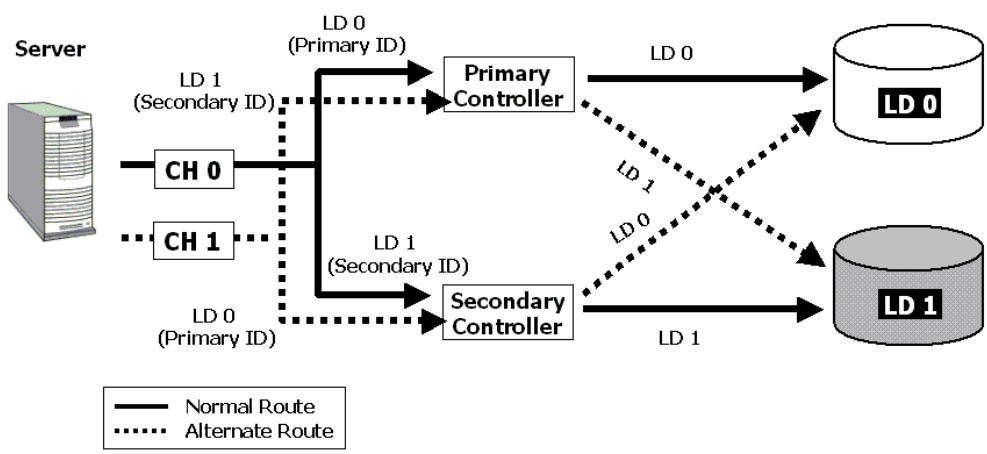

The host computer is connected to both the Primary and the Secondary controllers. Each controller has two of its I/O channels

assigned as the host ports, and the rest of I/O channels assigned to drive connection.

There are two logical drives. Logical drive 0 is assigned to the Primary controller (associated with the Primary ID), and logical drive 1 assigned to the Secondary controller (associated with the Secondary ID). If one controller fails, the surviving controller will manage the logical drive that was previously managed by the failed controller via the once inactive ID (the standby ID, see the IDs beside the dash lines).

ID mapping is synchronized between the controllers. If one controller fails, the surviving controller keeps a replica of all configuration data. The related ID mapping on the controllers looks like the tables below:

Table 13 - 1 ID Mapping Status (Normal Operation)

Channel	ID	Status	Target Chip
0	**0** (Primary ID)	Active	Primary Controller channel 0
	1 (Secondary ID)	Standby	Secondary Controller channel 0
1	**1** (Secondary ID)	Active	Secondary Controller channel 1
	0 (Primary ID)	Standby	Primary Controller channel 1

In the event of controller failure (say, the Primary controller fails), the once inactive IDs (chip) become active:

Table 13 - 2 ID Mapping Status (Single Controller Failure)

Channel	ID	Status	Target Chip
0	**0** (Primary ID)	~~Active –~~	~~Primary Controller channel 0 –~~ **Failed!**
	1 (Secondary ID)	Standby- becomes **Active!**	Secondary Controller channel 0
1	**1** (Secondary ID)	Active	Secondary Controller channel 1
	0 (Primary ID)	~~Standby~~	~~Primary Controller channel 1-~~ **Failed!**

Figure 13 - 9 Controller Failover

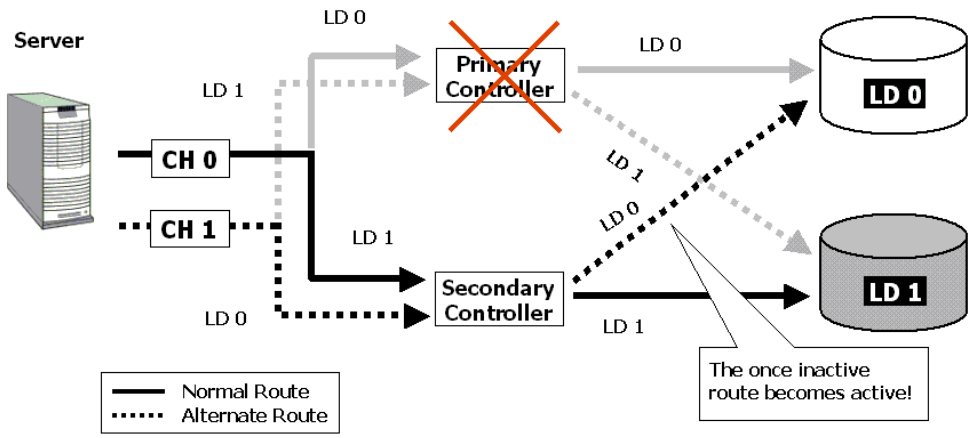

For every channel that is actively serving I/Os, there is another channel on the alternate controller that stays idle and will inherit the task should its counterpart fail.

As long as I/O bus bandwidth is not a concern, many active IDs can co-exist on single or multiple host channels. Standby chips may not be necessary.

B. 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.

Replacing a Failed Unit:

The replacement controller should have the same amount of memory and run the same version of firmware. However, it is inevitable a replacement controller will be running later revisions of firmware. To solve this problem, **you should manually downgrade the replacement's firmware.**

Your system vendor should be able to provide an appropriate version of firmware and replacement controller.

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 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 13 - 10 Controller Primary-Secondary Relative Positions

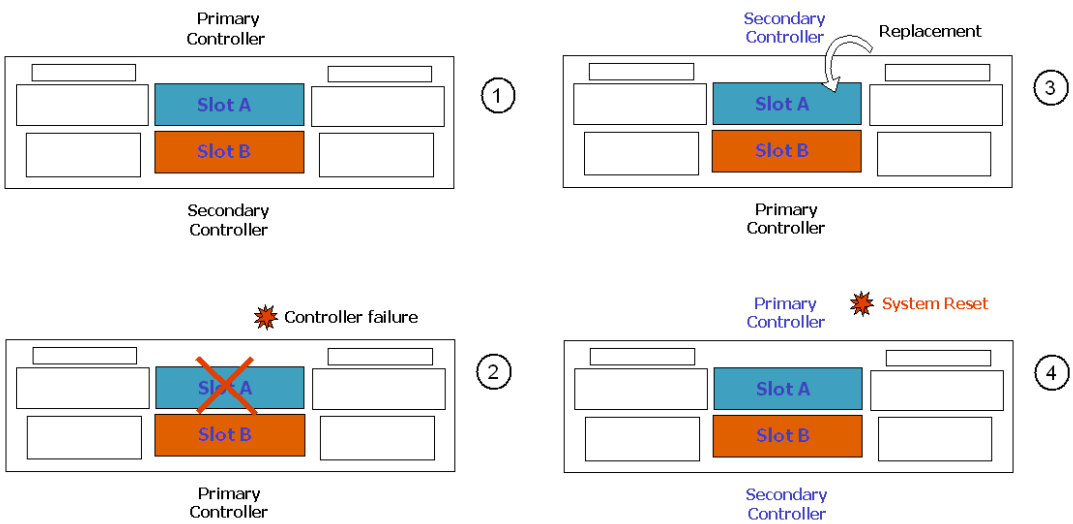

A typical controller failover and failback process is diagrammed above.

Status 1: Subsystem operating normally. Slot A controller is the Primary controller by factory default.

Status 2: Slot A controller fails. Slot B controller inherits the Primary role.

Status 3: The failed controller in Slot A is replaced by a healthy one. The replacement controller becomes the Secondary controller temporarily.

Status 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 the surviving controller. If the surviving controller fails to re-establish this connection, you can also choose to "de-assert" the replacement controller through the existing controller so that both will serve the original system drive mapping.

C. Active-to-Active Configuration:

Active-to-active configuration conducts all system resources to performance. Storage volumes can be equally assigned to both controllers so that both controllers are actively serving I/Os. This allows a flexible association between logical units and host ID/LUNs. Workload can then be manually distributed to different controllers.

D. Traffic Distribution

The diagram below illustrates a four-channel configuration using channel 0 as the communications path. Channel 1 serves as the host interface and multiple IDs have been created to facilitate active-active operation. Each controller occupies either a Primary ID or a Secondary ID on drive channels. There are two logical drives. One logical drive is assigned to the Primary Controller and the other the Secondary Controller.

In the event when one controller fails, the existing controller will inherit IDs from the failed controller and continue I/Os.

Figure 13 - 11 Traffic Distribution

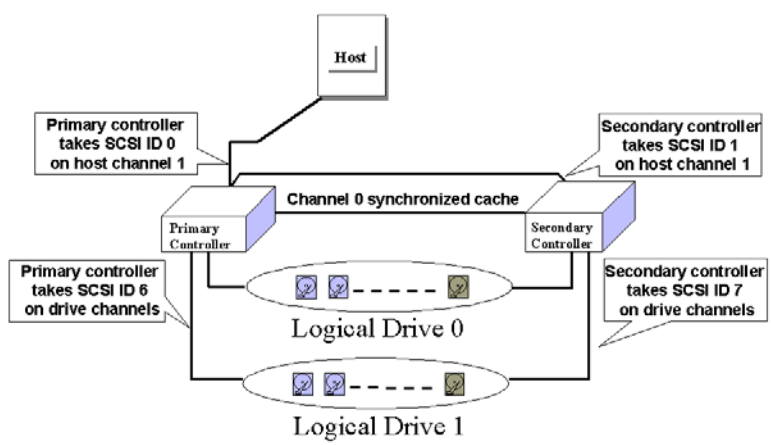

	Logical Drive 0	Logical Drive 1
Host LUN Mapping	ID0 / LUN* (PID)	ID1 / LUN* (SID)

Logical Drive Assignment	Primary	Secondary
Drive Channel	2	3

Users can assign a logical unit either to the Primary or to the Secondary Controller. Once the assignment is done, logical unit(s) assigned to the Primary Controller can only be mapped to the Primary IDs on the host channel; Logical unit(s) assigned to the Secondary Controller can only be mapped to the Secondary IDs on the host channel.

The channel ID (Primary/Secondary) assignment for a SCSI-based controller should look like this:

	Primary Controller ID	Secondary Controller ID
Host Chl SCSI ID	PID = 0	SID = 1
Drive Chl SCSI ID	7	6 suggested

NOTE: The EonStor series has preset drive-side IDs. No configuration change is required.

Figure 13 - 12 Controller Failover

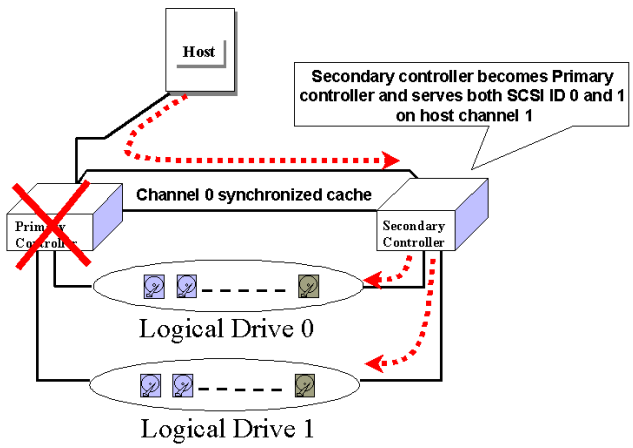

E. Controller Failure

Controller failure is managed by the surviving controller (regardless of its original role as Primary or Secondary). The surviving controller disables and 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 on the failed controller is off. The LCD on the surviving controller displays controller failure message.
- The surviving controller sounds an alarm.
- The "ATTEN" LED is flashing on the existing controller.
- The surviving controller sends event messages to notify of controller failure (indicating its partner has failed).

Connection

The channels of the two controllers that are connected together must be the same. For example, if controller A uses channel 2 to connect a group of drives, controller B must also use channel 2 to connect to the same group of drives.

13.2 Preparing Controllers

13.2.1 Requirements

Cabling Requirements

Communications Channels

- Controller Communications (Cache Synchronization) Paths:

Controller	RCC cable
SentinelRAID	A SCSI cable (CH 0)
EonRAID 2510FR	Dedicated RCC or RCC over drive loops
EonStor	Pre-configured RCC routes over the system backplane; may be SCSI, Fibre, or SATA data paths. These data paths cannot be changed!

- Using one or two of the I/O channels for controller communications (as listed above) is necessary when write-back caching is preferred. Each controller keeps in its memory a replica of cached writes directed to its counterpart.
- With dual-active controller configuration, write-back cache will be disabled if no RCC path can be found.
- If controllers are running in write-back mode, a battery module is recommended for each controller.
-

Out-of-Band Configuration

- RS-232C cable (for terminal interface operation) connection. Normally a Y-cable will be included with dual-controller subsystems.
- Ethernet connection: If management through Ethernet is preferred, connect the Ethernet interface from both controllers to ports on an Ethernet hub. In the event of controller failure, the IP address assigned to the Primary Controller will be inherited by the surviving controller. The Ethernet port connection (management session) will be continued. An operator may have to re-enter the IP address to re-connect the controller/subsystem.

Controller Settings: General Procedure for Stand-alone Controllers

1. ***Enable Redundant Controller Configuration***

   "Main Menu"→ "View and Edit Peripheral Devices"→ "Set Peripheral Device Entry"→ "Redundant Controller Enable/Disable"

   Note that some dual-controller models have this configuration set as "enabled" through the iappend utility. This feature cannot be disabled on the models running these specialized firmware. For example, if you are upgrading the single-controller EonRAID 2510FS to a dual controller configuration, the upgrade controller will arrive with appropriate preset configurations. No configuration is required.

2. ***Controller Unique Identifier***

   Set unique identifier for each controller. "View & Edit Peripheral Devices"→ "Controller Parameters"→ "Controller Unique Identifier." Enter a hex number between 0 and FFFFF (firmware 3.25 and above) for each controller.

3. ***Creating Primary and Secondary IDs on Drive Channels***

   "View and Edit SCSI Channels"→ Choose a Drive Channel→ "Primary/Secondary Controller SCSI ID."

4. ***Creating Primary and Secondary IDs on Host Channels***

   "View and Edit SCSI Channels"→ Choose a host channel→ "View and Edit SCSI ID"→ Choose a SCSI ID→ "Add/Delete Channel SCSI ID"→ "Primary/Secondary Controller"→ Add SCSI ID from the list. Reset the controller for the configuration to take effect.

5. ***Creating Logical Configurations of Drives and Assigning Each of Them Either to the Primary or the Secondary Controller***

   "View and Edit Logical Drives"→ Select a RAID level→ Select member drives→ "Logical Drive Assignments"→ Create Logical Drive.

6. ***Map Each Logical Configuration of Drives to the Primary/ Secondary ID on host channel(s)***

   "View and Edit Host LUN"→ Choose a "Host Channel-ID-Controller" Combination→ Choose Logical Drive/Logical Volume/Logical Partition→ Map to Host ID/LUN (Create Host LUN Entry).

NOTE:

- *The redundant controller function can be enabled via the LCD keypad panel or a terminal emulation program. Section **13.3 Configuration** describes the procedures for using terminal emulation and the LCD front panel. The same result can be achieved regardless of the interface used.*

13.2.2 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 installed with a battery module.

- The takeover process should take less than one second (using SCSI or Fibre for controller communications) to complete.

- In redundant mode, each controller takes an ID on each drive channel bus. This means the maximum number of disk drives on a SCSI bus is 14.

- Connection through Fibre switches can be necessary for joining host (Fibre) interfaces between controllers. The EonRAID 2510FS-XRH is an exception. Its type-1 ports come with an onboard hub (PBC).

- The controller defaults for ID settings are listed below:

Host interface	Host channel (Primary/Secondary)	Drive channel (Primary/Secondary)
SCSI	0 / 1…	7 / 6
Fibre	102 / 103…	119 / 120

- SCSI IDs 8 (PID) and 9 (SID) are sometimes used on the drive channels of the SCSI-based dual-controller subsystems.

13.2.3 Configurable Parameters

NOTE:

- *The following does not apply to the EonStor series which should run properly*

Primary or Secondary

Determinants:

- Upon initial setup, the controllers will decide between themselves which is the Primary and which is the Secondary by referring to each controller's serial no. One that has a larger serial number becomes the Primary.

- The controller firmware recognizes the two controllers used in a redundant configuration as Primary or Secondary. Two controllers share I/O loads while behave as one controller.

- Once the redundant configuration takes effect, user's configurations and settings can only be done through the Primary controller. The Secondary controller then synchronizes with the configuration of the Primary controller, making the configurations of two controllers exactly the same.

- The two controllers continuously monitor each other. When a controller detects that the other controller is not responding, the working controller will immediately take over and disable the failed controller. However, it is not predictable which one of the controllers will fail. It is necessary to connect all other interfaces, e.g., COM1 serial and Ethernet ports, to both controllers so that a surviving controller can readily continue all the services provided for the RAID system.

Active-to-Active Configuration

Users can freely assign any logical configuration of drives to both or either of the controllers, and then map the logical configurations to the host channel IDs/LUNs. I/O requests from host computer will then be directed to the Primary or the Secondary Controller accordingly. The total drive capacity can be divided and equally serviced by both controllers.

The active-to-active (or 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).

Active-to-Standby Configuration

By assigning all the logical configurations of drives to one controller, the other controller will stay idle and becomes active only when its counterpart fails.

Cache Synchronization

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 might occur when the surviving controller takes over and attempts to complete the writes.

Data inconsistency can be avoided using one or several of the I/O channels as the communications paths between the controllers. The cached data is always synchronized in each other's memory. Each controller saves an exact replica of the cache content on its counterpart. In the event of controller or power failure, the unfinished writes will be completed by the surviving controller.

Battery Support

Unfinished writes will be cached in memory in write-back mode. If power to the system is discontinued, data stored in the cache memory will be lost. Battery modules can support cache memory for a period of several days allowing the controller to keep the cached data. When two controllers are operating in write-back mode, it is recommended to install a battery module to each controller.

13.3 Configuration

General Configuration Steps:

Listed below are the steps necessary to configure a redundant controller subsystem:

1. The partner controllers have no preset roles. When the two controllers are powered on later, the firmware will negotiate and determine which of them is the Primary controller.

2. You may change the RCC communications paths between controllers using models such as the EonRAID 2510FS-4RH or the EonRAID 2510FS-6RH. If a channel is used as the communications channel, the firmware will display the channel status as "**RCCOM** (Redundant Communications Connection)" or "**Drive+RCC**." The dedicated RCC channel(s) will be excluded from the use of host or drive connection.

3. When powering on both controllers together, the LCD will display "**RC connecting**." After the controller negotiation is completed, the communications between controllers should be established.

4. Configure your SCSI/Fibre channels as host or drive. The default configuration for SCSI channel termination is "enabled." Please refer to your controller/subsystem ***Installation and Hardware Reference Guide*** and examine whether the termination jumpers on the controller backplane are shunted. If the associated jumpers are shunted, the SCSI channels will be terminated on the controller side no matter the firmware setting is "enabled" or "disabled."

5. Create both a "Primary ID" and a "Secondary ID" on every drive channel. If the controller comes with defaults, it is recommended to use them.

6. Reset the controller for the configuration to take effect.

7. Create Logical drives/logical volumes and assign each logical unit to the Primary or to the Secondary controller.

8. Proceed with host LUN mapping. After mapping each logical unit to one or more Primary or Secondary ID/LUN on the host channel(s), the redundant controller configuration is complete.

13.3.1 Setup Flowchart

Figure 13 - 13 Redundant Controller Configuration Flowchart

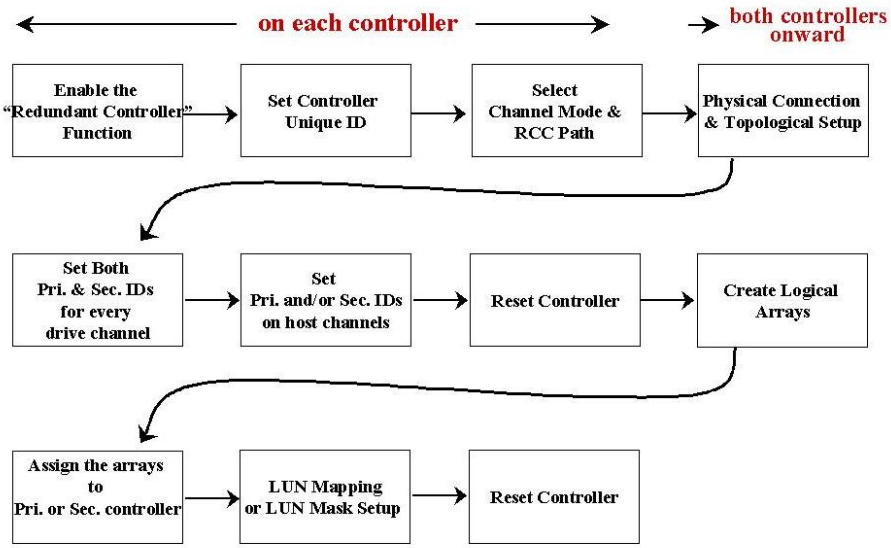

NOTE:

- *Some of Infortrend's dual controller configurations come with preset IDs for users' ease of configuration. It is, however, always best to check these IDs before proceeding with further configuration.*

13.3.2 Via Front Panel Keypad

Redundant Configuration Using Manual Setting

Power on Controller 1. Make sure Controller 2 is powered off.

1. Enable Redundant Controller

Press **[ENT]** for two seconds on the front panel of Controller 1 to enter the Main Menu. Use the up or down arrow keys to navigate through the menus. Choose "View and Edit Peripheral Dev," then press **[ENT]**.

```
View and Edit
Peripheral Dev
```

Choose "Set Peripheral Devices Entry," then press **[ENT]**.

```
Set Peripheral
Devices Entry
```

Choose "Redundant Ctlr Function__", and then press **[ENT]**. (**NOTE**: The current setting will be displayed on the LCD.) If this controller has never been set as a redundant controller before, the default setting of the redundant controller function is "Disabled." The message "Redundant Ctlr Function Disable" will be displayed on the LCD. Press **[ENT]** to proceed.

```
Redundant Ctlr
Function Disable
```

```
Enable Redundant
Ctlr          ?
```

2. Controller Unique ID

Enter "View and Edit Config Parms"-> "Controller Parameters". Use the up or down arrow keys to find "Ctlr Unique ID- xxxxx".

```
View and Edit
Config Parms
```

```
Controller
Parameters   ..
```

```
Ctlr  Unique
ID- 00012     ?
```

This value will be used to generate a controller-unique WWN node name and port names and to identify the controller during the failover process. Enter a hex number from 0 to FFFFF and press **[ENTER]**. The value you enter should be different for each controller.

Power off Controller 1, and then power on Controller 2. Set Controller 2 also to "**Primary**" as described previously. Power off Controller 2.

When both controllers' redundant controller function is set to the "Primary" setting, the controllers will decide between themselves which will be the Primary controller.

Starting the Redundant Controllers

Power on all hard drives and the two controllers. If drives are installed in a drive enclosure, wait for the drives to be ready, then power on the enclosure where the RAID controllers are installed.

```
RC connecting...
<ENT> to cancel
```

The message "RC (redundant controller) connecting... <ENT> to cancel" will appear on the LCD display of the two controllers. After a few seconds, the Primary Controller will startup with the model number and firmware version displayed on the LCD, while the Secondary Controller will display the message "RC Standing By.. <ENT> to Cancel" on its LCD.

A few seconds later, the LCD display on the Secondary Controller will be similar to the LCD display on the Primary Controller. The upper right corner of LCD will then display a "P" or "S," meaning "Primary" or "Secondary" respectively.

During normal operation, the controllers continuously monitor each other. Each controller is always ready to take over for the other controller in the unlikely event of a controller failure.

The Primary and Secondary Controllers synchronize each other's configurations at frequent intervals through the communications channel(s).

Creating Primary and Secondary ID

Drive Channel

Enter "View and Edit SCSI Channels." Press **[ENT]** and use the up or down arrow keys to select the host or drive channel on which you wish to create Primary/Secondary IDs.

```
View and Edit
Channels · ...
```

```
CH1=Drive PID=7
SID=NA SXF=80.0M
```

Press **[ENT]** to proceed.

Use the up or down arrow keys to select "Set SCSI Channel Pri. Ctlr ID .." or " "Set SCSI Channel Sec. Ctlr ID …" Press **[ENT]** to proceed.

```
Set SCSI Channel
Sec. Ctlr ID ..
```

Use the up or down arrow keys to select a SCSI ID and press **[ENT]** to confirm. The configuration change will take effect only after the controller has been reset.

```
Set Sec. Ctlr
ID:NA to ID: 6?
```

Host Channel

The process of creating Primary and Secondary IDs on host channels is basically the same.

```
CHL=0 ID=0
Primary Ctlr  ..
```

In "View and Edit SCSI Channels", press **[ENT]** to select a host channel. Use the up or down

arrow keys to select "Set SCSI Channel ID". A pre-configured ID will appear, press **[ENT]** to proceed. Use the up or down arrow keys to select "Add Channel SCSI ID" and then press **[ENT]** for two seconds on the "Primary Controller" or "Secondary Controller?" to proceed.

```
Add Channel
ID    ..
```

```
Primary
Controller ?
```

When prompted by this message, use the up or down arrow keys to select an ID. Press **[ENT]** to confirm.

```
Add CHL=0 ID=2
Primary Ctlr  ?
```

A message will prompt to remind you to reset the controller. Press **[ENT]** to reset the controller or press **[ESC]** to move back to the previous menu. The change of ID will only take effect after the controller has been reset.

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

Assigning a Logical Drive/Logical Volume to the Secondary Controller

A logical drive, logical volume, or any of its logical partitions can be assigned to the Primary or Secondary Controller. By default, a logical drive is automatically assigned to the Primary Controller. It can be assigned to the Secondary Controller if the host computer is also connected to the Secondary Controller.

Note that the partitions of a logical drive that has previously been assigned to the Secondary controller will automatically be assigned to the Secondary Controller.

Press **[ENT]** for two seconds on the front panel of the Primary controller to enter the Main Menu.

Use the up or down arrow keys to navigate through the menus. Choose "View and Edit Logical Drives", then press **[ENT]**.

```
View and Edit
Logical Drives
```

Create a logical drive or choose an existing logical drive, then press **[ENT]** to see the logical drive menu.

Choose "Logical Drive Assignment..," then press **[ENT]**.

```
Logical Drive
Assignment..
```

The message "Redud Ctlr LG Assign Sec Ctlr?" will appear. Press **[ENT]** for two seconds to confirm. The logical drive has now been assigned to the Secondary Controller.

```
Redud Ctlr LG
Assign Sec Ctlr?
```

Map the logical drive (or any logical unit) to a host ID or LUN number under the designated Secondary controller ID. The host channel must have a "Secondary" SCSI ID created. (Create the Secondary controller's SCSI ID on host channel and add a SCSI ID to every drive channel in "View and Edit Channels").

Mapping a Logical Drive/Logical Volume to the Host LUNs

Choose "View and Edit Host Luns" from Main Menu and press **[ENT]** to proceed.

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

Use the up or down arrow keys to navigate through the created IDs and press **[ENT]** to select one of them. Note that a logical unit previously assigned to a Primary Controller can only be mapped a Primary ID, and vice versa.

```
Map Sec Ctlr
CH=0 ID= 000  ?
```

Use the up or down arrow keys to choose mapping "Logical Drive" or "Logical Volume" to host LUNs. If the logical unit has been partitioned, map each partition to different ID/LUNs.

```
Map to
Logical Drive ?
```

Use the up or down arrow keys to choose a LUN number and press **[ENT]** to confirm.

```
CH0 ID0 LUN0
No Mapped
```

Press **[ENT]** again to confirm.

```
Map Host LUN   ?
```

Use the up or down arrow keys to select a logical drive/logical volume if there are many.

```
LG0 RAID5 DRV=3
  9999MB GD SB=0
```

Press **[ENT]** and choose a partition if the logical unit has been partitioned.

```
LG=0 PART=0
     999MB   ?
```

Press **[ENT]** again to confirm
Press **[ENT]** to confirm the mapping.
Press **[ENT]** to re-ensure.

```
CH0 ID9 LUN0 Map
to LG0 PRT0?
```

This message indicates that the logical unit has been successfully mapped to the ID/LUN combination. Use the up or down arrow keys to continue mapping other logical units or press **[ENT]** to delete the mapped LUN.

```
Map Host LUN   ?
```

```
Map Sec. Ctlr
CH=0 ID= 0   ?
```

Repeat the process to map all the logical units to host ID/LUNs.

```
CH0 ID9 LUN0
Mapto LG0 PRT0
```

Front Panel View of Controller Failure

What will happen when one of the controllers fails?

If one of the controllers fails, the existing controller will automatically take over within a few seconds.

The red ATTEN LED will light up, and the message "Redundant Ctlr Failure Detected" will appear on the LCD. Users will be notified by audible alarm.

```
Redundant Ctlr
Failure Detected
```

NOTE:

- *Although the surviving controller will keep the system working, you should contact your system vendor for a replacement controller as soon as possible. Your vendor should be able to provide an appropriate replacement unit.*
- *Some operating systems (SCO, UnixWare, and OpenServer, for example) will not attempt to retry accessing the hard disk drives while controller is taking over.*

When and how is the failed controller replaced?

Remove the failed controller **after** the "working" controller has taken over. For the ventilation concern in most enclosures, it is better to leave a failed controller in place before a replacement arrives. For a controller with hot-plug capability, all you have to do is to remove the failed controller.

Redundant controller subsystems are designed to withstand a single controller failure. To restore a dual-active operation, do the following:

1. Combine the controllers while the surviving controller is working. If you power down the surviving controller and power up both controller together, problems might occur for they might contend for the role of the Primary controller.

2. (The replacement controller provided by your supplier should have been configured as the Secondary controller. It is recommended to safety check the status of the replacement controller before installing it to your redundant system. Simply attach power to the replacement and check its status." When the safety check is done, remove the failed controller and install the replacement controller into its place.)

When the replacement is connected, the "Auto-Failback" process will start automatically. If the replacement controller does not initialize, you may

execute the following steps to bring the new controller online. Press **[ENT]** for two seconds on the existing controller to enter the Main Menu.

Use the up or down arrow keys to select "View and Edit Peripheral Dev..," then press **[ENT]**.

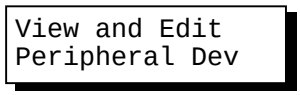

Choose "Set Peripheral Device Entry..," then press **[ENT]**.

Choose "Redundant Ctlr Function__," then press **[ENT]**.

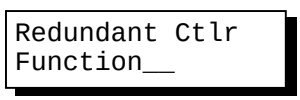

The message "Redundant Ctlr Primary/Secondary Degraded" will appear on the LCD.

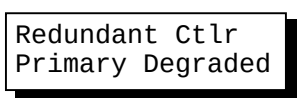

Press **[ENT]** and the message "Deassert Reset on Failed Ctlr?" will appear.

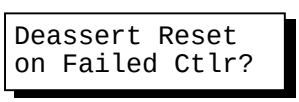

Press **[ENT]** for two seconds and the controller will start to scan for the new controller and bring it online.

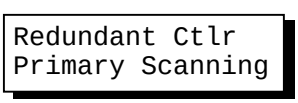

The new controller will then start to initialize.

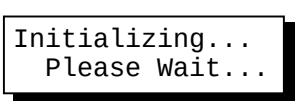

Once initialized, it will assume the role of the Secondary Controller.

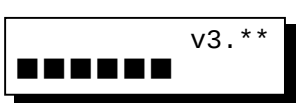

13.3.3 Via Terminal Emulation

Redundant Configuration Using Manual Setting

Power on Controller 1. Make sure Controller 2 is powered off.

Enter the Main Menu.

Use the arrow keys to navigate through the menus. Choose "View and Edit Peripheral Devices," then press **[ENTER]**.

Choose "Set Peripheral Devices Entry," then press **[ENTER]**. Choose "Redundant Controller [Function]," and then press **[ENTER]**. (**NOTE**: The current setting will be displayed on the screen. If this controller has never been set as a redundant controller before, the default setting is "Disabled." The message "Redundant Controller - Disabled" will be displayed on the screen. Press **[ENTER]** to proceed.)

The message "Enable Redundant Controller Function" will appear.

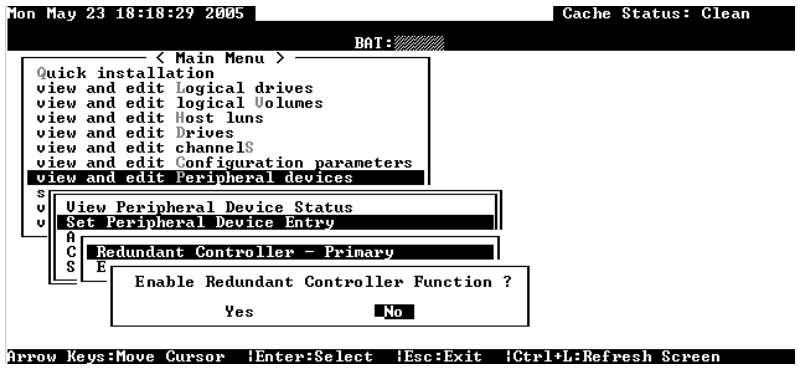

Use the arrow keys to scroll through the available options, then press **[ENTER]** to proceed. Choose **Yes** to continue**.**

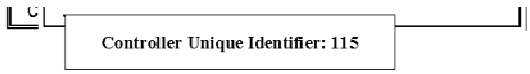

A "Controller Unique Identifier" box will appear. Enter a hex number from 0 to FFFFF, then press **[ENTER]** to proceed. The value you enter for controller unique ID should be different for each controller.

Power off Controller 1, and then power on Controller 2. Set Controller 2 to "Primary" as described in the steps mentioned above. Power off Controller 2.

When the redundant controller function on both of the partner controllers are set to the "Primary" setting, the controllers will decide between themselves which will be the Primary Controller. Primary is temporarily the only option.

Creating Primary and Secondary ID

Enter "View and Edit Channels." Press **[ENTER]** and select the host or drive channel on which you wish to create Primary/Secondary ID.

Drive Channel

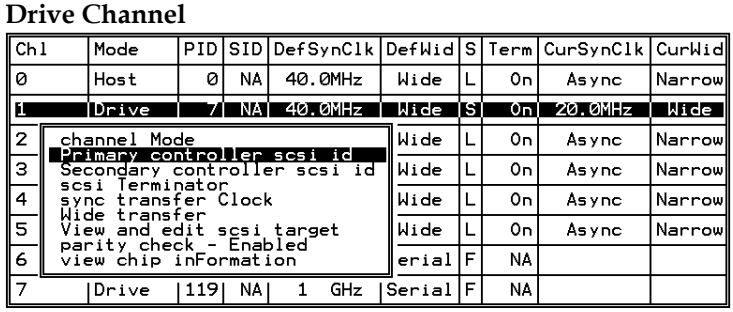

Host Channel

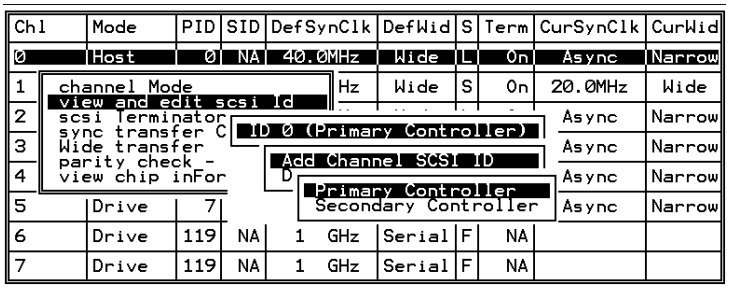

The configuration change will only take effect after the controller has been reset.

Assigning Logical Drives to the Secondary Controller

A logical drive can be assigned to the Primary or Secondary controller. By default, logical drives will be automatically assigned to the Primary controller. It can be assigned to the Secondary controller if the host computer is also connected to the Secondary controller. -

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. Move the cursor bar to "Logical Drive Assignments" and press **[ENTER]** if you want to assign logical drive to the Secondary controller.

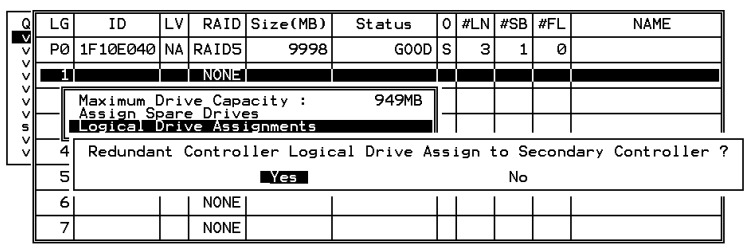

Logical drive assignment can also be changed after a logical drive is created. Create a logical drive or choose an existing logical drive, then press **[ENTER]** to see the logical drive menu. Choose "Logical Drive Assignments," then press **[ENTER]**. Choose **Yes** and press **[ENTER]** to confirm reassignment to the Secondary controller.

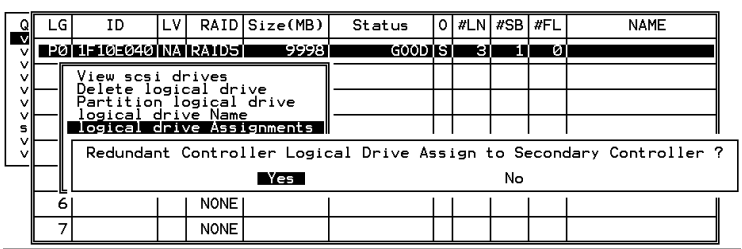

The reassignment is evident from the "View and Edit Logical Drives" screen. "S0" indicates that the logical drive is Logical Drive 0 assigned to the Secondary controller.

LG	ID	LV	RAID	Size(MB)	Status	0	#LN	#SB	#FL	NAME
S0	1F10E040	NA	RAID5	98	GOOD	S	3	1	0	
P1	4DB655C0	NA	RAID3	98	GOOD	S	3	0	0	
2			NONE							
3			NONE							
4			NONE							
5			NONE							
6			NONE							
7			NONE							

Mapping a Logical Drive/Logical Volume to the Host LUNs

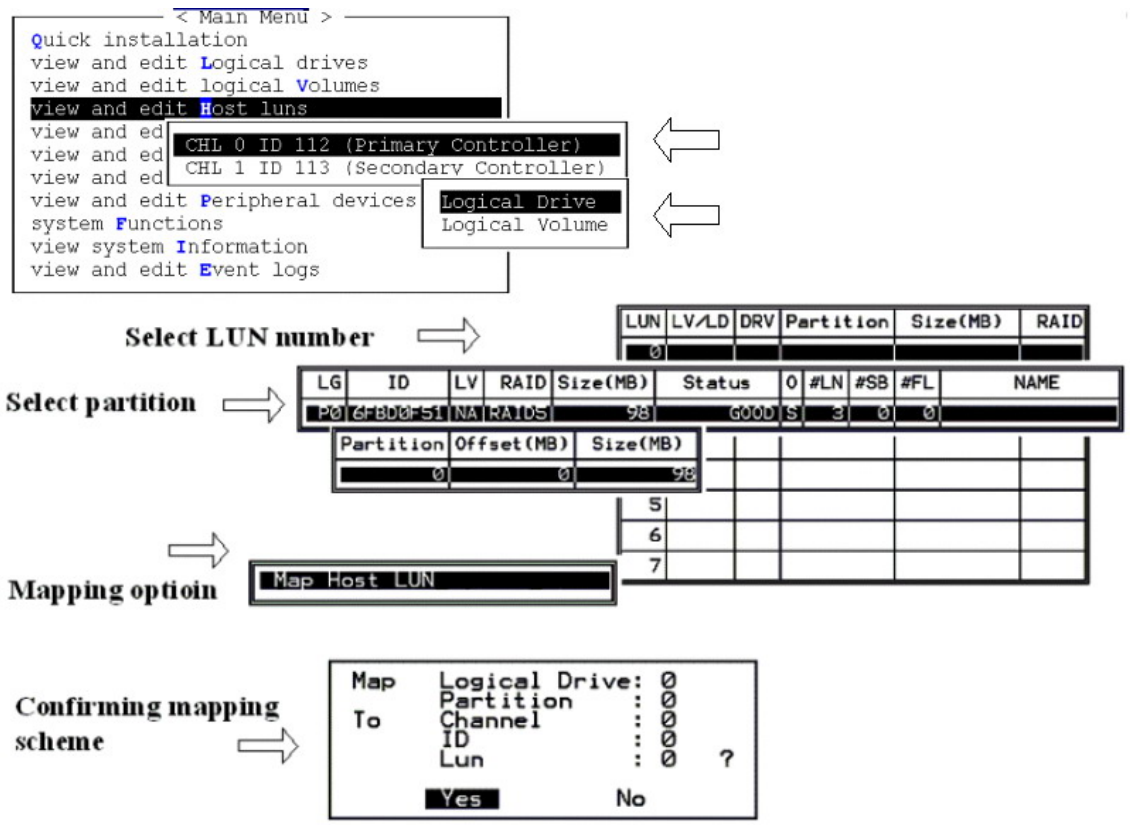

Terminal Interface View of Controller Failure

What will happen when one of the controllers fails?

When one of the controllers fails, the other controller will take over in a few seconds.

There will be an alert message that reads "Redundant Controller Failure Detected."

Users will be notified by audible alarm.

After a controller takes over, it will act as both controllers. If the Primary Controller failed, the Secondary Controller becomes the Primary Controller. If the failed controller is replaced by a new one later, the new controller will assume the role of the Secondary Controller.

NOTE:

- *Some operating systems (SCO, UnixWare, and OpenServer, for example) will not attempt to retry accessing the hard disk drives while the controller is taking over.*

13.3.4 When and How Is the Failed Controller Replaced?

Remove the failed controller **after** the take-over of the "working" controller has been completed. For a controller with hot-plug capability, all you have to do is to remove the failed controller. With ventilation considerations, it is recommended to leave a failed controller in most enclosures until the replacement controller arrives.

When the new controller is connected, the existing controller will automatically start initializing the replacement controller. If the existing controller does not initialize the replacement controller, execute the "Deassert Reset on Failed Controller" function.

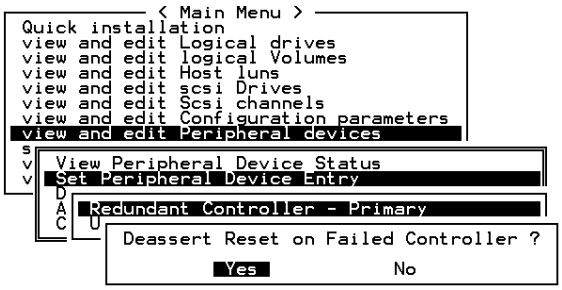

If the replacement has been initialized successfully, you may proceed to examine the system status. From the Main Menu, select "View and Edit Peripheral Devices" and then "View Peripheral Device Status" to see that the new controller is being scanned.

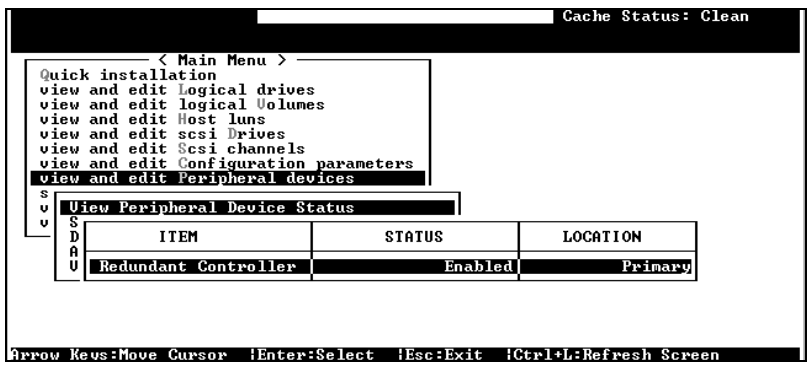

When the scanning is completed, the status will change to "Enabled."

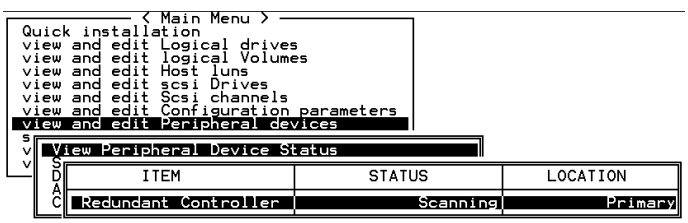

Forcing Controller Failover for Testing

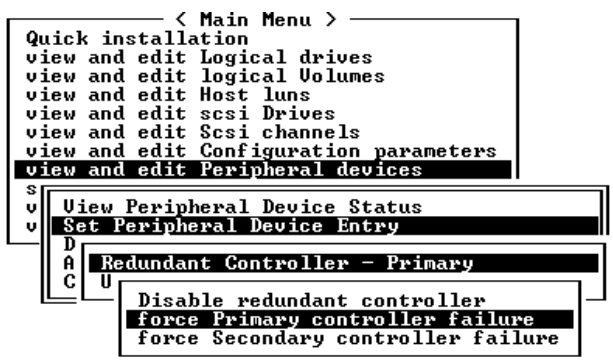

This function is reserved for de-bugging purposes.

Testing the failover functionality can be performed using the following methods.

1. Pulling out one of the controllers to simulate controller failure

Pull out either the Primary or the Secondary Controller An error message will display immediately and the alarm will sound. The existing controller takes over the workload within a second. Clear all errors by pressing the **ESC** key. You may re-install the removed controller after all activities have been taken over by the existing controller. It may take a while for the controllers to finish re-initialization and assuming their load.

2. "Forcing controller failure"

Select "View and Edit Peripheral Devices," "Set Peripheral Device Entry," and "Redundant Controller Primary/Secondary."

Select "Force Primary/ Secondary Controller Failure." You may now pull out the controller you had just disabled. I/Os should be continued by the existing controller. Continue the aforementioned procedure to complete the test.

WARNING!

- *This function should only be performed for testing the redundant controller functionality before any critical data is committed to the drives. Although the controller is designed to be hot-swappable, unpredictable failures may occur during the process, e.g., improper handling of the module or improper handling of PCB boards while replacing the controller.*

RCC Status (Redundant Controller Communications Channel)

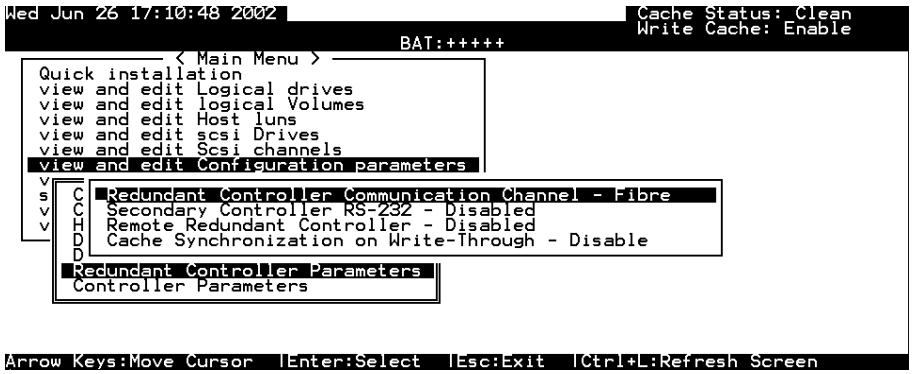

This item is for display only, showing the current communications route.

Secondary Controller RS-232

This is an option reserved for debugging purposes. When enabled, you can access the Secondary Controller through its serial port. In a redundant controller system, only status display is available through the terminal session with a Secondary Controller. No configuration change can be done through a Secondary Controller.

Remote Redundant Controller

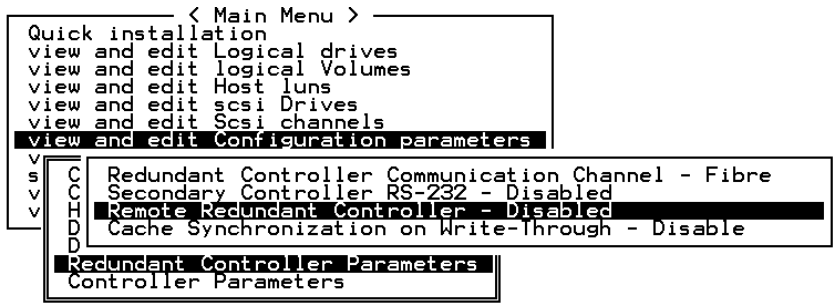

This function enables two partner controllers to be connected by FC links over an extended distance, e.g., between two campus buildings. This is an advanced option reserved for system integrators. For more details, please contact Infortrend Technical Support.

Cache Synchronization on Write-Through

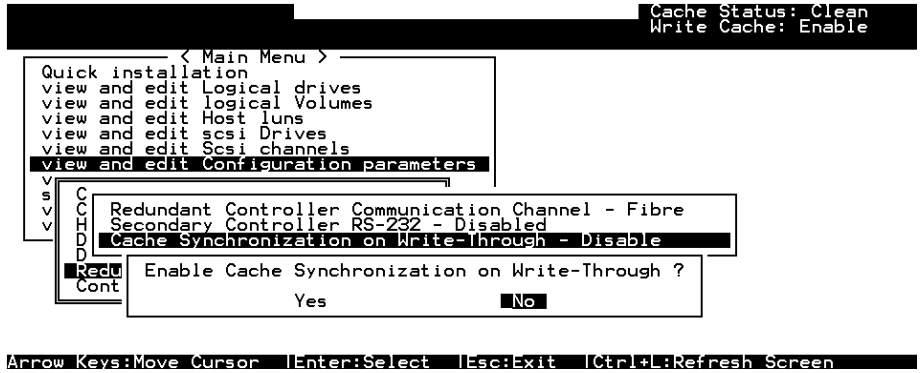

If your redundant controller system is not operating with Write-back caching, you can disable synchronized cache communications. Your system can be spared the effort to mirror and transfer data between partner controllers. This increases array performance for subsystems that operate without write-caching.

Chapter

14 Record of Settings

In addition to saving the configuration data in NVRAM to disk, keeping a hard copy of the controller configuration is also recommended. This will speed the recreation of the RAID in the event of a disaster.

The following tables are provided as a model for recording the configuration data.

As a general rule, the configuration data in the NVRAM should be saved to disk or as a file (using RAIDWatch Manager) whenever a configuration change is made.

14.1 View and Edit Logical Drives

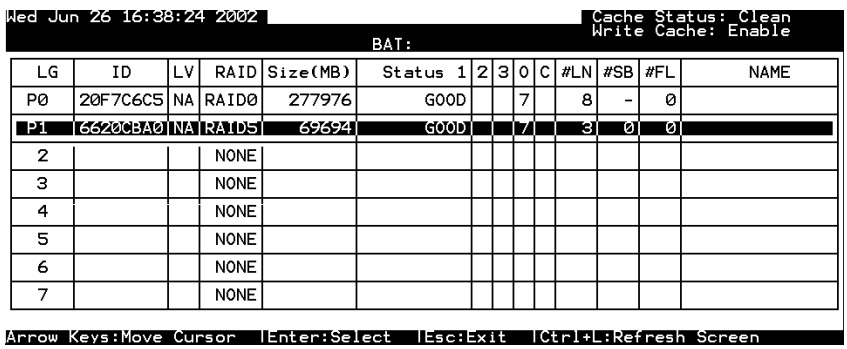

Logical Drive Information

LG	ID	LV	RAID Level	Size (MB)	Status 1	2	3	O	C

#LN	#SB	#FL	NAME

Partition Information

LG	Partition	Size (MB)	LG	Partition	Size (MB)

14.2 View and Edit Logical Volumes

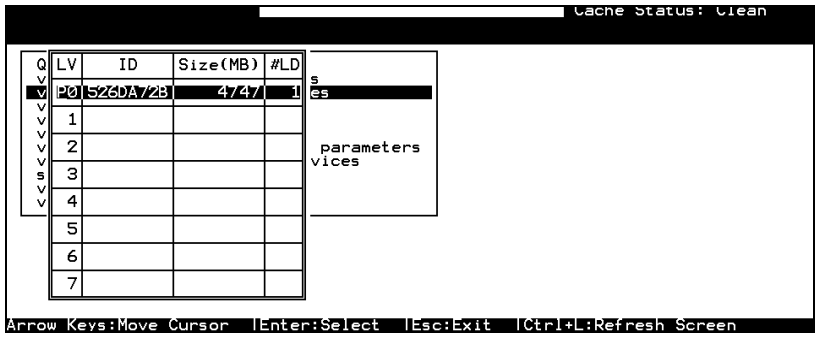

Logical Volume Information

LV	ID	Size	#LD	Stripe Size

Partition Information

LV	Partition	Size (MB)	LV	Partition	Size (MB)

14.3 View and Edit Host LUN's

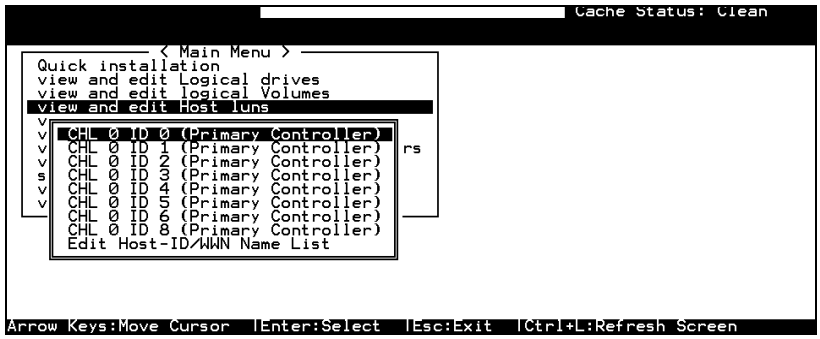

LUN Mappings

Host Channel	Pri. / Sec. Controller	SCSI ID	LUN	Logical Drive / Logical Volume	Partition	Size

Host-ID/WWN Name List

Host-ID/WWN	Name List

Access Restriction Setting

Logical Drive / Logical Volume	Partition	HBA WWN List

14.4 View and Edit Drives

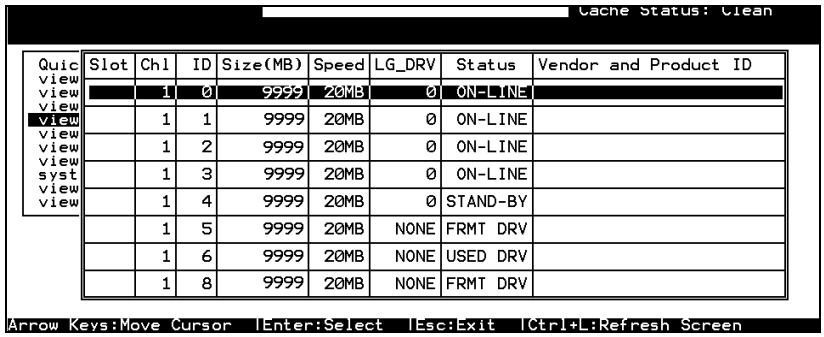

Slot	Chl	ID	Size (MB)	Speed	LG DRV? Global Spare? Local Spare?		Vendor & Product ID
					LG		
					LG		
					LG		
					LG		
					LG		
					LG		
					LG		
					LG		
					LG		
					LG		
					LG		
					LG		
					LG		
					LG		

14.5 View and Edit Channels

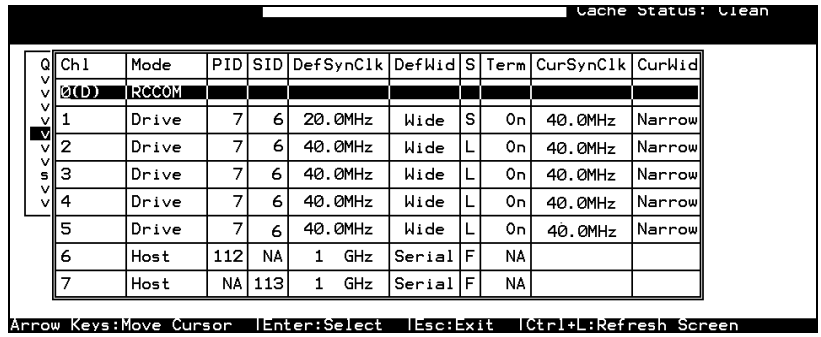

Chl	Mode (Host / Drive)	Primary Controller SCSI ID(s)	Secondary Controller SCSI ID(s)	Default Sync Clock	Default Wide	Terminator Diff/Enable/ Disable/	Current Sync Clock	Current Width

Parity Check	View Channel Host-ID/WWN	View Device Port Name List (WWPN)

14.6 View and Edit Configuration Parameters

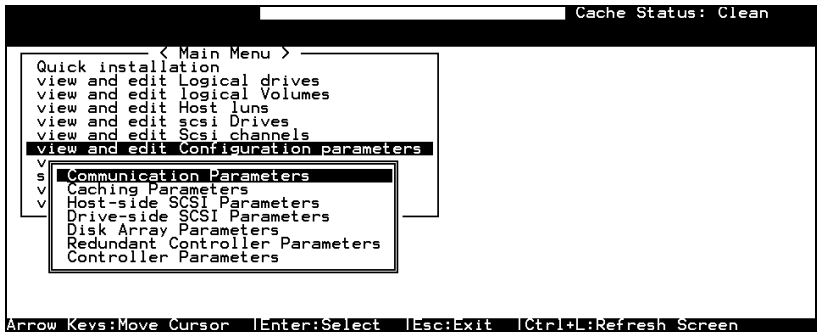

Communication Parameters

RS-232 Port Configuration

COM 1 (RS-232 Port)

Baud Rate	☐2400 ☐4800 ☐9600 ☐19200 ☐38400
Terminal Emulation	☐Enabled ☐Disabled

COM 2 (Redundant Controller Port)

Baud Rate	☐ 2400 ☐4800 ☐9600 ☐ 19200 ☐38400
Data Routing	☐Direct to Port ☐Through Network
Terminal Emulation	☐Enabled ☐Disabled

Ethernet Configuration

IP address	_____
NetMask	_____
Gateway	_____

Caching Parameters

Write-back Cache	☐Enabled ☐Disabled
Periodic Cache Flush Time	☐Enabled ☐Disabled
Optimization for	☐Random I/O ☐Sequential I/O

Host Side Parameters

Maximum Queued I/O Count	☐ Auto _____
LUNs per Host SCSI ID	☐ LUNs
Number of Tags Reserved for each Host-LUN connection	_____
Peripheral Device Type Parameters	Peripheral Device Type - Device Qualifier - Removable Media -

	LUN Applicability -
Host Cylinder/Head/Sector Mapping configuration	Cylinder - Head - Sector -
Max Number of Concurrent Host-LUN Connection	_____
Fibre Connection Options	_____

Drive Side Parameters

SCSI Motor Spin-up (SCSI drives only)	☐Enabled ☐Disabled
SCSI Reset at Power Up (SCSI drives only)	☐Enabled ☐Disabled
Disk Access Delay Time	☐No Delay _____ Seconds
SCSI I/O Timeout	☐Default _____
Maximum Tag Count	☐Disabled _____
Periodic Drive Check Time	☐Disabled _____
Periodic SAF-TE and SES Device Check Time	☐Disabled _____
Periodic Auto-Detect Failure Drive Swap Check Time	☐Disabled _____
Drive Predictable Failure Mode	☐Disabled ☐Detect Only ☐Detect and Perpetual Clone ☐Detect and Clone + Replace
Auto-assign Global Spare Drive	☐Enabled ☐Disabled
Fibre Channel Dual Loop	☐Enabled ☐Disabled

Disk Array Parameters

Rebuild Priority	☐Low ☐Normal ☐Improved ☐High
Verifications on Writes	
Verifications on LD Initialization Writes	☐Enabled ☐Disabled
Verifications on LD Rebuild Writes	☐Enabled ☐Disabled
Verifications on Normal Drive Writes	☐Enabled ☐Disabled
Max. Drive Response Timeout	☐Disabled _____

Redundant Controller Parameters

Secondary Controller RS-232	☐Enabled ☐Disabled
Cache Synchronization on Write-through	☐Enabled ☐Disabled

Controller Parameters

Controller Name	☐Not Set _____

LCD Tile Display	☐Controller Logo ☐Controller Name
Password Validation Timeout	☐Disabled ☐1 minute ☐2 minutes ☐5 minutes ☐Always Check
Controller Unique Identifier	_____
SDRAM ECC	☐Enabled ☐Disabled
Set Controller Date and Time	_____

14.7 View and Edit Peripheral Devices

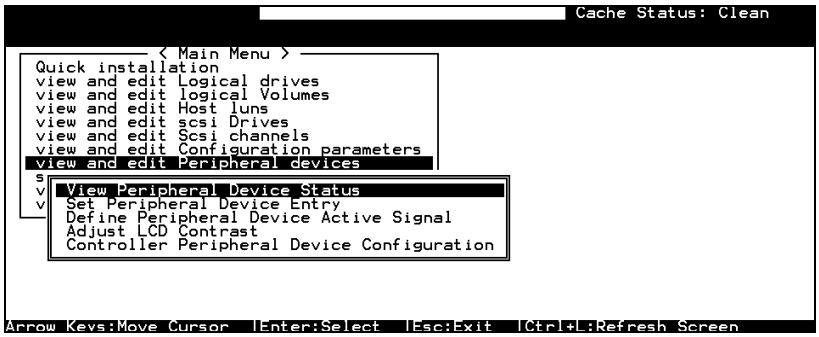

Set Peripheral Device Entry

Redundant Controller	☐ Enabled ☐ Disabled
Power Supply Status	☐ Enabled ☐ Disabled
Fan Status	☐ Enabled ☐ Disabled
Temperature Status	☐ Enabled ☐ Disabled
UPS Status	☐ Enabled ☐ Disabled

Event Trigger Options

Controller Failure (dual-active configurations)	☐ Enabled ☐ Disabled
BBU Low/Failed	☐ Enabled ☐ Disabled
UPS AC Power Loss	☐ Enabled ☐ Disabled
Power Supply Failed	☐ Enabled ☐ Disabled
Fan Failed	☐ Enabled ☐ Disabled
Temperature Exceeds Limits	☐ Enabled ☐ Disabled

Define Peripheral Device Active Signal

Power Supply Fail Signal	☐ Active High ☐ Active Low
Fan Fail Signal	☐ Active High ☐ Active Low
Temperature Alert Signal	☐ Active High ☐ Active Low
UPS Power Fail Signal	☐ Active High ☐ Active Low
Drive Failure Outputs	☐ Active High ☐ Active Low

View System Information

CPU Type	
Total Cache Size	☐ SDRAM _____ MB
Firmware Version	
Bootrecord Version	

FW Upgradability	
Serial Number	
Battery Backup	☐ On ☐ Off
Base Board Rev. ID	
Base Board ID	
ID of NVRAM Defaults	
Controller Position	

Event Threshold Parameters

Thresholds for +3.3V	Upper _____ Lower _____
Thresholds for +5V	Upper _____ Lower _____
Thresholds for +12V	Upper _____ Lower _____
Thresholds for CPU temperature	Upper _____ Lower _____
Thresholds for Board Temperature	Upper _____ Lower _____

14.8 Save NVRAM to Disk, Restore from Disk

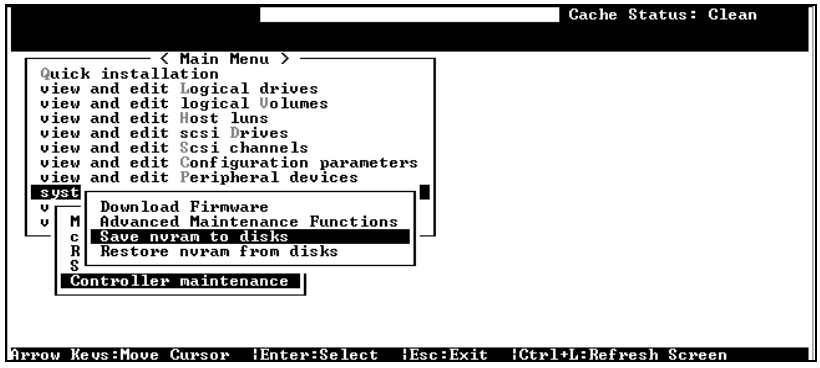

Update Firmware	Date	Save NVRAM to Disk or File	Date/Location	Restore NVRAM from Disk	Date

14.9 RAID Security: Password

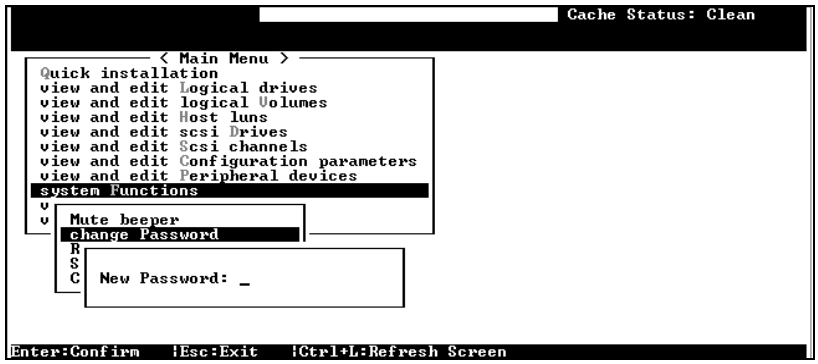

RAID Security

Controller Name	Password _____

Chapter

15 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.

This chapter is organized as follows:

15.1 Overview: Notes on using the expansion functions

15.2 Mode 1 Expansion: Theory and configuration procedure: expansion by adding drives

15.3 Mode 2 Expansion: Theory and configuration procedure for expansion by copying and replacing drives

15.4 Making Use of the Added Capacity: Expand Logical Drive: Configuration procedure of the Expand function for a logical drive

15.5 Expand Logical Volume: Configuration procedure for the Expand function for a logical volume

15.6 Configuration Example: Volume Extension in Windows 2000®

15.1 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.

Notes on Expansion

1. Added 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).

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, and 5 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 15 - 1 Logical Drive Expansion

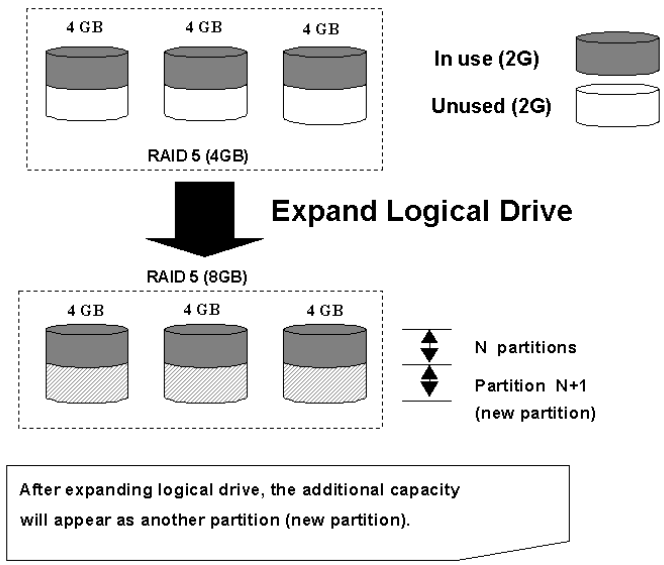

RAID levels supported: RAID 0, 1, 3, and 5
Expansion can be performed on logical drives or logical volumes under the following conditions:

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.

15.2 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 15 - 2 Expansion by Adding Drive

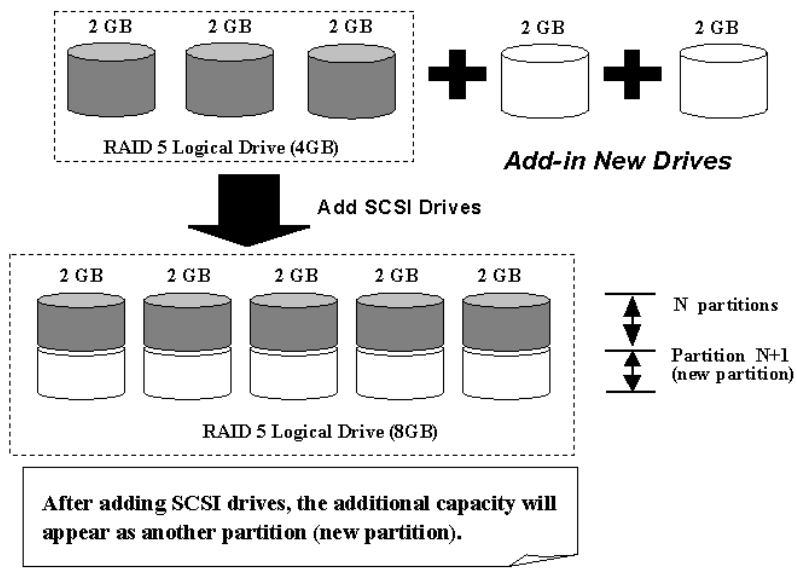

RAID levels supported: RAID 0, 1, 3, and 5.

The new partition must be made available through a host ID/LUN.

Add Drive Procedure

First select from the Main Menu, "View and Edit Logical Drive," and select a logical drive to add a new drive to. 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.

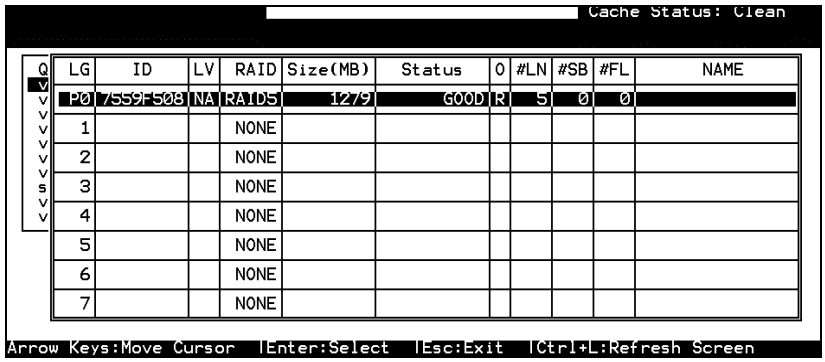

Press **[ENTER]** to select a logical drive and choose "Add Drives" from the submenu. Proceed with confirming the selection.

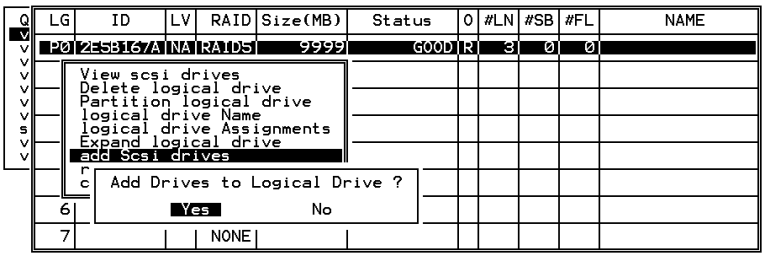

Available drives will be listed. Select one or more drive(s) to add to the target logical drive by pressing **[ENTER].** The selected drive will be indicated by an asterisk "*" mark.

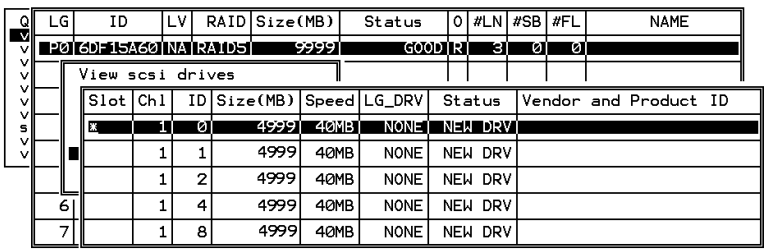

Press **[ESC]** to proceed and the notification will prompt.

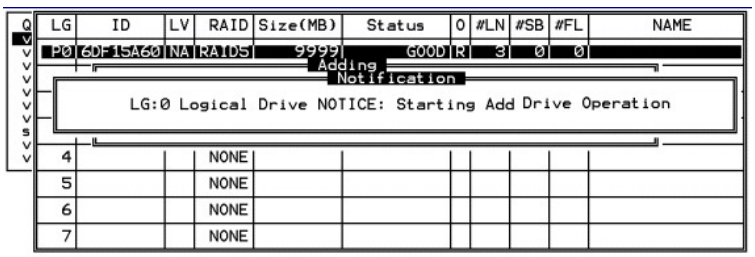

Press **[ESC]** again to cancel the notification prompt; a status bar will indicate the percentage of progress.

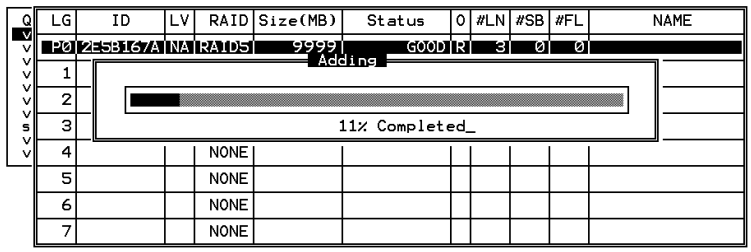

Upon completion, there will appear a confirming notification. The capacity of the added drive will appear as an unused partition.

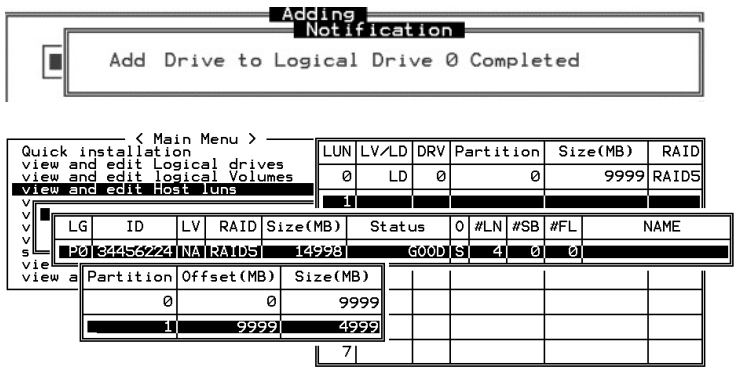

The added capacity will be automatically included, meaning that you do not have to "expand logical drive" later. Map the added capacity to another host ID/LUN to make use of it.

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.*

15.3 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 added capacity.

Figure 15 - 3 Expansion by Copy & Replace

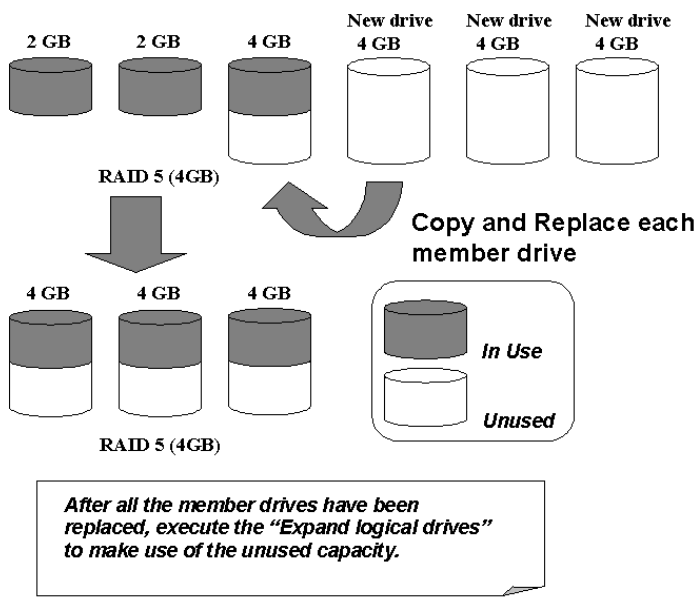

RAID levels supported: RAID 0, 1, 3, and 5

Copy and Replace Procedure

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.

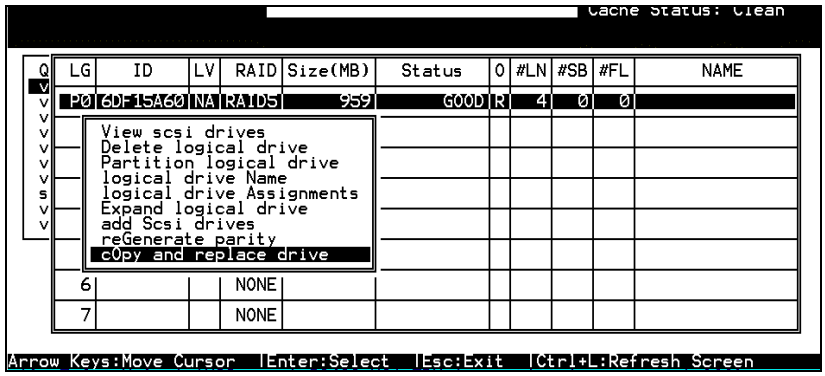

The array members will be listed. Select the member drive (the source drive) you want to replace with a larger one.

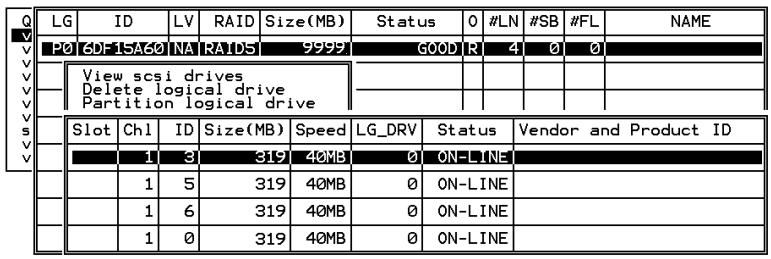

Select one of the members as the "**source drive**" (status indicated as ON-LINE) by pressing **[ENTER]**; a table of available drives will prompt. Select a "**new drive**" to copy the capacity of the source drive onto. The channel number and ID number of both the "Source Drive" and the "Destination Drive" will be indicated in the confirming box.

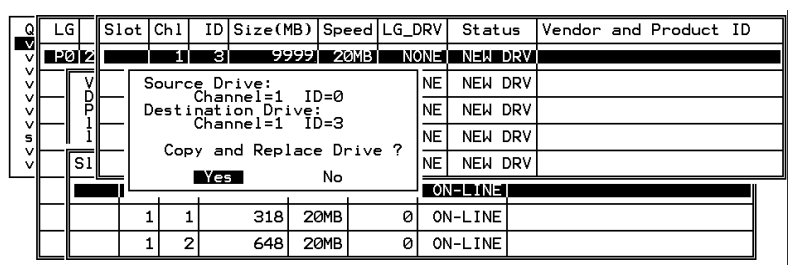

Choose **Yes** to confirm and proceed.

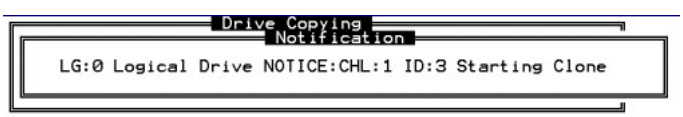

Press **[ESC]** to view the progress.

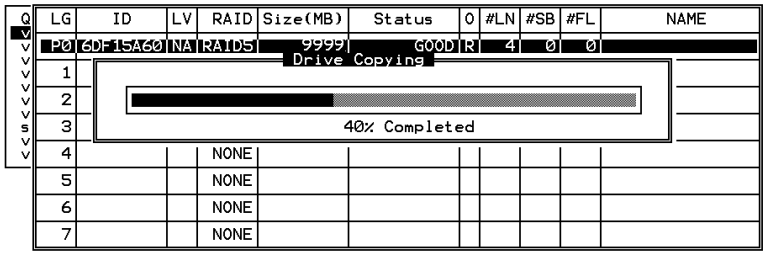

Completion of the Copy and Replace process will be indicated by a notification message. Follow the same method to copy and replace every member drive. You may now perform "Expand Logical Drive" to make use of the added capacity, and then map the additional capacity to a host LUN.

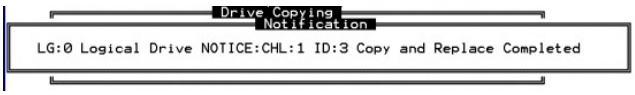

15.4 Making Use of the Added Capacity: Expand Logical Drive

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.

1. Select "View and Edit Logical Drives" from the Main Menu and select the logical drive with its members copied and replaced.

2. Select "Expand Logical Drive" in the sub-menu and press **[ENTER]** to proceed. A confirming box will appear.

3. Proceed by pressing **[ENTER]** or entering any value no larger than the "maximum drive expand capacity" and press **[ENTER]**.

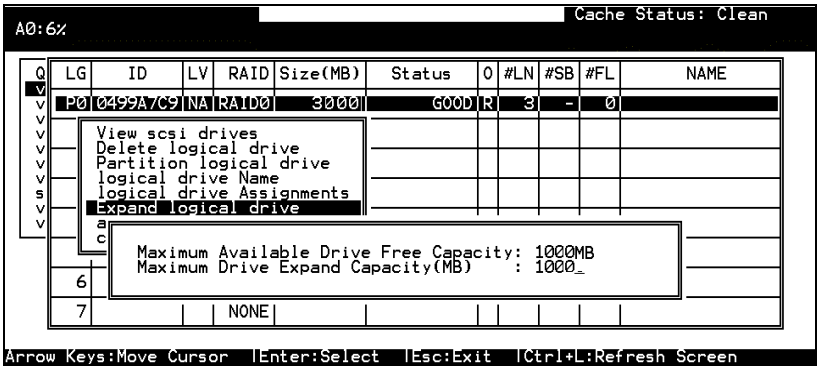

Choose **Yes** to confirm and proceed.

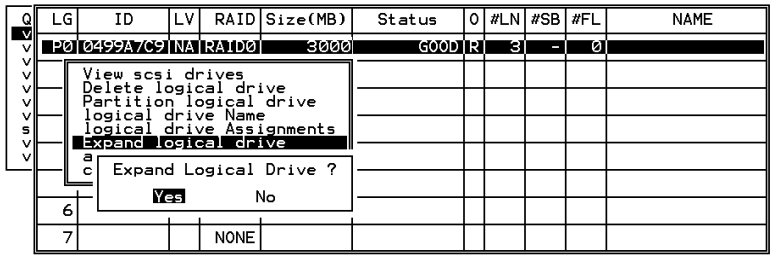

Upon completion, you will be prompted by the notification message.

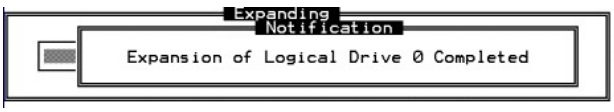

Press **[ESC]** to return to the previous menu screen.

The total capacity of logical drive has been expanded to 6GB.

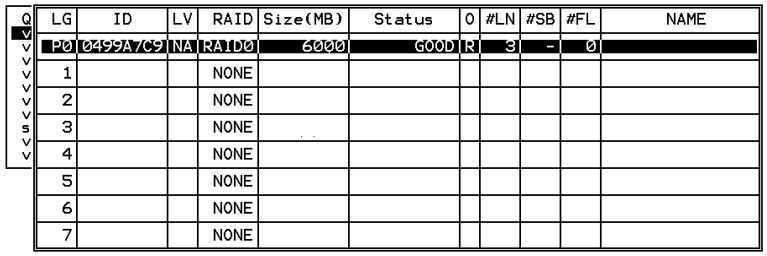

15.5 Expand Logical Volume

To expand a logical volume, expand its logical drive member(s) and then perform "expand logical volume."

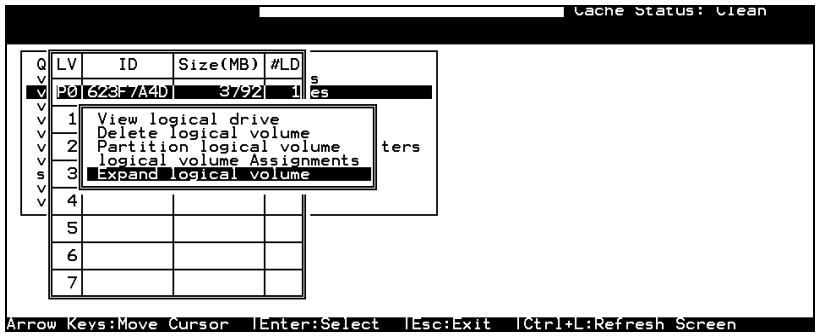

When prompted by "Expand Logical Volume?", choose **Yes** to confirm and the process will be completed immediately.

15.6 Configuration Example: Volume Extension in Windows 2000®

Limitations When Using Windows 2000

1. This 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 existing files 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.

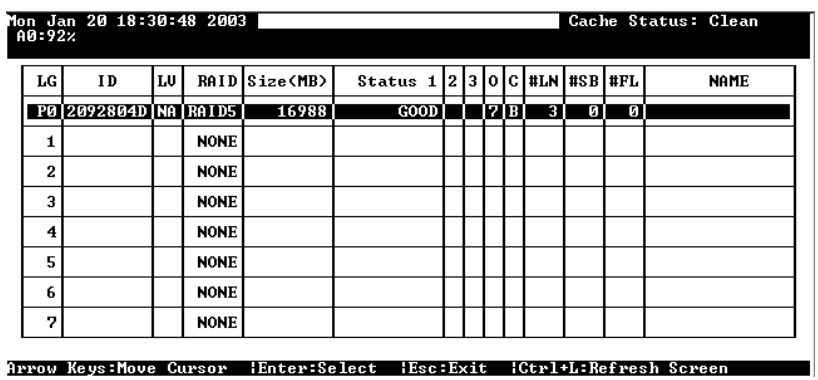

You can view information about this drive in the Windows 2000 Server's Computer Management -> Storage -> Disk Management.

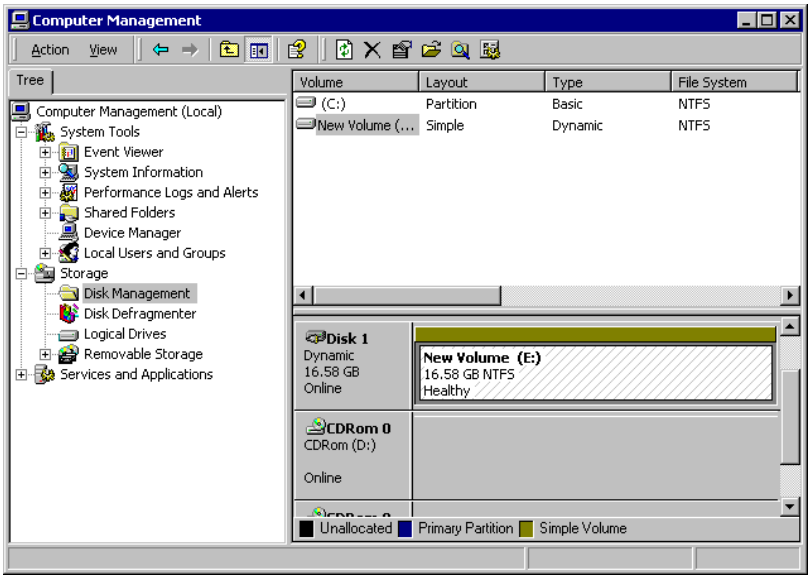

Place the cursor on Disk 1, right-click your mouse, and select "Properties." You will see that the total capacity for the Drive E: is about 16.5GB.

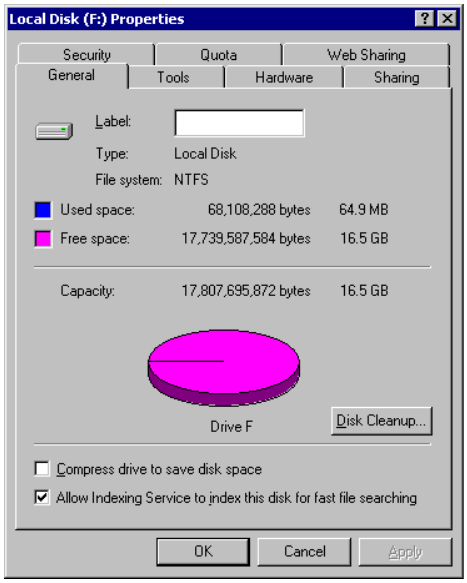

Follow the steps described in the previous section to "add" or "copy & replace" SCSI disk drives and perform "Logical Drive Expansion."

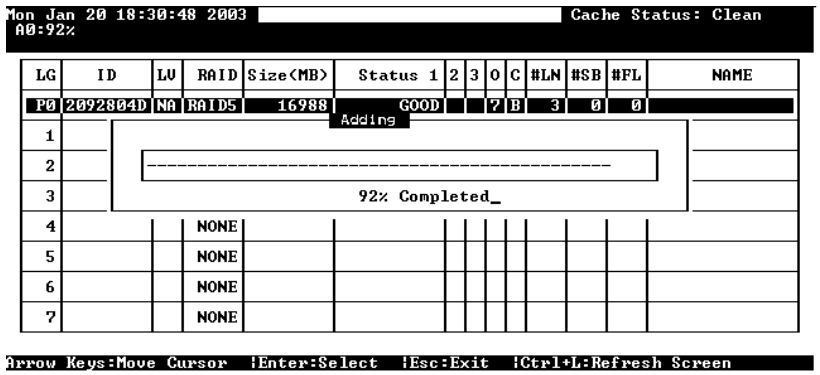

The 16.5GB logical drive has become a 25GB logical drive. Place the cursor on that logical drive, and then press **[ENTER]**.

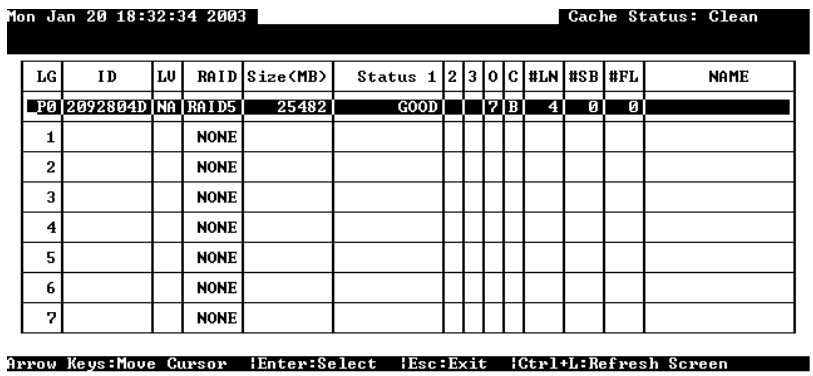

From the menu, select "Partition Logical Drive." You will see that the 25GB logical drive is composed of a 17GB partition and an 8.4GB partition.

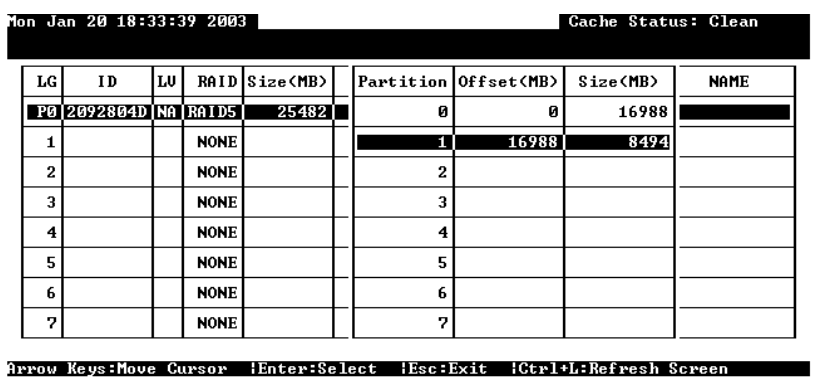

Follow the directions in ***Chapter 5*** and ***Chapter 7*** to map the new partition to a host LUN. The new partition must be "**mapped**" to a host LUN in order for the HBA (host-bus adapter) to see it. Once you have mapped the partition, reboot your Windows server. The HBA should be able to detect 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.

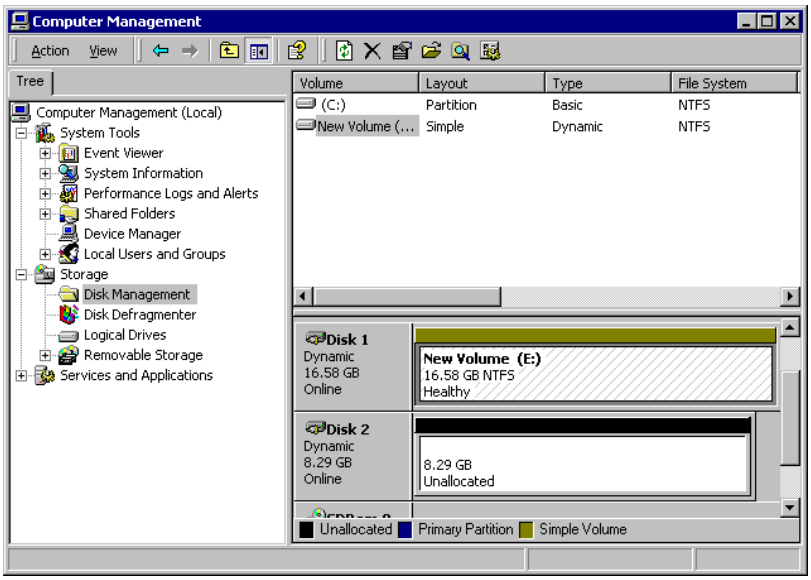

Select an existing volume (Disk1) and then right-click on the disk column. Select "Extend Volume" to proceed.

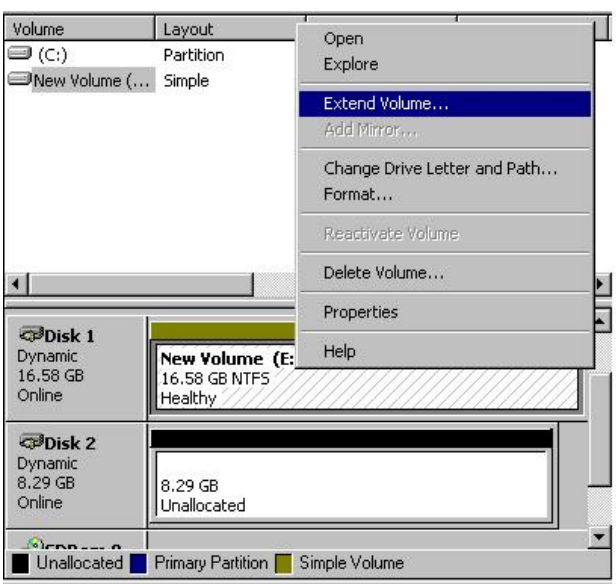

The Extend Volume Wizard will guide you through the rest of the process.

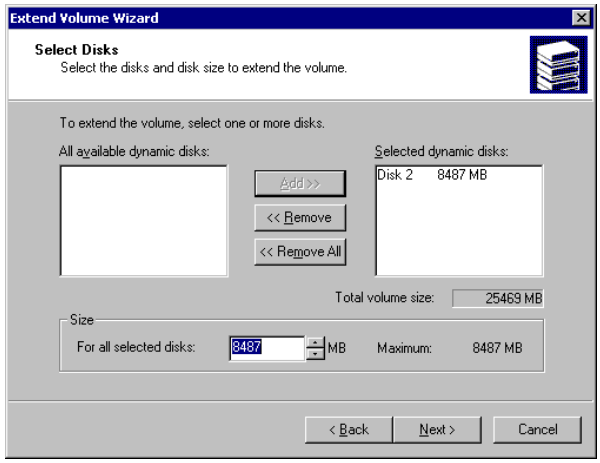

The screen will display that the volume set of Drive E: has been extended into a spanned volume by the 8.3GB in Disk2.

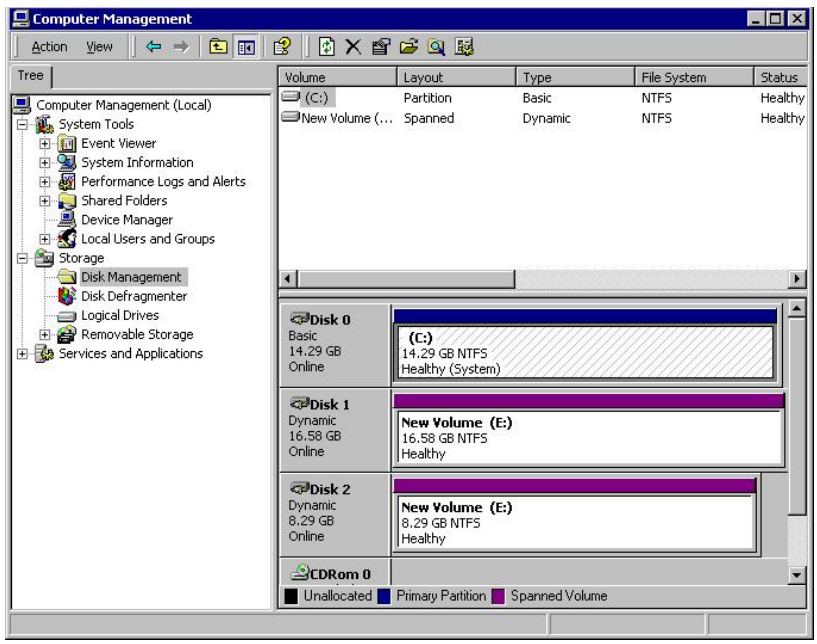

Logical Drive E: is now composed of two partitions with a total volume of 2500MB. To see this, hold down on the <Ctrl> key and select both Disk 1 and Disk2; then right-click your mouse and select "Properties."

Drive E: now has a capacity of about 25GB.

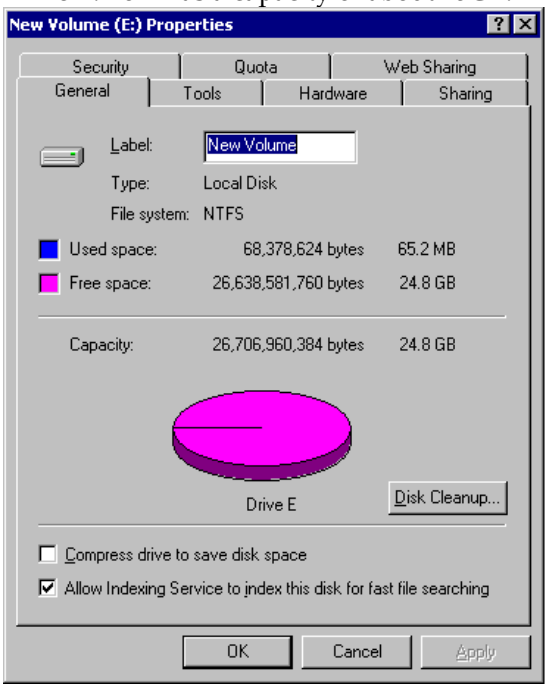

Chapter

16 S.M.A.R.T. Configuration

16.1 Overview

S.M.A.R.T.

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.

16.1.1 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."

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.

Locate the logical drive to which the drive with predictable errors belongs. Select the "Clone Failing Drive" function.

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.

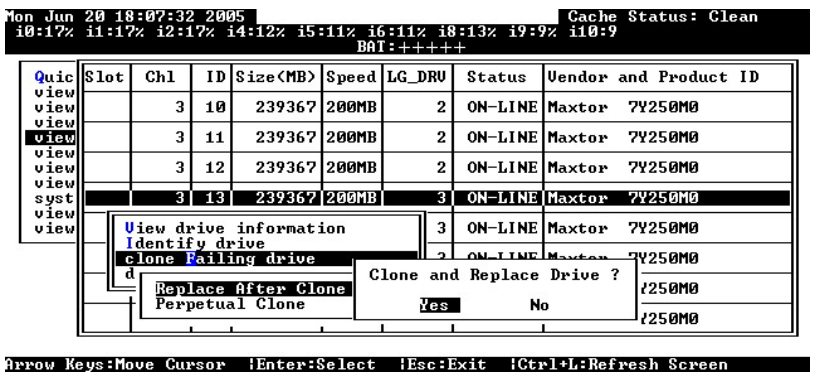

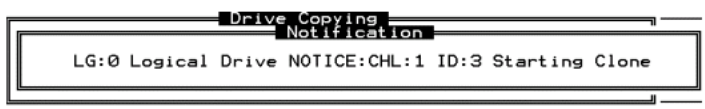

The cloning process will begin with a notification message. Press **[ESC]** to proceed.

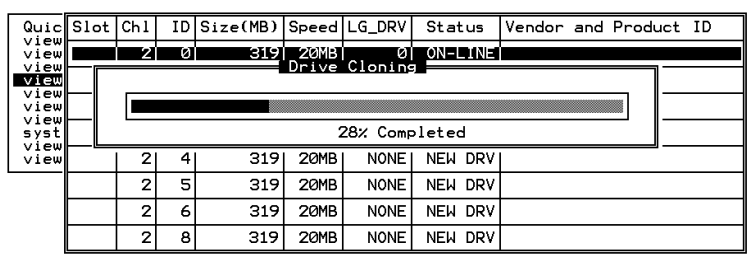

The cloning process will be indicated by a status bar.

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]**.

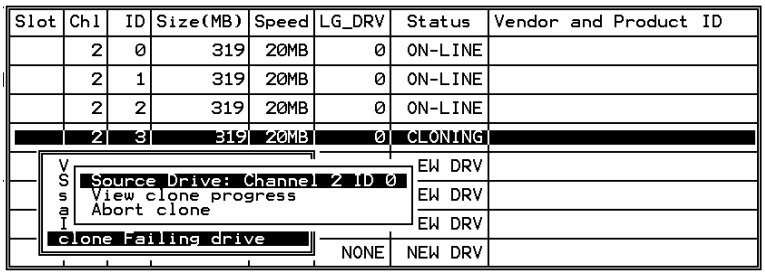

Select "Clone Failing Drive" again to view the current status. You may identify the source drive and choose to "View Clone Progress," or "Abort Clone" if you happen to have selected the wrong drive.

When the process is completed, you will be notified by the following message.

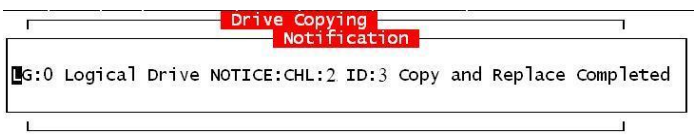

Perpetual Clone:

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.

In "View and Edit Drives," locate the member drive that shows predicted errors. Select "Clone Failing Drive," and choose "Perpetual Clone."

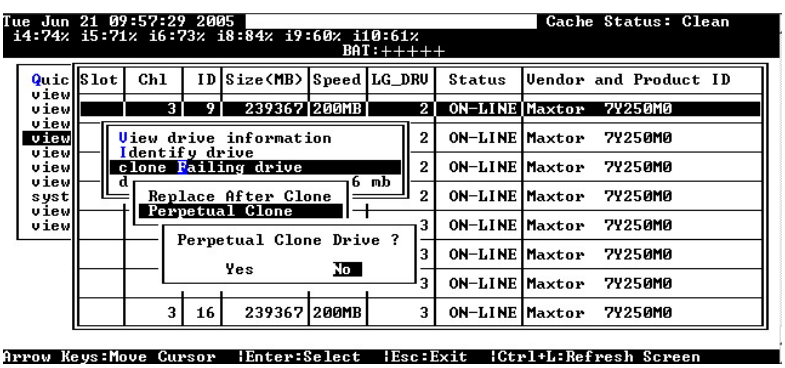

The controller will automatically start the cloning process using the existing "stand-by" (local/global spare drive) to clone the source drive (the target member drive).

The cloning process will begin with a notification message:

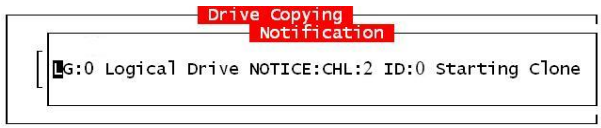

Press **[ESC]** to view the current progress:

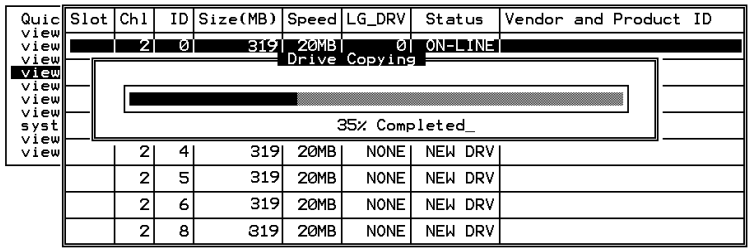

You may quit viewing the status bar by pressing **[ESC]** to return to the previous menu. Select the drive indicated as "CLONING" by pressing **[ENTER]**. Select "Clone Failing Drive" again to view the progress. You may identify the source drive and choose to "View Clone Progress" or "Abort Clone" if you happen to have selected the wrong drive.

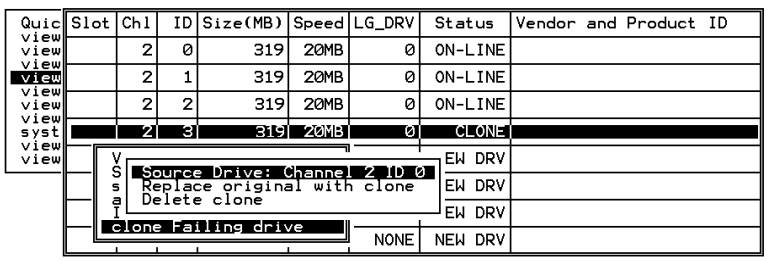

The cloning progress will be completed by a notification message as displayed below:

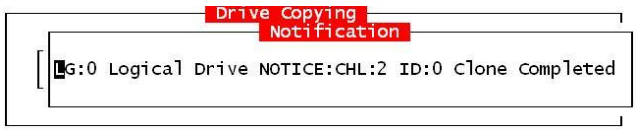

You may press **[ESC]** to clear the notification message to see the drives' status after the cloning process. The source drive (Channel 1 ID 5) remains as a member of logical drive "0," and the "stand-by" drive (Channel 1 ID 2, the dedicated/global spare drive) has become a "CLONE" drive.

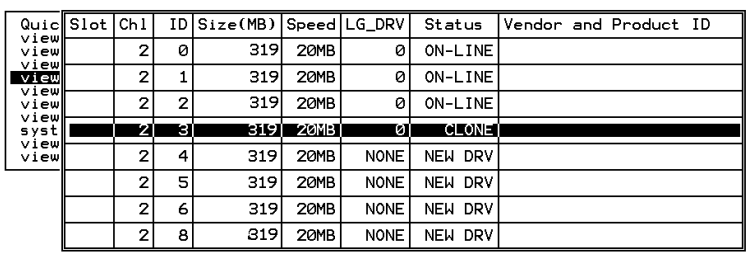

16.1.2 S.M.A.R.T. (Self-Monitoring, Analysis and Reporting Technology)

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 drive 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.

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.*

16.2 Configuration Procedure

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.

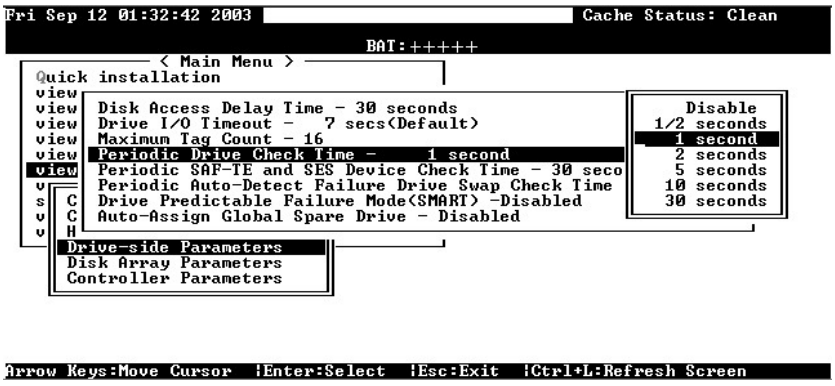

Step 2. In \View and Edit Configuration Parameters\Drive-side Parameters\Drive Predictable Failure Mode <SMART>, choose one from "Detect Only," "Detect, Perpetual Clone" and "Detect, Clone+Replace."

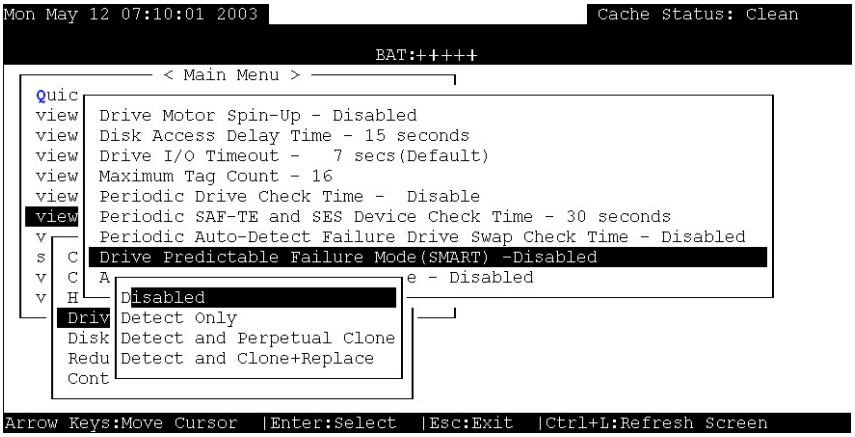

Examining Whether Your Drives Support S.M.A.R.T.

To see if your drive supports S.M.A.R.T., follow the steps below:

Step 1. Enable "S.M.A.R.T." for your drives in the RAID system.

Step 2. In "View and Edit Drives," choose one drive to test to. Press **[ENTER]** on the drive; a sub-menu will appear.

Step 3. Note that a new item "Predictable Failure Test" appears in the sub-menu. If the S.M.A.R.T." feature is not properly enabled, this item will not appear in the sub-menu.

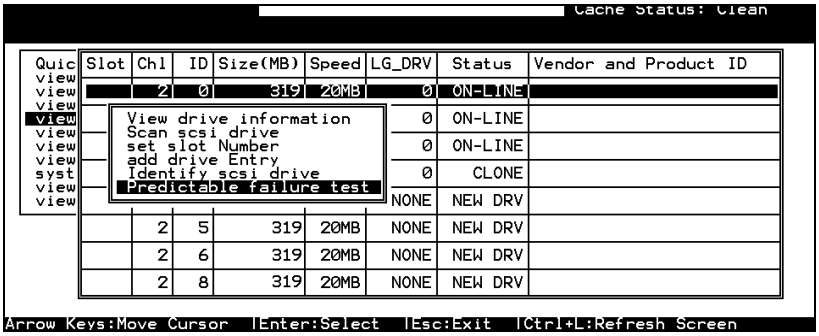

Step 4. Choose "Predictable Failure Test.," The controller will force the drive to simulate predictable drive errors.

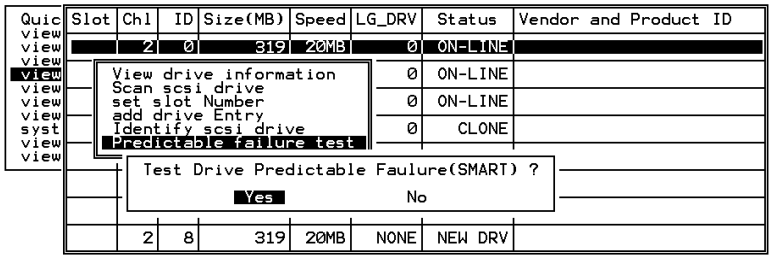

Step 5. Press **[ENTER],** and after a while (the next time the controller performs "Periodic Drive Check"), the controller will detect the errors simulated by the drive. An error message displays like this: "SMART-CH:? ID:? Predictable Failure Detected (TEST)." If this error message appears, it means your drive supports S.M.A.R.T. features.

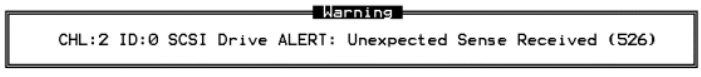

Step 6. If the error message does not appear, you may simply refer to related documentation or contact your drive manufacturer for information about whether the drive model and drive firmware version support S.M.A.R.T.

Using S.M.A.R.T. Functions

Step 1. Enable "S.M.A.R.T." on the RAID controller.

Step 2. Make sure your drives do support S.M.A.R.T. so that your system will work properly.

Step 3. **The "Detect Only" Setting**

3a. In \View and Edit Configuration Parameters\Drive-side Parameters\Drive Predictable Failure Mode <SMART>, choose "Detect Only."

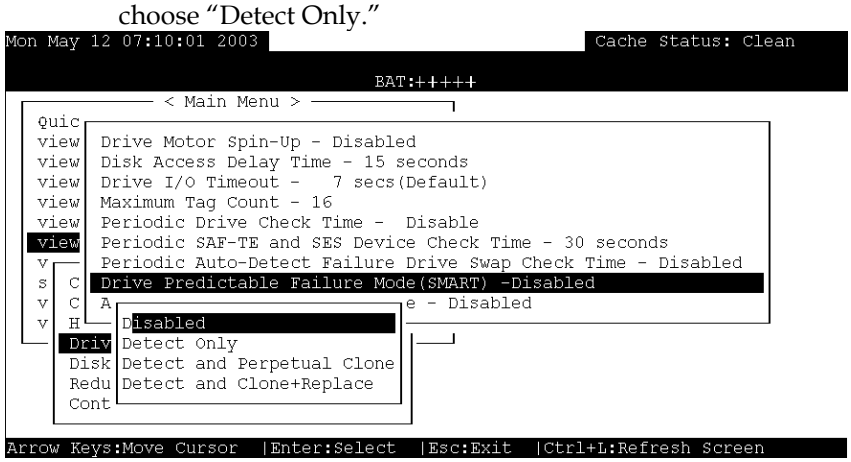

3b. Whenever a drive detects symptoms of predictable drive failure, the controller will issue an error message.

Step 4. **The "Detect, Perpetual Clone" Setting**

4a. In \View and Edit Configuration Parameters\Drive-side Parameters\Drive Predictable Failure Mode <SMART>, choose "Detect, Perpetual Clone."

4b. You should have at least one spare drive for the logical drive (either Local Spare or Global Spare Drive).

4c. 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.

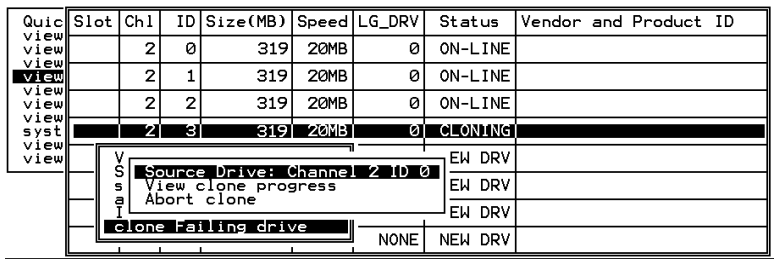

NOTE:

- *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.*

4d. 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 5. The "Detect, Clone + Replace" Function

5a. In \View and Edit Configuration Parameters\Drive-side Parameters\Drive Predictable Failure Mode <SMART>, choose "Detect, Clone+Replace."

5b. Make sure you have at least one spare drive to the logical drive. (Either Local Spare Drive or Global Spare Drive)

5c. 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."

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.

Chapter

17

Implementations for AV Applications

Maximum Drive Response Time

This option prepares the array for applications featuring “no drop-frame” operations and continuous reading such as video streaming. More options will be available for AV applications with future firmware releases.

In situations such as drive failure or the occurrence of bad blocks, a read returned after several hundreds milliseconds would prove too long and cause choppy audio or dropped video frames.

The maximum response time option, as displayed below, provides a timeout value for processing read requests. If timeout is reported on a specific member of an array, the subsystem immediately retrieves data from the parity data and other members of the array. In this way, causes of delay can be eliminated.

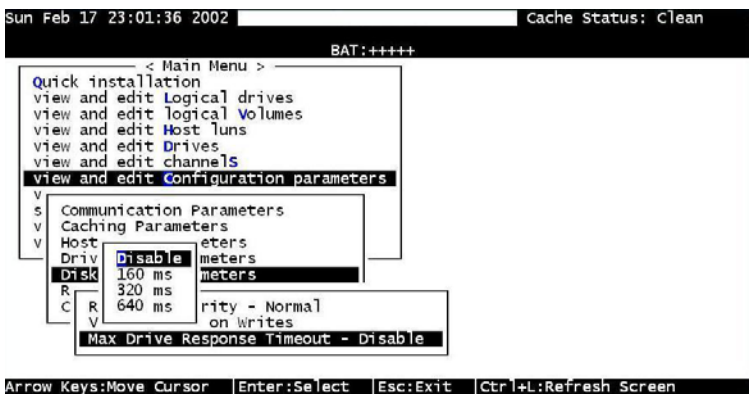

To prepare the array for read-intensive applications, the following are recommended:

1. Performance using the write-through caching mode is better than that using the write-back mode on subsystems equipped with redundant RAID controllers.

2. Arrays should not be partitioned.

3. The priorities for Rebuild or Add-drive 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.

The maximum response time option comes available with an OEM append file. With the AV version of firmware, some of the internal settings are also adjusted:

1. Read requests have higher priority.

2. Firmware's Smart Sorting algorithm is disabled to avoid latency that may arise due to sorting. Firmware is also forced to receive ordered read requests, the original receiving order from host.

Appendix

A

Firmware Functionality Specifications

Basic RAID Management:

Specification	Feature
RAID levels	0, 1(0+1), 3, 5, 10, 30, and 50. Levels 10, 30, and 50 are the multi-level RAID with the logical volume implementations; logical volumes consist of logical drives of different RAID levels that are striped together.
Maximum number of logical drives	8 as default, up to128 via OEM configuration utility
Maximum logical drive capacity	64TB – when optimization mode is set to Sequential 16TB – when optimization mode is set to Random
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 - 512MB memory size 112 - 256MB memory size
Configurable stripe size	4KB, 16KB, 32KB, 64KB, 128KB, or 256KB per logical drive ▪ 128KB is the default for earlier firmware with optimization for Sequential I/Os ▪ 32KB is the default for earlier firmware with optimization for Random I/Os
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 multiple arrays configuration
Maximum number of partitions for each logical drive	8 as default, up to 128 via OEM configuration utility
Maximum number of logical drives in a logical volume	8
Maximum number of logical volumes	8

Maximum number of LUNs Mappable	Default is 128, up to 1024 via OEM configuration utility and dependable on available cache 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
Global spare drive	Supported, the spare drive serving all logical drives
Global spare auto-assign	Supported, applies to non-configured drive(s); safeguards the array if a spare has been used in array rebuild and users forget to configure a new drive as a spare.
Co-existing dedicated 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 controller 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. Administrators may find appropriate time to reset controller later.
Auto recovery from logical drive failure (configuration on drives)	Supported. When 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 controller to reaccept the logical drive by switching off the controller, installing the drive back to its original drive slot, and then power on the controller. The logical drive will be restored to the one-drive-failed status.

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). 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 that 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 passed to host. 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 "Iappend" utility.
Drive low-level format	Supported
RAID parity update tracking and recovery	Yes
Host-side Ordered Tag support	Supports write commands with Ordered Tags embedded
Drive identification (flash drive function)	Supported. Force the drive to light on the activity indicator for user to recognize the correct drive.
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	Supported. The logical drive information is recorded on drive media. The logical drives can still be accessed if using different Infortrend RAID controllers/subsystems.

Save/ restore NVRAM to / from disks	Supported. Save all the settings stored in the controller NVRAM to the logical drive members
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.
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)	1. Capacity <64GB: Head=63, Sector=32, Cylinder=? (depends on capacity) 2. 64GB<capacity<128GB:Head=64, Sector=64, Cylinder=? (depends on capacity) 3. 128GB<capacity<256GB: Head=127, Sector=64, Cylinder=? (depends on capacity) 4. 256GB<capacity<512GB: Head=127, Sector=127, Cylinder=? 5. 512GB<capacity<1TB: Head=255, Sector=64, Cylinder=? (depends on capacity) 6. 1TB<capacity: Head=225, Sector=225, Cylinder=? (depends on capacity) 7. For capacity larger than 1TB, please refer to ***Chapter 9***
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 64
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.
Drive I/O timeout	User adjustable
IO channel diagnostics	Supported; please contact your dealer for more details
Maximum Drive Response Time (Guaranteed Latency I/O)	User adjustable from 160 to 640ms. 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.
Drive roaming	Supported; array configuration stored on disk drives; also a copy of system event logs

Caching Operation:

Write-back cache	Supported.		
Write-through cache	Supported.		
Supported memory type	SDRAM memory for enhanced performance Fast Page Memory with Parity for enhanced data security		
Read-ahead operation	Intelligent Dynamic read-ahead operation for sequential data accessing		
Multi-threaded operation	Yes		
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-cache and the synchronized cache operation between partner controllers. Restores write-back mode when later encountering random and small writes.		
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		**Opt. for Sequential I/Os**	**Opt. For Random I/Os**
	RAID0	128	32
	RAID1	128	32
	RAID3	16	4
	RAID5	128	32
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 IO threads (host command queuing)			

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. Add Drive can even be performed in 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. No need to add another enclosure for the extra drives.
Operating system support for RAID expansion	No. No operating system driver required. No software has to be installed for this purpose.

Fibre Channel Support:

Channel mode	All channels configurable to Host or Drive mode, RCC or Drive+RCC, user configurable.
Redundant controller	Redundancy using FC controllers supported.
Host-side loop failure detection	Supported. The LIPs on the host channels will not be displayed to users.
Drive-side loop failure detection	Supported.
Point-to-point topology (N_Port or NL_Port)	Supported.
Arbitrated loop topology (NL_Port)	Supported.
Fabric topology	Supported.
Host redundant loop / dual-loop topology	Supported. (Also requires the host computer Fibre HBA driver support)
Drive side redundant loop load-sharing	Workloads can be automatically balanced between member loops for performance optimization.
Fibre channel ID	User selectable from ID 0 to 125.
Fibre channel CRC	Supported
Point-to-point and FC-AL protocol	User configurable.
WWN table stored in NVRAM	Each WWN number can be assigned with a nick name for ease of identification
Issue LIP command	LIP manually evoked to bring in FC ports on a restored link
Sync. cache channel over Fibre loops	Supported, no extra cabling between two controllers; communications data can be distributed to one or two dedicated channels or over all drive loops.

S.M.A.R.T. Support:

Copy & replace drive	Supported. User can choose to clone a member drive before drive failure.
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

Redundant Controller:

Active-active redundant controller	Supported
Synchronized cache	Supported. Through single or redundant, dedicated synchronizing channels. Synchronized cache over Fibre loops is supported. Synchronized cache can be disabled when using write-through mode in redundant controllers to prevent performance trade-offs.
Write-back cache enabled in redundant controller mode	Yes; with synchronized cache connection between controllers.

Automatic failover	Yes for all PowerPC controllers (user's interaction necessary)
Automatic failback	Yes for all PowerPC controllers (user's interaction necessary)
Fibre channel redundant controller	Supported.
Controller hot-swap	▪ No need to shut down the failed controller before replacing the failed controller. (Customer's design-in hot-swap mechanism necessary) ▪ Support on-line 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. (Customer's design-in hot-swap mechanism necessary)
Redundant controller communication channel	SCSI; RCC Reset signals built-in Fibre channel(s)
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
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); 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 existing controller.
Cache battery backup	Supported. Battery backup solutions for cache memory are supported in both single controller and redundant controller modes.
Load sharing	Supported. Workload can be flexibly divided between different controllers by assigning logical configurations of drives (LDs/LVs) to different controllers.
User configurable channel mode	Supported. Channel modes configurable (SCSI or Fibre) as HOST or DRIVE in both single controller and redundant controller mode.
Require a special firmware for redundant controller?	No.

Data Safety:

Regenerate parity of logical drives	Supported. Can be performed by users to ensure that bad sectors do not cause data loss in the event of drive failure.
Scheduled Media Scan	Revision 3.34 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 and RAIDWatch revision 2.0 will support this functionality.
Bad block auto-reassignment	Supported. Automatic reassignment of bad block
Battery backup for cache memory	Supported. The battery backup solutions provide long-lasting battery support to the 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 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 starts a shutdown sequence (entering an idle state receiving no I/O requests) upon the detection of high-ambient temperature for an extended period of time.

System Security:

Password protection	Supported. All settings require the correct password (if set) to ensure system security.
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.

Environment Management:

SAF-TE/S.E.S. support	Supported. The SAF-TE/S.E.S. modules can be connected to the drive channel, the 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 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.
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
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 module (if available).
Fault-bus support	Provides the simplest implementation for the enclosure management. All fault-bus input/output signals are active-high/active-low user adjustable.
On-board controller voltage monitors	Supported. Monitors the 3.3V, 5V, and 12V 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

User Interface:

RAIDWatch on-board	Out-of-band configuration via LAN. Browser accessible configuration option by installing RAIDWatch to reserved space on drive via ftp.
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 In-band SCSI, In-band Fibre or SNMP (Windows-based GUI).
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.

High Availability:

Custom inquiry serial number	Custom Inquiry Serial Number (for support of multi-pathing software like Veritas, QLogic, etc)

JBOD-Specific:

Format	Restore controller-maintained defect list to default
Reassign blocks	Add entry to the defect list maintained by controller on disk drives
Write-verification	Writes followed by a verify
SMART	Sense data and mode parameters support
Special mode parameters	Error handling page – Enable/Disable retry Caching page – Enable/Disable Read/Write caching SMART enabled parameters Geometry – saved on format command completion

Appendix

B System Functions: Upgrading Firmware

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.

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.

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.*

- *While the firmware is new, the boot record that comes with it may be the same version as the one in the controller. If this is the case, there is no need to upgrade the Boot Record Binary.*

NOTE:

- *Controller serial port COM 2cannot be used to download firmware.*

Upgrading Firmware Using RAIDWatch Manager

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.

Upgrade Both Boot Record and Firmware Binaries

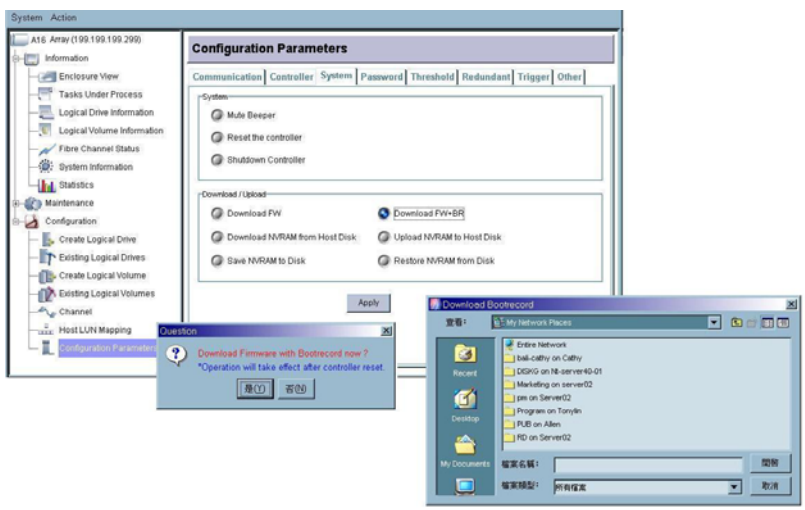

1. Connect to the RAID system locally or from a remote computer using RAIDWatch Manager. While connected to the RAID system, there will be icon(s) with IP address specified on the left of the menu screen. Select by double-clicking the icon of

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.

Upgrade the Firmware Binary Only

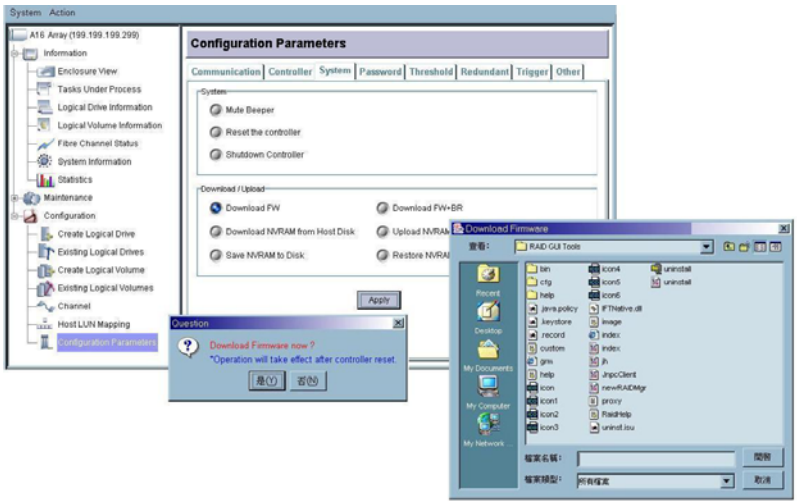

1. Connect to the RAID system locally or from a remote computer using RAIDWatch Manager. While connected to the RAID system, there will be icon(s) with IP address specified on the left of the menu screen. Select by double-clicking the icon of the RAID subsystem which firmware is to be upgraded. Select 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.

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.

Establishing the connection for the RS-232C Terminal Emulation

Please refer to chapter 4, “Connecting to Terminal Emulation,” and also your hardware manual for details on establishing the connection.

Upgrading Both Boot Record and Firmware Binaries

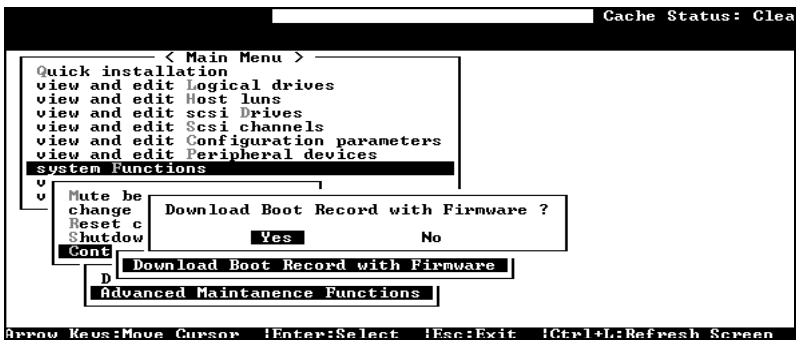

1. From the Main Menu, scroll down to "System Functions."

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.

Upgrading the Firmware Binary Only

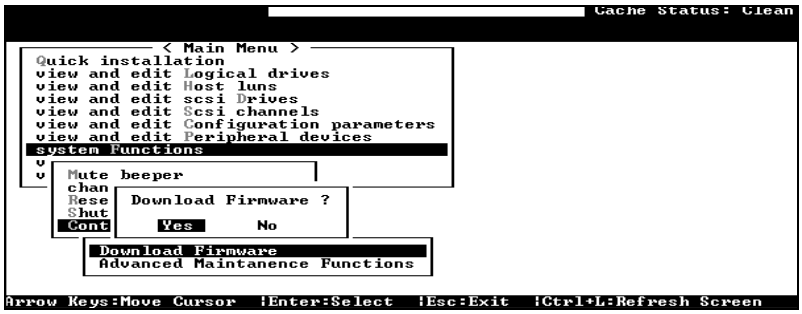

1. From the Main Menu, scroll down to "System Functions."

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 HyperTerminal, choose "Upload" or "Send" (depending on the software).

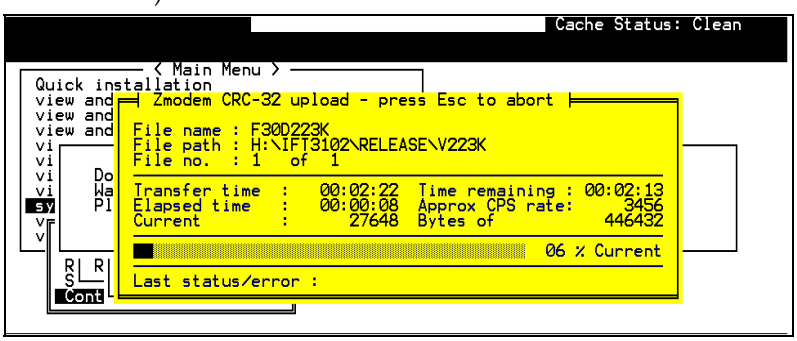

6. When the Firmware completes downloading, the controller will automatically reset itself. If the subsystem/controller is running firmware later than version 3.21, you need to manually reset the subsystem/controller for the new firmware to take effect.

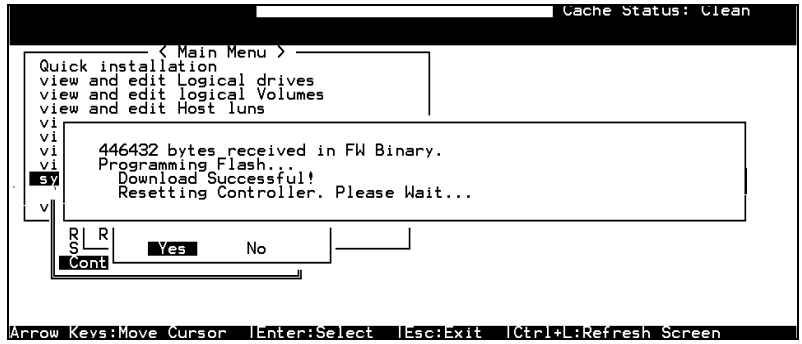
Cache Status: Clean
< Main Menu >
Quick installation
view and edit Logical drives
view and edit logical Volumes
view and edit Host luns
446432 bytes received in FW Binary.
Programming Flash...
Download Successful!
Resetting Controller. Please Wait...
Yes
No
Cont
Arrow Keys:Move Cursor
Enter:Select
Esc:Exit
Ctrl+L:Refresh Screen

Appendix

C

Event Messages

The controller events can be categorized as follows:

Alert	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, it can record up to 1,000 events. To power off or to reset the controller will cause an automatic deletion of all the recorded event logs.

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 LCD screen or terminal.

Event Index

Controller Events

Alert:

Controller SDRAM ECC Multi-bits Error Detected
Controller SDRAM ECC Single-bit Error Detected
Controller SDRAM Parity Error Detected
Controller ALERT: Power Supply Unstable or NVRAM Failed
Controller ALERT: Redundant Controller Failure Detected
Controller BBU Absent or Failed!
Controller BBU Failure Detected!
Controller BBU Thermal Shutdown/Enter Sleep-Mode!
Controller PCI Bus Parity Error Detected
BBU Absent or Failed! Correct It and Reset Ctlr to Take Effect
Force Controller Write-Through on Trigger Cause
ALERT: Controller FAN 0 Low Speed Detected (XXX RPM)
ALERT: +_V Low Voltage Detected(_._V)

Warning:

Force Controller Write-Through on Trigger Cause
Controller BBU Not Fully Charged!
Memory Not Sufficient to Fully Support Current Config.

Notification:

Controller NOTICE: NVRAM Factory Defaults Restored
Controller NOTICE: Redundant Controller Firmware Updated
Memory is Now Sufficient to Fully Support Current Config.
Controller BBU is Charging !
NOTICE: Controller BBU Back On-Line !
NOTICE: Controller BBU Fully Charged !
NOTICE: Controller BBU Present !
NOTICE: Controller FAN _ Back On-Line(_ RPM)
Board_ Cold Temperature Back To Non-Critical Levels(_C))
+_V Upper Voltage Back within Acceptable Limits(_._V)
+_V Lower Voltage Back within Acceptable Limits(_._V)
Controller NOTICE: Redundant Controller Firmware Updated

Channel and Drive Error

Drive:

Alert:

CHL:_ ID:_ Drive ALERT: Aborted Command (B)
CHL:_ ID:_ Drive ALERT: Bad Block Encountered - 0 (10B)
CHL:_ ID:_ Drive ALERT: Block Reassignment Failed - 0 (10B)
CHL:_ ID:_ Drive ALERT: Block Successfully Reassigned - 0 (10B)
CHL:_ ID:_ Drive ALERT: Unexpected Sense Received (B)
CHL:_ ID:_ Drive ALERT: Unexpected Drive Not Ready (__)
CHL:_ ID:_ Drive ALERT: Drive HW Error (__)
CHL:_ ID:_ Target ALERT: Data Overrun/Underrun Detected
CHL:_ ID:_ Target ALERT: Invalid Status/Sense Data Received (10B)
CHL:_ ID:_ Target ALERT: Negotiation Error Detected
CHL:_ ID:_ Target ALERT: Unit Attention Received (10B)
CHL:_ ID:_ ALERT: Redundant Path for Chl:_ ID:_ Failure Detected
CHL:_ ID:_ ALERT: Media Scan Bad Block Unrecoverable-0x0
CHL:_ ID:_ Target ALERT: Gross Phase/Signal Error Detected
CHL:_ ID:_ Target ALERT: Timeout Waiting for I/O to Complete
CHL:_ ID:_ Target ALERT: Unexpected Disconnect Encountered
CHL:_ ID:_ Target ALERT: Unexpected Select Timeout
CHL:_ ID:_ Drive ALERT: Bad Block Encountered - 0 (184) W
CHL:_ ID:_ ALERT: Media Error Recovered-_x__
CHL:_ ID:_ ALERT: Media Error Unrecoverable-_x__
Slot_ IDE Drive ALERT: Bad Block Encountered-__
Slot_ IDE Drive ALERT: Drive HW Error
Slot_ IDE Drive ALERT: Block Successfully Reassigned-__
Slot_ IDE Drive ALERT: Block Reassignment Failed-__
SMART-Slot_ Predictable Failure Detected

SMART-Slot_ Predictable Failure Detected-Clone Failed

Warning:

[] SMART-CH:_ ID:_ Predictable Failure Detected (TEST)
[] SMART-CH:_ ID:_ Predictable Failure Detected
[] SMART-CH:_ ID:_ Predictable Failure Detected – Starting Clone
[] SMART-CH:_ ID:_ Predictable Failure Detected – Clone Failed
CHL:_ ID:_ ALERT: Redundant Path for Chl:_ ID:_ Failure Detected

Notification:

CHL:_ ID:_ Drive NOTICE: Scan Drive Successful
CHL:_ ID:_ NOTICE: Redundant Path for Chl:_ ID:_ Restored
Slot_ NOTICE: Media Scan Bad Block Recovered-__
SMART-Slot_ Predictable Failure Detected-Starting Clone

Channel:

Alert:

CHL:_ FATAL ERROR(_)
CHL:_ Host Channel ALERT: Bus Reset Issued
CHL:_ Host Channel ALERT: Channel Failure
CHL:_ ALERT: Fibre Channel Loop Failure Detected
CHL:_ ALERT: Redundant Loop Connection Error Detected on ID:11
CHL:_ ALERT: Redundant Loop for Chl:_ Failure Detected
CHL:_ ALERT: Redundant Path for Chl:_ ID:_ Expected but Not Found
CHL:_ Drive Channel ALERT: Data Overrun/Underrun Detected
CHL:_ Drive Channel ALERT: Invalid Status/Sense Data Received
CHL:_ Drive Channel ALERT: Negotiation Error Detected
CHL:_ Drive Channel ALERT: Unit Attention Received
CHL:_ RCC Channel ALERT: Data Overrun/Underrun Detected
CHL:_ RCC Channel ALERT: Invalid Status/Sense Data Received
CHL:_ RCC Channel ALERT: Negotiation Error Detected
CHL:_ RCC Channel ALERT: Unit Attention Received
CHL:_ Host Channel ALERT: Parity/CRC Error Detected
CHL:_ Drive Channel ALERT: Gross Phase/Signal Error Detected
CHL:_ Drive Channel ALERT: Timeout Waiting for I/O to Complete
CHL:_ Drive Channel ALERT: Unexpected Disconnect Encountered
CHL:_ Drive Channel ALERT: Unexpected Select Timeout
CHL:_ ID:1 Host Channel ALERT: Parity/CRC Error Detected
CHL:_ RCC Channel ALERT: Gross Phase/Signal Error Detected
CHL:_ RCC Channel ALERT: Parity/CRC Error Detected
CHL:_ RCC Channel ALERT: Timeout Waiting for I/O to Complete
CHL:_ RCC Channel ALERT: Unexpected Disconnect Encountered
CH:_ LIP (__ _) Detected

Warning:

CHL:_ ALERT: Fibre Channel Loop Failure Detected
CHL:_ ALERT: Redundant Loop for Chl:_ Failure Detected
CHL:_ ALERT; Redundant Path for Chl:_ OD:_ Expected but not Found
CHL:_ ID:_ ALERT: Redundant Path for Chl:_ ID:_ Failure Detected

Notification:

CHL:_ Host Channel Notification: Bus Reset Issued
CHL:_ NOTICE: Fibre Channel Loop Connection Restored
CHL:_ LIP(FF B) Detected
CHL:_ ID:_ NOTICE: Redundant Path for Chl:_ ID:_ Restored

Logical Drive Events

Alert:

LG:_ ALERT: CHL:_ ID:_ Media Scan Aborted
LG:_ ALERT: CHL:_ ID:_ Media Scan Failed
LG:_ ALERT: Logical Drive BAD Block Marked 000000084
LG:_ ALERT: Logical Drive Media Error Recovered 000000084
LG:_ Logical Drive ALERT: CHL:_ ID:_ Drive Failure
LG:_ Logical Drive ALERT: CHL:_ ID:_ Drive Missing
LG:_ Logical Drive ALERT: Creation Aborted
LG:_ Logical Drive ALERT: Creation Failed
LG:_ Logical Drive ALERT: Expansion Failed
LG:_ Logical Drive ALERT: Expansion Aborted
LG:_ Logical Drive ALERT: Initialization Failed
LG:_ Logical Drive ALERT: Initialization Aborted
LG:_ Logical Drive ALERT: Logical Drive Bad Block Table BAD
LG:_ Logical Drive ALERT: Logical Drive Bad Block Table FULL
LG:_ Logical Drive ALERT: Logical Drive On-Line Init Table BAD
LG:_ Logical Drive ALERT: Parity Regeneration Failed
LG:_ Logical Drive ALERT: Parity Regeneration Aborted
LG:_ Logical Drive ALERT: Rebuild Aborted
LG:_ Logical Drive ALERT: Rebuild Failed
LG:_ Logical Drive ALERT:CHL:_ ID:_ Clone Failed
LG:_ Logical Drive ALERT:CHL:_ ID:_ Clone Aborted
LG:_ ALERT: Inconsistent Parity Encountered Block 20000000B
ALERT: CHL:_ ID:_ Media Scan Failed
LG:_ Logical Drive ALERT: Logical Drive Block Marked__

Notification:

Creation of Logical Drive _ Completed
LG:_ Logical Drive NOTICE: Add Drive Operation Paused
LG:_ Logical Drive NOTICE: Continue Add Drive Operation

LG:_ Logical Drive NOTICE: Continue Parity Regeneration Operation
LG:_ Logical Drive NOTICE: Continue Rebuild Operation
LG:_ Logical Drive NOTICE: Continue Clone Operation
LG:_ Logical Drive NOTICE: Starting Add Drive Operation
LG:_ Logical Drive NOTICE: Starting Creation
LG:_ Logical Drive NOTICE: Starting Off-Line Expansion
LG:_ Logical Drive NOTICE: Starting Off-Line Initialization
LG:_ Logical Drive NOTICE: Starting On-Line Expansion
LG:_ Logical Drive NOTICE: Starting On-Line Initialization
LG:_ Logical Drive NOTICE: Starting Parity Regeneration
LG:_ Logical Drive NOTICE: Starting Rebuild
LG:_ Logical Drive NOTICE:CHL:_ ID:_ Clone Completed
LG:_ Logical Drive NOTICE:CHL:_ ID:_ Copy and Replace Completed
LG:_ Logical Drive NOTICE:CHL:_ ID:_ Starting Clone
LG:_ NOTICE: CHL:_ ID:_ Starting Media Scan
LG:_ NOTICE: Media Scan of CHL:_ ID:_ Completed
LG:_ NOTIFY: Logical Drive BAD Block Encountered ____
Off-Line Expansion of Logical Drive _ Completed
Off-Line Initialization of Logical Drive _ Completed
On-Line Expansion of Logical Drive _ Completed
On-Line Initialization of Logical Drive _ Completed
Parity Regeneration of Logical Drive _ Completed
Rebuild of Logical Drive _ Completed
Add Drive to Logical Drive_ Completed
NOTICE: Media Scan of CHL:_ ID:_ Completed
NOTICE: CHL:_ ID:_ Starting Media Scan

General Target Events

Alert:

SAF-TE Devices:

SAF-TE Device(_)ALERT: Cooling Fan Failure Detected(_:_)
SAF-TE Device(_)ALERT: Cooling Fan Not Installed(_:_)
SAF-TE Device(_)ALERT: Elevated Temperature Alert
SAF-TE Device(_)ALERT: Power Supply Failure Detected(_:_)
SAF-TE Device(_)ALERT: UPS Power Failure Detected

RAID Controller:

Peripheral Device ALERT: CPU Cold Temperature Detected(0.0(C))
Peripheral Device ALERT: Elevated Temperature Alert
BBU Absent or Failed! Correct It and Reset Ctlr to Take Effect
Controller FAN _ <high/low threshold> Speed Detected (_RPM)
+3.3V <upper/lower threshold> Voltage Detected (_)

Enclosure Devices:

Peripheral Set _ Device ALERT: Cold Temperature _ Detected(_(C))
Peripheral Set _ Device ALERT: FAN _ Failure Detected
Peripheral Set _ Device ALERT: FAN _ Not Present
Peripheral Set _ Device ALERT: Low FAN _ Speed Detected(_ RPM)
Peripheral Set _ Device ALERT: Low Power Supply _ Voltage Detected(_._V)
Peripheral Set _ Device ALERT: Temperature Sensor _ Failure Detected
Peripheral Set _ Device ALERT: Temperature Sensor _ Not Present
Peripheral Set _ Device ALERT: Cold Temperature _ Detected (_(C))
Peripheral Set _ Device ALERT: Power Supply _ Failure Detected
Peripheral Set _ Device ALERT: Power Supply _ Not Present
Peripheral Set _ Device ALERT: UPS _ AC Power Failure Detected
Peripheral Set _ Device ALERT: UPS _ Battery Failure Detected
Peripheral Device ALERT: Elevated Temperature Alert
UPS Connection Is Absent

SES Devices:

SES(C_ I_)Cooling Fan _: Failure Detected !
SES(C_ I_)Power Supply_: Failure Detected !
SES(C_ I_)Power Supply_: DC Failed!
SES(C_ I_)Power Supply_: AC Failed!
SES(C_ I_)Power Supply_: Over Temp Failed!
SES(C_ I_)Power Supply_: DC Over Voltage
SES(C_ I_)Power Supply_: DC Under Voltage
SES(C_ I_)Power Supply_: DC Over Current
SES(C_ I_)Power Supply_: Over Temp Warn
SES(C_ I_)Power Supply_: Power-Off
SES(C_ I_)Power Supply_: Not Turned on
SES(C_ I_)Temp Sensor _: Device Not Supported !
SES(C_ I_)UPS _: Device Not Supported !
SES(C_ I_)UPS_: Battery Failed!
SES(C_ I_)UPS_: UPSFailure Detected!
SES(C_ I_)UPS_: DC Failed!
SES(C_ I_)UPS_: AC Failed!
SES(C_ I_)UPS_: Interface Failed!
SES(C_ I_)UPS_: Batt Predict Fail!
SES(C_ I_)UPS_: Low Battery Warn
SES(C_ I_)UPS_: AC Quality
SES(C_ I_)UPS_: AC High Voltage
SES(C_ I_)UPS_: AC Low Voltage
SES(C_ I_)Voltage_: Under Volt Crit
SES(C_ I_)Voltage_: Over Volt Crit
SES(C_ I_)Voltage_: Under Volt Warn
SES(C_ I_)Voltage_: Over Volt Warn
SES(C_ I_)TEMP_: Under Temp Fail
SES(C_ I_)TEMP_: Over Temp Fail
SES(C_ I_)TEMP_: Under Temp Warn

SES(C_ I_)TEMP_: Over Temp Warn
SES(C_ I_)CURRENT_: OverCureent Crit
SES(C_ I_)CURRENT_: OverCurrent Warn

(**NOTE**: The display of these SES events is partially prepared for OEM enclosure implementations. The SES monitoring events may not all comply if using Infortrend's FC JBODs)

General Peripheral Devices:

ALERT: UPS AC Power-Loss Detected
ALERT: UPS Power Failure Detected
ALERT: UPSBattery Low _%

Warning:

Warning: UPS AC Power-Loss Detected
Warning: UPS Battery Low - _%

Notification:

SAF-TE Devices:

SAF-TE Device(_) NOTICE: Power Supply Back On-Line(_:_)
SAF-TE Device(_) NOTICE: Temperature Back To Non-Critical Levels
SAF-TE Device(_) NOTICE: Fan Back On-Line(_:_)
SAF-TE Device(_) NOTICE: UPS Power Back On-Line

SES Devices:

SES(C_ I_)Cooling Fan _:Fan Back On-Line
SES(C_ I_)Power Supply _:Power Supply Back On-Line
SES(C_ I_)Temp Sensor _:Temperature Back To Non-Critical Levels
SES(C_ I_)UPS _:UPS Power Back On-Line

Controller Self Diagnostics:

Board _ <high/low threshold> Temperature Back To Non-Critical Levels (_ C)

Enclosure Devices:

Peripheral Device NOTICE: Fan Back On-Line
Peripheral Set _ Device NOTICE: FAN_ Back On-Line
Peripheral Set _ Device NOTICE: FAN _ Back On-Line(_ RPM)
Peripheral Set _ Device NOTICE: FAN _ is Present
Peripheral Set _ Device NOTICE: Temperature _ Back To Non-Critical Levels
Peripheral Set _ Device NOTICE: Temperature _ Back To Non-Critical Levels(_C)

Peripheral Set _ Device NOTICE: Temperature _ is Present
Peripheral Set _ Device NOTICE: Power Supply _ Back On-Line
Peripheral Set _ Device NOTICE: Power Supply _ Back On-Line(_._V)
Peripheral Set _ Device NOTICE: Power Supply _ is Present
Peripheral Set _ Device NOTICE: UPS _ AC Power Back On-Line
Peripheral Set _ Device NOTICE: UPS _ Battery Back On-Line
Power Supply _ Back-Online
UPS AC Power Restored
UPS Battery Restored To Safe Level
UPS Connection Detected
NOTICE: FAN Module _ Back On-Line(FAN_ _RPM)

Controller Events

Alert:

Message	Controller SDRAM ECC <multi-bits/single-bit> Error Detected
What Happens?	A DRAM ECC detected error encountered.
What to Do?	Contact your RAID system supplier and replace with new module(s) if necessary.

Message	Controller SDRAM Parity Error Detected
What Happens?	A DRAM parity error encountered.
What to Do?	Contact your RAID system supplier and replace with new module(s) if necessary.

Message	Controller ALERT: Power Supply Unstable or NVRAM Failed
What Happens?	The output voltage drops below preset thresholds or NVRAM component failure.
What to Do?	Contact your RAID system supplier and replace with new module(s) if necessary.

Message	Controller ALERT: Redundant Controller Failure Detected
What Happens?	One of the RAID controllers has failed.
What to Do?	Contact your RAID system supplier for a replacement controller.

Message	CHL:_ FATAL ERROR (_)
What Happens?	One channel has failed.
What to Do?	Check if cable connectors are firmly seated and SCSI buses are properly terminated. With Fibre channels, disconnection may happen on the host side, hub or switch, etc. In redundant mode, the counterpart controller will take over and you may ask your system provider to remove the controller with a failed channel for a repair.

Message	Controller BBU Absent or Failed!
What Happens?	BBU failed or disconnected.
What to Do?	BBU is mounted on controller main board and is not accessed by users. Contact your RAID system supplier for a replacement unit.

Message	Controller BBU Failure Detected
What Happens?	BBU has failed.
What to Do?	Contact your RAID system supplier for a replacement controller.

Message	Controller PCI Bus Parity Error Detected
What Happens?	May be triggered by component failure or conditions like elevated temperature.
What to Do?	Contact your RAID system supplier for an onsite support.

Message	Force Controller Write-Through on Triggered Cause
What Happens?	Critical working conditions, e.g., component failure, forced the subsystem to adopt a safer caching mode.
What to Do?	Restore normal working condition, or contact your RAID system supplier for help.

Warning:

Message	Controller BBU Not Fully Charged!
What Happens?	BBU under charge and is unable to support some configuration change, e.g., changing the caching mode from write-through to write-back.
What to Do?	Wait for the battery charge to complete or replace battery if it can no longer hold the charge.

Message	Controller BBU Thermal Shutdown/Enter Sleep-Mode!
What Happens?	Elevated temperature detected on the charge board. Fast charge halted.
What to Do?	Wait for the battery charge to complete and make sure system cooling is operating normally.

Message	Memory Not Sufficient to Fully Support Current Config.
What Happens?	The installed size of memory does not support current configuration. Try increase memory size.
What to Do?	Contact your RAID system supplier and replace with new module(s) if necessary.

Message	Controller BBU Not Fully Charged!
What Happens?	BBU under charge and is unable to support some configuration change, e.g., changing the caching mode from write-through to write-back.
What to Do?	Wait for the battery charge to complete or replace battery if it can no longer hold the charge.

Notification:

Message	CONTROLLER notice: NVRAM Factory Defaults Restored
What Happens?	Firmware settings have been restored to factory defaults. Options for restoring defaults are not available to users and are only reserved for qualified engineers.
What to Do?	Press ESC to clear the message.

Message	Controller Initialization Completed
What Happens?	Controller initialization completed .
What to Do?	Contact your RAID system supplier and replace with new module(s) if necessary.

Message	Controller NOTICE: Redundant Controller Firmware Updated
What Happens?	RAID controllers have finished shifting I/Os, resetting, and have come online with new version of firmware.
What to Do?	Press ESC to clear the message.

Message	Memory is Now Sufficient to Fully Support Current Config.
What Happens?	Memory size has been expanded.
What to Do?	Press ESC to clear the message.

Message	NVRAM Restore from Disk is Completed
What Happens?	Firmware configuration data previously saved to disk is restored.
What to Do?	Press ESC to clear the message.

Message	NVRAM Restore from File is Completed
What Happens?	Firmware configuration data previously saved as a file is restored.
What to Do?	Press ESC to clear the message.

Message	NOTICE: Controller BBU Back On-Line!
What Happens?	BBU connected or restored..
What to Do?	Press ESC to clear the message.

Message	NOTICE: Controller BBU Fully Charged!
What Happens?	BBU charging completed.
What to Do?	Press ESC to clear the message.

Message	NOTICE: Controller BBU Present!
What Happens?	BBU once absent is now restored.
What to Do?	Press ESC to clear the message.

Message	NOTICE: Controller FAN On-Line(_RPM)

What Happens?	Controller fan back on-line.
What to Do?	Press ESC to clear the message.

Channel/Disk Drive

Disk Drive

Alert:

Message	CHL:_ ID:_ SCSI Target ALERT: Unexpected Select Timeout
What Happens?	Drive SCSI target select timeout. The specified hard drive cannot be selected by the controller. Whether the drive has been removed, or the cabling/termination/canister is out of order.
What to Do?	Check drive-side SCSI cable/termination and drive canister connections.

Message	CHL:_ ID:_ SCSI Target ALERT: Gross Phase/Signal Error Detected
What Happens?	Drive-side SCSI phase/signal abnormality detected.
What to Do?	Press <ESC> to clear the message.

Message	CHL:_ ID:_ SCSI Target ALERT: Unexpected Disconnect Encountered
What Happens?	Drive-side SCSI target unexpected disconnect detected.
What to Do?	Check cabling/termination and canister connections.

Message	CHL:_ ID:_ SCSI Drive ALERT: Negotiation Error Detected
What Happens?	Drive-side SCSI target sync/wide negotiation abnormality detected.
What to Do?	Contact your RAID system supplier.

Message	CHL:_ ID:_ SCSI Target ALERT: Timeout Waiting for I/O to Complete
What Happens?	Drive-side SCSI 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 connections and hard drive.

Message	CHL:_ ID:_ SCSI Target ALERT: SCSI Parity/CRC Error Detected
What Happens?	Drive-side SCSI channel parity or CRC error detected to the specified hard drive.
What to Do?	Check drive-side cable/termination or drive canister connection.

Message	CHL:_ ID:_ SCSI Target ALERT: Data Overrun/Underrun Detected
What Happens?	Drive-side SCSI target data overrun or underrun detected.
What to Do?	Check drive-side cabling/termination/canister connections and hard drive.

Message	CHL:_ ID:_ SCSI Target ALERT: Invalid Status/Sense Data Received (*Sense_key Sense_code*)
What Happens?	Drive-side SCSI invalid status/sense data received from target
What to Do?	Check cabling/termination/canister connections.

Message	CHL:_ ID:_ SCSI Drive ALERT: Drive HW Error (*Sense_key Sense_code*)
What Happens?	Drive-Side SCSI drive unrecoverable hardware error reported
What to Do?	Replace hard drive and the rebuild may begin with a hot-spare or a replacement drive

Message	CHL:_ ID:_ SCSI Drive ALERT: Bad Block Encountered - *Block_number* (*Sense_key Sense_code*)
What Happens?	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:_ ID:_ SCSI Drive ALERT: CHL:_ ID:_ Clone Failed
What Happens?	Drive installed does not respond with "Ready"
What to Do?	Check hard drive and drive-side cabling/termination/canister connections.

Message	Slot _ _ Drive ALERT: Bad Block Encountered - * * * * * * * * *
What Happens?	Occurrences of bad blocks found by Media Scan or during the rebuild process, in a RAID 1/3/5 array. If you see "Bad Block Encountered" but there are no subsequent events, it means data affected by the bad blocks have been reconstructed by the controller and committed to the original drive sector or to other healthy sector.
What to Do?	Press [ESC] to clear this error message.

Message	CHL:_ ID:_ SCSI Target ALERT: Unit Attention Received (*Sense_key Sense_code*)
What Happens?	Drive-side SCSI target unit attention received.
What to Do?	Check hard drive and drive-side cabling/termination/canister connections.

Message	CHL:_ ID:_ SCSI Drive ALERT: Unexpected Sense Received (*Sense_key Sense_code*)
What Happens?	Drive-side SCSI drive unexpected sense data received.
What to Do?	Checking drive-side cabling/termination/drive canister

	connections. This might result from a bad signal quality of poor connection, etc.

Message	CHL:_ ID:_ SCSI Drive ALERT: Block Reassignment Failed - *Block_number* (*Sense_key Sense_code*)
What Happens?	Drive-side block reassignment failed. Drive will be considered failed.
What to Do?	Press [ESC] to clear this error message.

Message	CHL:_ ID:_ SCSI Drive ALERT: Aborted Command (*Sense_key Sense_code*)
What Happens?	SCSI drive aborted command reported
What to Do?	Press [ESC] to clear the error message.

Message	CHL:_ ID:_ ALERT: Media Scan Bad Block Unrecoverable-0x0
What Happens?	Media Scan unable to repair bad blocks.
What to Do?	Press [ESC] to clear the error message.

Warning

Message	SMART-CH:_ ID:_ Predictable Failure Detected (TEST)
What Happens?	This message appears when simulating the SMART detect function. This message shows that your drives support SMART functions.
What to Do?	Press [ESC] to clear the error message.

Message	SMART-CH:_ ID:_ Predictable Failure Detected
What Happens?	SMART-related errors detected. This message will only be displayed when SMART detect is enabled.
What to Do?	Contact your RAID system supplier.

Message	SMART-CH:_ ID:_ Predictable Failure Detected-Starting Clone
What Happens?	SMART errors detected, a spare is conducted to rebuild and to replace the faulty drive. This is a response to the preset scheme.
What to Do?	Contact your RAID system supplier.

Message	SMART-CH:_ ID:_ Predictable Failure Detected-Clone Failed
What Happens?	SMART errors detected and a spare is conducted to rebuild. The cloning process is halted due to power interruption and some other reasons.
What to Do?	Contact your RAID system supplier.

Notification:

Message	CHL:_ ID:_ SCSI Drive ALERT: Block Successfully Reassigned – *Block_number* (*Sense_key Sense_code*)
What Happens?	Bad blocks have been reassigned successfully
What to Do?	Press [ESC] to clear this message.

Message	CHL:_ ID:_ SCSI Drive NOTICE: Scan SCSI Drive Successful
What Happens?	Scanning a new drive from on a SCSI drive successful.
What to Do?	Contact your RAID system supplier.

Channel:

Alert:

Message	CHL:_ ALERT: Redundant Loop Connection Error Detected on ID:_
What Happens?	One of the dual loop members may have failed or been disconnected. Make sure all channels are properly connected and topological configuration properly set.
What to Do?	Check the redundant Fibre channel loop connection is right.

Message	CHL:_ Host Channel ALERT: Channel Failure
What Happens?	Specific host channel may have failed or disconnected.
What to Do?	Check cable connection, Fibre hub, or switch connection. Contact your supplier if no connection errors could be found.

Message	CHL:_ Drive Channel ALERT: Channel Failure
What Happens?	Specific drive channel may have failed or disconnected.
What to Do?	Check cable connection, Fibre hub, or switch connection. Contact your supplier if no connection errors could be found.

Message	CHL:_ ALERT: Fibre Channel Loop Failure Detected
What Happens?	Fibre channel loop failure is detected.
What to Do?	Check cable connection, Fibre hub, or switch connection. Contact your supplier if no connection errors could be found.

Message	CHL:_ ALERT: Redundant loop for Chl:_ Failure Detected
What Happens?	The pair loop of channel:_ has failed.
What to Do?	Check cable connection, Fibre hub, or switch connection. Contact your supplier if no connection errors could be found.

Message	CHL:_ ALERT: Redundant Path for Chl:_ ID:_ Expected but Not Found
What Happens?	Disconnection with the pair loop may have occurred.
What to Do?	Check cable connection, Fibre hub, or switch connection. Contact your supplier if no connection errors could be found.

Message	CHL:_ ID:_ ALERT: Redundant Path for Chl:_ ID:_ Failure Detected
What Happens?	Disconnection with the pair loop may have occurred.
What to Do?	Check cable connection, Fibre hub, or switch connection. Contact your supplier if no connection errors could be found.

Message	CHL:_ Host Channel ALERT: Bus Reset Issued
What Happens?	Bus reset issued on CHL:_ host channel.
What to Do?	Contact your RAID system supplier.

Message	CHL:_ Drive Channel ALERT: Data Overrun/Underrun Detected
What Happens?	Drive-side SCSI target data overrun or underrun detected..
What to Do?	Contact your RAID system supplier.

Message	CHL:_ Drive Channel ALERT: Invalid Status/Sense Data Received
What Happens?	Drive-side SCSI invalid status/sense data received from target.
What to Do?	Contact your RAID system supplier.

Message	CHL:_ Drive Channel ALERT: Negotiation Error Detected
What Happens?	Drive-side SCSI target sync/wide negotiation abnormality detected.
What to Do?	Contact your RAID system supplier.

Message	CHL:_ Drive Channel ALERT: Unit Attention Received
What Happens?	Drive-side SCSI target unit attention received.
What to Do?	Contact your RAID system supplier.

Message	CHL:_ FATAL ERROR (_)
What Happens?	CHL:_ failed.
What to Do?	Contact your RAID system supplier.

Message	CHL:_ RCC Channel ALERT: Data Overrun/Underrun Detected
What Happens?	RCC channel Drive-side SCSI target data overrun or underrun detected.
What to Do?	Contact your RAID system supplier.

Message	CHL:_ RCC Channel ALERT: Invalid Status/Sense Data Received
What Happens?	RCC channel invalid status/sense data received from target.
What to Do?	Contact your RAID system supplier.

Message	CHL:_ RCC Channel ALERT: Negotiation Error Detected
What Happens?	RCC channel target sync/wide negotiation abnormality detected..
What to Do?	Contact your RAID system supplier.

Message	CHL:_ RCC Channel ALERT: Unit Attention Received
What Happens?	RCC channel SCSI target unit attention received.
What to Do?	

Message	CHL:_ Host Channel ALERT: Parity/CRC Error Detected
What Happens?	Host-side parity or CRC error detected to the specified hard drive..

What to Do?	Contact your RAID system supplier.

Message	CHL:_ Drive Channel ALERT: Gross Phase/Signal Error Detected
What Happens?	Drive-side SCSI phase/signal abnormality detected..
What to Do?	Contact your RAID system supplier.

Message	CHL:_ Drive Channel ALERT: Timeout Waiting for I/O to Complete
What Happens?	Drive-side SCSI target I/O timeout. Possible drive-side cabling/termination and canister connection abnormal or drive malfunctioning.
What to Do?	Contact your RAID system supplier.

Message	CHL:_ Drive Channel ALERT: Unexpected Disconnect Encountered
What Happens?	Drive-side SCSI target unexpected disconnect detected.
What to Do?	Contact your RAID system supplier.

Message	CHL:_ Drive Channel ALERT: Unexpected Select Timeout
What Happens?	Drive SCSI 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?	Contact your RAID system supplier.

Message	CHL:_ RCC Channel ALERT: Gross Phase/Signal Error Detected
What Happens?	RCC channel phase/signal abnormality detected..
What to Do?	Contact your RAID system supplier.

Message	CHL:_ RCC Channel ALERT: Parity/CRC Error Detected
What Happens?	RCC channel parity or CRC error detected to the specified hard drive..
What to Do?	Contact your RAID system supplier.

Message	CHL:_ RCC Channel ALERT: Timeout Waiting for I/O to Complete
What Happens?	RCC channel I/O timeout. Possible cabling and canister connection abnormal or backplane malfunctioning.
What to Do?	Contact your RAID system supplier.

Message	CHL:_ RCC Channel ALERT: Unexpected Disconnect Encountered
What Happens?	RCC channel unexpected disconnect detected.
What to Do?	Contact your RAID system supplier.

Notification:

Message	CHL:_ NOTICE: Fibre Channel Loop Connection Restored
What Happens?	Fibre loop connection restored.
What to Do?	Press <ESC> to clear the message.

Message	CHL:_ ID:_ NOTICE: Redundant Path for Chl:_ ID:_ Restored
What Happens?	The connection with pair loop regained.
What to Do?	Press <ESC> to clear the message.

Message	CHL:_ SCSI Drive Channel Notification: SCSI Bus Reset Issued
What Happens?	SCSI bus reset issued.
What to Do?	Contact your RAID system supplier.

Message	CHL:_ Host Channel Notification: SCSI Bus Reset Issued
What Happens?	Host channel bus reset issued.
What to Do?	Contact your RAID system supplier.

Message	CHL:_ LIP(__) Detected
What Happens?	Fibre Loop LIP issued.
What to Do?	Press [ESC] to clear the error message.

Logical Drive Events:

Alert:

Message	LG: _ ALERT: CHL:_ ID:_ Media Scan Aborted
What Happens?	Media Scan has been forced to stop for certain reasons, e.g., critical system events occurred.
What to Do?	Media Scan can be manually started once the normal system condition is restored.

Message	LG: _ ALERT: CHL:_ ID:_ Media Scan Failed
What Happens?	Media Scan has been failed to carry on for certain reasons, e.g., logical drive has fatally failed.
What to Do?	Contact your RAID supplier for further information.

Message	LG:_ Logical Drive ALERT: Logical Drive Block Marked _____
What Happens?	Bad blocks have been located by comparing and recalculating parity. The host computer will receive media error messages if host accesses fall within the affected blocks.
What to Do?	An attempt to recover data from the associated data blocks will be performed.

Message	LG:_ Logical Drive ALERT: Logical Drive Block Recovered _____
What Happens?	Data in bad blocks has been recovered by controller comparing and recalculating parity.
What to Do?	Press ESC to clear the message.

Message	LG:_ Logical Drive ALERT: Logical Drive Block Marked BAD
What Happens?	This message means data in bad blocks can not be recovered by controller comparing and recalculating parity. The host computer will receive media error messages if host accesses fall within the affected blocks.
What to Do?	Contact your RAID system supplier.

Message	LG: Logical Drive ALERT: CHL:_ ID:_ Drive Failure
What Happens?	The specified hard drive in the specified logical drive has failed.
What to Do?	If a spare is available, the controller will automatically start rebuild. If there is no spare, replace the faulty drive and rebuild will be automatically initiated.

Message	LG: Logical Drive ALERT: CHL:_ ID:_ Drive Missing
What Happens?	The specified hard drive in the specified logical drive is missing.
What to Do?	The drive could have been accidentally removed or drive connection problem occurred. Check drive connection and

	related enclosure status.

Message	LG: Logical Drive ALERT: Creation Aborted
What Happens?	System forced to abort logical drive creation process.
What to Do?	Check proper system working conditions.

Message	LG: Logical Drive ALERT: Creation Failed
What Happens?	System forced to abort logical drive creation process. Drive abnormality could have occurred.
What to Do?	Check proper system working conditions and drive connection.

Message	LG: Logical Drive ALERT: Creation Aborted
What Happens?	System forced to abort logical drive creation process.
What to Do?	Check proper system working conditions.

Message	LG:_ Logical Drive ALERT: Expansion Failed
What Happens?	Logical drive expansion failed. It could result from one of the following reasons: 1. The expansion has been canceled by user. 2. The drive used for expansion might have failed during rebuild. 3. Bad blocks are encountered on another member drive during the rebuild. 4. One member drive has failed.
What to Do?	Carefully identify and replace the faulty drive and perform logical drive initialization again.

Message	LG:_ Logical Drive ALERT: Initialization Failed
What Happens?	Logical drive initialization failed. It could result from one of the following reasons: 1. One of the drives used for a logical drive might have failed during rebuild. 2. Bad blocks are encountered on another member drive during the rebuild. 3. Drive hardware faults might have occurred.
What to Do?	Carefully identify and replace the faulty drive and perform logical drive initialization again.

Message	LG: Logical Drive ALERT: Logical Drive Bad Block Table BAD
What Happens?	System failed to generate a bad block table. Logical drive may have serious integrity problem.
What to Do?	Check proper system working conditions.

Message	LG: Logical Drive ALERT: Logical Drive Bad Block Table FULL
What Happens?	Available bad block table entries full.
What to Do?	Check proper system working conditions. Contact system vendor for help.

Message	LG: Logical Drive ALERT: Logical Drive On-Line Init Table BAD

What Happens?	Serious system faults might have occurred and forced the initialization to halt.
What to Do?	Check proper system working conditions. Contact system vendor for help.

Message	LG: Logical Drive ALERT: Parity Regeneration Failed
What Happens?	System faults or integrity problems forced system to abandon the operation.
What to Do?	Check proper system working conditions. Contact system vendor for help.

Message	LG:_ Logical Drive ALERT: Rebuild Aborted
What Happens?	Logical drive rebuild aborted. It could result from one of the following reasons: 1. The rebuild has been canceled by user. 2. The drive used for rebuild might have failed during rebuild. 3. Bad blocks are encountered on another member drive during the rebuild. 4. System faults occurred.
What to Do?	Carefully identify and replace the faulty drive and perform logical drive rebuild. It is best to stop host I/Os temporarily to this logical drive.

Message	LG:_ Logical Drive ALERT: Rebuild Failed
What Happens?	Logical drive rebuild failed. It could result from one of the following reasons: 1. The rebuild has been canceled by user. 2. The drive used for rebuild might have failed during rebuild. 3. Bad blocks are encountered on another member drive during the rebuild. 4. System faults occurred.
What to Do?	Carefully identify and replace the faulty drive and perform logical drive initialization again.

Message	LG:_ Logical Drive ALERT: CHL:_ ID:_ Clone Failed
What Happens?	Drive cloning failed. It could result from one of the following reasons: 1. Cloning has been canceled by user. 2. The drive used for cloning might have failed during rebuild. 3. Bad blocks are encountered on another member drive during the cloning. 4. System faults occurred.
What to Do?	Carefully identify and replace the faulty drive and perform drive cloning again.

Message	LG:_ ALERT: Inconsistent Parity Encountered Block _____
What Happens?	Parity inconsistency encountered on block address _____.
What to Do?	Nay perform regenerate parity function to correct the parity

	faults.

Notification:

Message	Creation of Logical Drive_ Completed
What Happens?	Creation of logical drive completed; the initialization may take place some time later if on-line mode is chosen.
What to Do?	Press <ESC> to clear the message.

Message	LG:_ Logical Drive NOTICE: Add Drive Operation Paused
What Happens?	The add drive expansion process is halted by: 1. Logical drive expansion canceled by user. 2. On of the member drives failed during logical drive initialization. 3. One of the member drive encountered bad block 4. Hardware failure
What to Do?	If the target logical drive has failed, try to rebuild the logical drive.

Message	LG:_ Logical Drive NOTICE: Continue Add Drive Operation
What Happens?	The target logical drive has been restored to its previous status, and the add drive operation may continue.
What to Do?	Press <ESC> to clear the message.

Message	LG:_ Logical Drive NOTICE: Continue Parity Regeneration Operation
What Happens?	The target logical drive has been restored to its previous status, and the regeneration operation can continue.
What to Do?	Press <ESC> to clear the message.

Message	LG:_ Logical Drive NOTICE: Continue Parity Rebuild Operation
What Happens?	The target logical drive has been restored to its previous status, and the rebuild operation can continue.
What to Do?	Press <ESC> to clear the message.

Message	LG_ Logical Drive NOTICE: Starting Creation
What Happens?	Logical drive initial configuration completed, starting the creation process.
What to Do?	Press <ESC> to clear the message.

Message	LG_ Logical Drive NOTICE: Starting Off-Line Expansion
What Happens?	Start expanding the logical drive immediately after selecting to expand.
What to Do?	Press <ESC> to clear the message.

Message	LG_ Logical Drive NOTICE: Starting Off-Line Initialization
What Happens?	Start initializing the logical drive immediately after creating the logical drive.
What to Do?	Press <ESC> to clear the message.

Message	LG_ Logical Drive NOTICE: Starting On-Line Expansion
What Happens?	Start expanding the logical drive when system find appropriate time after selecting to expand.
What to Do?	Press <ESC> to clear the message.

Message	LG_ Logical Drive NOTICE: Starting On-Line Initialization
What Happens?	Start initializing the logical drive immediately when system find appropriate time after creating the logical drive.
What to Do?	Press <ESC> to clear the message.

Message	LG=_ Logical Drive NOTICE: Starting Parity Regeneration
What Happens?	Start regenerating parity of a logical drive.
What to Do?	Press <ESC> to clear the message.

Message	LG_ Logical Drive NOTICE: Starting Rebuild
What Happens?	The rebuild process has begun.
What to Do?	This is the message displayed when a stand-by spare is available or when a faulty drive is replaced. The controller automatically detects a drive for rebuild if the hot swap check time is set.

Message	LG:_ Logical Drive NOTICE: CHL:_ ID:_ Clone Completed
What Happens?	The clone process has been completed.
What to Do?	Press <ESC> to clear the message. When cloning is completed, carefully identify and replace the faulty/replaced drive.

Message	LG_ Logical Drive NOTICE: CHL:_ ID:_ Copy and Replace Completed
What Happens?	Copy and Replace on drive CHL:_ ID:_ completed.
What to Do?	Press <ESC> to clear the message.

Message	LG_ Logical Drive NOTICE: CHL:_ ID:_ Starting Clone
What Happens?	This message is displayed when a member drive is manually cloned to a spare, or that a spare is automatically applied to clone a faulty member according to the preset scheme.
What to Do?	Press <ESC> to clear the message. When cloning is completed, carefully identify and replace the faulty/replaced drive.

Message	LG_ NOTICE: CHL:_ ID:_ Starting Media Scan
What Happens?	Starting Media Scan on a specific hard drive.
What to Do?	Press <ESC> to clear the message.

Message	LG_ NOTICE: Media Scan of CHL:_ ID:_ Completed

What Happens?	Media Scan on a specific hard drive completed.
What to Do?	Press <ESC> to clear the message.

Message	Off-Line Expansion of Logical Drive_ Completed
What Happens?	Logical drive expansion off-line completed.
What to Do?	Press <ESC> to clear the message.

Message	On-Line Expansion of Logical Drive_ Completed
What Happens?	Logical drive expansion on-line completed.
What to Do?	Press <ESC> to clear the message.

Message	Off-Line Initialization of Logical Drive_ Completed
What Happens?	The off-line initialization process of LG_ has been completed.
What to Do?	Press <ESC> to clear the message.

Message	On-Line Initialization of Logical Drive_ Completed
What Happens?	The on-line initialization process of LG_ has been completed.
What to Do?	Press <ESC> to clear the message.

Message	Parity Regeneration of Logical Drive_ Completed
What Happens?	The parity regeneration process on logical drive_ completed.
What to Do?	Press <ESC> to clear the message.

Message	Rebuild of Logical Drive_ Completed
What Happens?	The controller/subsystem has successfully rebuilt a logical drive.
What to Do?	Press <ESC> to clear the message.

Message	Add Drive to Logical Drive_ Completed
What Happens?	The expansion "by adding new drive" is completed.
What to Do?	Press <ESC> to clear the message.

Message	NOTICE: CHL:_ ID:_ Starting Media Scan
What Happens?	Media Scan started on drive CHL:_ ID:_.
What to Do?	Press <ESC> to clear the message.

General Target Events:

Alert:

SAF-TE Devices:

Message	SAF-TE Device (_) ALERT: Cooling Fan Failure Detected (_)
What Happens?	The cooling fan has failed.
What to Do?	Contact your system supplier for a replacement and further diagnosis.

Message	SAF-TE Device (_) ALERT: Cooling Fan Not Installed (_)
What Happens?	The installed fan once installed (_) is missing.
What to Do?	See if the fan has been removed or a general failure has occurred. Contact your system supplier.

Message	SAF-TE Device (_) ALERT: Power Supply Failure Detected (_)
What Happens?	Power supply failure detected by SAF-TE enclosure management.
What to Do?	Check the power supply module, contact your RAID system supplier.

Message	SAF-TE Device (_) ALERT: Elevated Temperature Alert (_)
What Happens?	High temperature detected.
What to Do?	High temperature may lead to malfunctioning and system failure. The most probable cause is the cooling system failure. Contact your system provider immediately.

Message	SAF-TE Device (_) ALERT: UPS Power Failure Detected (_)
What Happens?	UPS Power Failure.
What to Do?	If UPS protection is lost, inconsistency may occur to cached data upon power interruption. Auto-switch to write-through cache upon the detection of UPS failure will be activated if this option has been enabled. Check connection with the UPS and contact your system provider for help.

RAID Controller:

Alert:

Message	Peripheral Device ALERT: CPU Temperature <high/low threshold> Temperature Detected (_._C)
What Happens?	The detected CPU temperature is higher or lower than the preset threshold.
What to Do?	Check the enclosure ventilation condition. If necessary, temperature thresholds can be modified to suit different working conditions.

Message	Peripheral Device ALERT: Temperature <Elevated/Low> Temperature Alert (_._C)
What Happens?	The detected main circuit board temperature is higher or lower than the preset threshold.
What to Do?	Check the enclosure ventilation condition. If necessary, temperature thresholds can be modified to suit different working conditions.

Message	BBU Absent or Failed!Correct It and Reset Ctlr to Take Effect
What Happens?	As the message suggests.
What to Do?	BBU is not a field replaceable module, contact your system provider for a replacement.

Message	Controller FAN_ <high/low threshold> Speed Detected (_RPM)
What Happens?	This event refers to the cooling fan in front panel. Higher or Lower rotation speed detected.
What to Do?	Contact your system vendor for replacing the cooling fan if necessary.

Message	+3.3V <upper/lower threshold> Voltage Detected (_)
What Happens?	The detected +3.3V voltage source is now higher or lower than the preset voltage threshold.
What to Do?	Check power supply condition, voltage threshold settings and contact the your system supplier.

Message	Peripheral Device ALERT: Controller FAN_ Not Present or Failure Detected
What Happens?	This event refers to the cooling fan in front panel. Check cable connection and see if the fan(s) has failed.
What to Do?	Check cable connection and see if the fan(s) is rotating. Some OEM solutions may have removed front panel fans and the "fan detect" signals should be disabled by setting related jumpers. Please refer to your Hardware Manual for more details.

Message	Peripheral Device ALERT: +5V <upper/lower threshold> Voltage Detected (_)

What Happens?	The detected +5V voltage source is now higher or lower than the preset voltage threshold.
What to Do?	Check power supply condition, voltage threshold settings and contact your system supplier.

Message	Peripheral Device ALERT: +12V <upper/lower> Voltage Detected (_)
What Happens?	The detected +12V voltage source is higher or lower than the preset voltage threshold.
What to Do?	Check power supply condition, voltage threshold settings and contact your system supplier.

Enclosure Devices:

Alert:

Message	Peripheral Set_ Device ALERT: Cold Temperature_ Detected (_._C)
What Happens?	The designated temperature sensor has detected temperature lower than preset threshold. This may be caused by mistakes with device target setting or extreme environmental condition.
What to Do?	Check device target setting and improve condition of the installation site.

Message	Peripheral Set_ Device ALERT: FAN_ Failure Detected
What Happens?	Cooling fan failure detected by enclosure management.
What to Do?	Check cooling fan(s) status, and contact your system supplier.

Message	Peripheral Set_ Device ALERT: FAN_ Not Present
What Happens?	The controller failed to detect the presence of the designated cooling fan.
What to Do?	Check enclosure connection, peripheral device setting, fan module, and contact your system supplier.

Message	Peripheral Set_ Device ALERT: <high/low threshold> Fan_ Speed Detected (_RPM)
What Happens?	Enclosure fan higher or lower rotation speed detected.
What to Do?	Contact your system vendor for a replacement fan.

Message	Peripheral Set_ Device ALERT: <high/low threshold> Power Supply_ Voltage Detected (_._V)
What Happens?	Voltage exceeding preset thresholds
What to Do?	Contact your system supplier.

Message	Peripheral Set_ Device ALERT: Temperature Sensor_ Failure Detected

What Happens?	The designated temperature sensor has failed. This may be caused by mistakes with device target setting or a device failure.
What to Do?	Check device set connection and contact your system supplier.

Message	Peripheral Set_ Device ALERT: Power Supply_ Failure Detected
What Happens?	Power supply failure detected.
What to Do?	Contact your system provider for help.

Message	Peripheral Set_ Device ALERT: Power Supply_ Not Present
What Happens?	Could not detect power supply.
What to Do?	Check device set connection, peripheral device setting, and contact your system supplier.

Message	Peripheral Set_ Device ALERT: UPS_ AC Power Failure Detected
What Happens?	UPS power failure.
What to Do?	If UPS protection is lost, inconsistency may occur to cached data upon power interruption. Contact your system provider for help.

Message	Peripheral Device ALERT: UPS_ Battery Failure Detected
What Happens?	UPS battery failure.
What to Do?	If UPS protection is lost, inconsistency may occur to cached data upon power interruption. Contact your system provider for help.

Message	UPS Connection Is Absent
What Happens?	The diagnose line RS-232C connection from UPS to your array is lost.
What to Do?	Check RS-232C cable connection and UPS working status, and contact your system supplier if the fault can not be corrected.

SES Devices:

Alert:

Message	SES (C_I_) Power Supply_: <Vendor descriptor strings/Device Not Supported>!
What Happens?	Unrecognizable device type.
What to Do?	Press <ESC> to clear the message. Contact your RAID system supplier.

Message	SES (C_I_) Power Supply_: <Vendor descriptor strings/Device Not Installed>!
What Happens?	The installed power supply is missing.
What to Do?	Check loop connection and contact your system provider for help.

Message	SES (C_I_) Power Supply_: <Vendor descriptor strings/Device Unknown Status>!
What Happens?	Device reports unknown status strings.
What to Do?	Check loop connection and contact your system provider for help.

Message	SES (C_I_) Cooling Fan_: <Vendor descriptor strings/Device Not Supported>!
What Happens?	Unrecognizable device type
What to Do?	Check loop connection and contact your system provider for help.

Message	SES (C_I_) Cooling Fan_: <Vendor descriptor strings/Device Not Installed>!
What Happens?	The installed device is missing
What to Do?	Check loop connection and contact your system provider for help.

Message	SES (C_I_) Cooling element_: <Vendor descriptor strings/Device Unknown Status>!
What Happens?	Device reports unknown status strings.
What to Do?	Press <ESC> to clear the message. Contact your RAID system supplier.

Message	SES (C_I_) Temperature Sensor_: <Vendor descriptor strings/Device Not Supported>!
What Happens?	Unrecognizable device type.
What to Do?	Press <ESC> to clear the message. Contact your RAID system supplier.

Message	SES (C_I_) Temperature Sensor_: <Vendor descriptor strings/Device Not Installed>!
What Happens?	The installed device is missing.
What to Do?	Press <ESC> to clear the message. Contact your RAID system supplier.

Message	SES (C_I_) Temperature Sensor_: <Vendor descriptor strings/Device Unknown Status>!
What Happens?	Device reports unknown status strings.
What to Do?	Press <ESC> to clear the message. Contact your RAID system supplier.

Message	SES (C_I_) UPS_: <Vendor descriptor strings/Device Not Supported>!
What Happens?	Unrecognizable device type.
What to Do?	Press <ESC> to clear the message.

Message	SES (C_I_) UPS_: <Vendor descriptor strings/Device Not

	Installed>!
What Happens?	The installed device is missing.
What to Do?	Press <ESC> to clear the message. Contact your RAID system supplier.

Message	SES (C_I_) UPS_: <Vendor descriptor strings/Device Unknown Status>!
What Happens?	Device reports unknown status strings.
What to Do?	Press <ESC> to clear the message. Contact your RAID system supplier.

Message	SES (C_I_) Voltage Sensor_: <Vendor descriptor strings/Device Not Supported>!
What Happens?	Unrecognizable device type.
What to Do?	Press <ESC> to clear the message. Contact your RAID system supplier.

Message	SES (C_I_) Voltage Sensor_: <Vendor descriptor strings/Device Not Installed>!
What Happens?	The installed device is missing.
What to Do?	Press <ESC> to clear the message. Contact your RAID system supplier.

Message	SES (C_I_) Voltage Sensor_: <Vendor descriptor strings/Device Unknown Status>!
What Happens?	Device reports unknown status strings.
What to Do?	Press <ESC> to clear the message. Contact your RAID system supplier.

Message	SES (C_I_) Current Sensor_: <Vendor descriptor strings/Device Not Supported>!
What Happens?	Unrecognizable device type.
What to Do?	Press <ESC> to clear the message. Contact your RAID system supplier.

Message	SES (C_I_) Current Sensor_: <Vendor descriptor strings/Device Not Installed>!
What Happens?	The installed device is missing.
What to Do?	Press <ESC> to clear the message. Contact your RAID system supplier.

Message	SES (C_I_) Current Sensor_: <Vendor descriptor strings/Device Unknown Status>!
What Happens?	Device reports unknown status strings.
What to Do?	Press <ESC> to clear the message. Contact your RAID system supplier.

General Peripheral Devices:

Alert:

Message	ALERT: UPS AC Power-Loss Detected
What Happens?	Mains power is interrupted.
What to Do?	Controller/subsystem can be configured to commence an auto cache flush upon the detection of mains power. See Event Triggered operation.

Message	Peripheral Device ALERT: Power Supply Failure Detected
What Happens?	Power supply failure detected
What to Do?	Press <ESC> to clear the message. Contact your RAID system supplier.

Message	Cooling Fan Not Installed
What Happens?	
What to Do?	Press <ESC> to clear the message. Contact your RAID system supplier.

Message	Cooling Fan Failure Detected
What Happens?	Cooling fan failure.
What to Do?	Press <ESC> to clear the message.

Message	Elevated Temperature Alert
What Happens?	Upper temperature threshold violated.
What to Do?	Press <ESC> to clear the message. Contact your RAID system supplier.

SAF-TE Devices:

Notification:

Message	SAF-TE (_) NOTICE: Fan Back On-Line
What Happens?	Failed fan back to on-line state.
What to Do?	Press <ESC> to clear the message.

Message	SAF-TE Device (_) NOTICE: Temperature Back to Non-Critical Levels (_)
What Happens?	Temperature back to non-critical level.
What to Do?	Press <ESC> to clear the message.

Message	SAF-TE Device (_) NOTICE: Power Supply Back On-Line (_)
What Happens?	Power supply restored.
What to Do?	Press <ESC> to clear the message.

Message	SAF-TE Device (_) NOTICE: UPS Power Back On-Line
What Happens?	UPS power restored..
What to Do?	Press <ESC> to clear the message.

Controller Self-Diagnostics:

Notification:

Message	CPU <high/low threshold> Temperature Back to Non-Critical Levels
What Happens?	CPU operating temperature back to non-critical level.
What to Do?	Press <ESC> to clear the message.

Message	Board_ <high/low> Temperature Back To Non-Critical Levels
What Happens?	Board_ temperature back to non-critical level.
What to Do?	Press <ESC> to clear the message.

Message	+3.3V <high/low> Voltage Back within Acceptable Limits
Event Type	☐Critical ☐Warning ☑Notification
What Happens?	+3.3V voltage source back within acceptable limits.
What to Do?	Press <ESC> to clear the message.

Message	+5V <high/low> Voltage Back within Acceptable Limits
What Happens?	+5V voltage source back within acceptable limits.
What to Do?	Press <ESC> to clear the message.

Message	+12V <high/low> Voltage Back within Acceptable Limits
What Happens?	+12V voltage source back within acceptable limits.
What to Do?	Press <ESC> to clear the message.

Message	NOTICE: Controller FAN_ Back On-Line (_RPM)
What Happens?	Controller fan operating status back to normal
What to Do?	Press <ESC> to clear the message.

I^2C Device:

Notification:

Message	NOTICE: Temperature_ Back to Non-Critical Levels
What Happens?	Detected temperature back to non-critical levels.
What to Do?	Press <ESC> to clear the message.

Message	NOTICE: Temperature_ is present
What Happens?	Temperature sensor_ detected.
What to Do?	Press <ESC> to clear the message.

Message	NOTICE: FAN_ Back On-Liine
What Happens?	FAN_ back online.
What to Do?	Press <ESC> to clear the message.

Message	NOTICE: FAN_ is present
What Happens?	FAN_ detected.
What to Do?	Press <ESC> to clear the message.

Message	NOTICE: Power Supply_ Back On-Line
What Happens?	Power supply back online.
What to Do?	Press <ESC> to clear the message.

Message	NOTICE: Power Supply_ is present
What Happens?	Power supply_ detected.
What to Do?	Press <ESC> to clear the message.

Message	Peripheral Device NOTICE: UPS_ AC Power Back On-Line
What Happens?	UPS_ AC Power Back On-Line.
What to Do?	Press <ESC> to clear the message.

Message	Peripheral Device NOTICE: UPS_ Battery Back On-Line
What Happens?	UPS_ battery back online.
What to Do?	Press <ESC> to clear the message.

SES Device:

Alert:

Message	SES (C_I_) Power Supply_: Power Supply Failure Detected
What Happens?	Power supply failure detected.
What to Do?	Check power module status and contact your supplier for a replacement unit.

Message	SES (C_I_) Cooling element_: Cooling Fan Not Installed
What Happens?	Cooling fan missing or not detected.
What to Do?	Check proper fan installation or contact your supplier to replace a failed unit.

Message	SES (C_I_) Cooling element_: Cooling Fan Failure Detected
What Happens?	Contact your system provider for an immediate replacement of fan modules.
What to Do?	Contact your system provider for an immediate replacement of fan modules.

Message	SES (C_I_) Temperature Sensor_: Elevated Temperature Alert
What Happens?	Detected temperature exceeding safety range.
What to Do?	Check cooling fan status. Contact your system provider for an immediate replacement of fan modules.

Message	SES (C_I_) UPS_: UPS Power Failure Detected
What Happens?	UPS power failure detected.
What to Do?	Check UPS status. If power should fail and UPS is not able to sustain power, data loss might occur.

General Peripheral Device:

Alert:

Message	Peripheral Device ALERT: Power Supply Failure Detected
What Happens?	Power Supply Failure Detected.
What to Do?	Check power module status and contact your supplier for a replacement unit.

Message	Cooling Fan Not Installed
What Happens?	Cooling fan missing or not detected.
What to Do?	Check proper fan installation or contact your supplier to replace a failed unit.

Message	Cooling Fan_ Failure Detected
What Happens?	Cooling fan failure detected.
What to Do?	Contact your system provider for an immediate replacement of fan modules.

Message	Elevated Temperature Alert
What Happens?	General overheating warning.
What to Do?	Check cooling fan status and proper installation of dummy plate. Consult your enclosure vendor's document for probable cause.

Message	UPS Power Failure Detected
What Happens?	UPS device failure detected.
What to Do?	Check UPS status. If power should fail and UPS is not able to sustain power, data loss might occur.

Index

- D -

- E -

- F -

- G -

- H -

- I -

- J -

- L -

- M -

- N -

- O -

- P -

- Q -

- R -

- S -

- T -

- U -

- V -

- W -

- Z -